

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Creating, imaging, and exploiting collective excitations of a multicomponent Bose-Einstein condensate

Permalink

<https://escholarship.org/uc/item/6qb1t6b9>

Author

Olf, Ryan

Publication Date

2015

Peer reviewed|Thesis/dissertation

**Creating, imaging, and exploiting collective excitations of a multicomponent
Bose-Einstein condensate**

by

Ryan Ewy Olf

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Dan M. Stamper-Kurn, Chair

Professor Alexander Pines

Associate Professor Hartmut Häffner

Fall 2015

**Creating, imaging, and exploiting collective excitations of a multicomponent
Bose-Einstein condensate**

Copyright 2015
by
Ryan Ewy Olf

Abstract

Creating, imaging, and exploiting collective excitations of a multicomponent Bose-Einstein condensate

by

Ryan Ewy Olf

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Dan M. Stamper-Kurn, Chair

Ultracold atomic gas systems provide a remarkably versatile platform for studying a wide range of physical phenomena, from analogue particle physics and gravity, to the emergence of subtle and profound order in many body and condensed matter systems. In addition, ultracold atomic gas systems can be used to perform a range of precision measurements, from time keeping to variations in the fine structure constant. In this dissertation, I describe our efforts to build a new apparatus capable testing a range of techniques for performing precision measurements in a magnetic storage ring for cold, possibly Bose-condensed, lithium and rubidium atoms. Next, I briefly touch upon our explorations of spin vortices in a ferromagnetic rubidium Bose-Einstein condensate before presenting an exhaustive account of our work using free-particle-like magnon excitations of the ferromagnetic gas to cool it and measure its temperature in a never-before-seen regime of low entropy. Using magnons as a thermometer, we measure temperatures as low as one nanokelvin in gases with an entropy per particle of about one thousandth of the Boltzmann constant, $0.001 k_B$. I conclude by presenting the details of our procedure for calculating the entropy of our coldest, lowest entropy gases in the regime where the local density approximation does not apply.

To my family, for every advantage they have afforded me

Contents

Contents	ii
List of Figures	v
List of Tables	viii
1 Introduction	1
1.1 Purpose of this dissertation	1
1.2 Prologue	1
1.3 Why ultracold AMO?	2
1.4 Outline: An empty room	4
2 Dual-species apparatus	6
2.1 Overview	6
2.2 Dual species oven	6
2.2.1 Pressures and conductance	8
2.2.2 Temperature control	10
2.2.3 Iterating on the design	12
2.3 Dual species Zeeman slower	14
2.3.1 Optimizing the slower length	14
2.3.2 Dual-species slower design	18
2.3.3 Designing the slower windings	21
2.3.4 Winding the slower	24
2.3.5 Testing the slower	25
3 Microfabricating a magnetic ring trap	28
3.1 Electroplating	30
3.2 Planarization	32
3.3 Singularization, Annealing, and Testing	35
3.4 Chip stack bonding	37
3.5 Mounting, thermal testing, and <i>in situ</i> monitoring	39
3.6 Moving on	41

4	Spin vortices	43
5	Thermometry and cooling with magnons	45
5.1	The challenge of low entropy	49
5.2	Cooling at low entropy	51
5.3	First cooling: a benchmark for improvements	53
5.4	Momentum space focusing	55
5.4.1	Experimental procedure	57
5.4.2	Simulating the magnetic focusing lens	58
5.4.3	Calibrating MSF	60
5.4.4	Practical considerations and limits of MSF	62
5.5	Imaging	65
5.5.1	Focusing the imaging system	67
5.5.2	Calibrating the imaging system	67
5.5.3	Magnification	68
5.5.4	Gain and efficiency	69
5.5.5	Polarization, cross-section, and optical pumping	69
5.5.6	Pulse duration and intensity	71
5.5.7	Dealing with high OD	73
5.5.8	Tracking photons	73
5.6	Characterizing the optical trap	76
5.7	Measuring the trap depth	77
5.8	Leveling the ODT	80
5.9	Setting up a state-dependent trap	82
5.10	Noise performance and trap depth limits	83
5.11	Magnetic field gradient control	85
5.12	Purging magnons from the trap (for cycled cooling)	89
5.13	State purification and preparation	91
5.14	Data and data management	93
5.15	Temperature fitting	95
5.15.1	Masking the condensate	96
5.15.2	Selecting a fit	98
5.16	Thermalization	101
5.16.1	Real-space magnon thermalization	106
5.17	State dependent trap	107
5.18	Error analysis	109
5.18.1	Statistical error	109
5.18.2	Systematic error	112
5.19	Temperature extrapolation	113
5.20	Optimizing cycled decoherence cooling	115
5.21	Two types of cooling	115
5.22	Benefits of a square well potential	117

5.23	Thermometry and cooling results	118
6	Calculating the entropy per particle	124
6.1	Numerical objective	124
6.2	Discretization strategy	125
6.3	The Hermite mesh	129
6.4	Solving a simple 1D problem	130
6.5	Higher dimensions	131
6.6	Comparing to the Thomas-Fermi limit in 2D	134
6.7	Simulating our (harmonic) trap	137
6.8	Entropy in our harmonic trap	141
6.9	The trap is not harmonic	143
6.10	Adjustments to the chemical potential	147
7	The future	148
A	Winding the slower	149
B	Chip fabrication proposal	152
	Bibliography	156

List of Figures

1.1	Lab then and now.	4
2.1	The vacuum chamber	7
2.2	Dual-species oven and parameters	8
2.3	Self-latching relay	11
2.4	Original oven design	12
2.5	Schematic of zeeman slower	13
2.6	Capturable fraction in simulated slower	16
2.7	Estimated slow flux in single-species slower	18
2.8	Dual-species slower design with three sections	19
2.9	Dual-species slower trade-offs	20
2.10	Slower winding and target field	22
2.11	Design field and acceleration required	24
2.12	Close-up of the slower windings	25
2.13	Measured slower field	26
2.14	Numerical optimization of slower parameters	26
3.1	Overview of the microfabricated chip design	28
3.2	Photos of the curvature coil electroplating	32
3.3	Planarization scheme	32
3.4	The chip after wet etch	34
3.5	CMP chuck and machine in action	34
3.6	CMP results	36
3.7	Chip resistance pre- and post-anneal.	36
3.8	Resistance curves without and with a short	37
3.9	Chip stack bonding	38
3.10	Our best, fully-assembled ring trap coils.	38
3.11	Mounting and heat sinking the chip stack	39
3.12	Thermal cycling curves for the chips	40
5.1	What does entropy look like?	48
5.2	Effective potentials for various spin excitations	50

5.3	Magnon cooling and thermometry protocols	52
5.4	Preliminary cycled decoherence cooling data	53
5.5	Principle of momentum space focusing	56
5.6	Experimental sequence for momentum space focusing	58
5.7	Simulating momentum space focusing	59
5.8	Calibrating momentum space focusing	61
5.9	The expansion of the condensate	63
5.10	Schematic of optical trap and imaging system	66
5.11	Calibrating the top imaging magnification	69
5.12	Side imaging light polarization and pulse duration	70
5.13	Features of the repump beams can appear on the atoms.	72
5.14	Saturation of the optical depth	75
5.15	A new, more harmonic, trap	76
5.16	Cross section through the center of combined optical/gravitational potential	77
5.17	Fitting optical trap parameters	78
5.18	Contour plots of tilted trap potentials	80
5.19	Calibrating the circularly polarized trap	82
5.20	Heating rate and trap lifetime in the optical dipole trap	84
5.21	Magnetic field gradients and the elevator	86
5.22	Stabilized magnetic fields	88
5.23	Initial scheme for purging magnons	89
5.24	Purging magnons from the trap	90
5.25	Our scheme to remove majority atoms from MSF images	92
5.26	Data structures for data analysis	94
5.27	Masking the condensate automatically	97
5.28	Assortment of fits varying the mask size	98
5.29	The “fit maker” interface	100
5.30	Thermalization of magnons in a magnetic gradient	102
5.31	Thermalization of magnons created at various trap depths	103
5.32	Thermalization of magnons created in a non-degenerate sample before a several second ramp to a final trap depth	104
5.33	Saturation of the thermal gas	105
5.34	Magnons imaged <i>in situ</i>	107
5.35	Temperatures in a state-dependent trap	108
5.36	Temperatures from fits with and without uncertainty weighting	110
5.37	Comparing expected and actual errors	111
5.38	Reduced Chi-square statistics of the fits	112
5.39	Temperatures in the zero-magnon limit	114
5.40	Optimizing cycled decoherence cooling	116
5.41	Combining magnon-assisted evaporation and cycled decoherence cooling	117
5.42	The critical number for magnon condensation	118
5.43	Thermometry and cooling results	119

5.44	Lowest temperatures versus effective trap depth	120
5.45	A closer look at decoherence cooling	123
6.1	The Hermite basis Lagrange functions and mesh points	129
6.2	The energy and excitations in a harmonic trap	132
6.3	Sample kinetic energy matrices in 1D and 2D	133
6.4	The Thomas-Fermi limit with transverse harmonic confinement	135
6.5	Spectrum of the cylindrical Thomas-Fermi condensate in three dimensions	136
6.6	The ground state in our trap	137
6.7	Excitation spectrum with zero longitudinal wave vector in our trap, modeled as harmonic	138
6.8	Excitation spectrum versus longitudinal wave vector in our trap, modeled as harmonic	139
6.9	Entropy at the longitudinal center of the trap at the lowest depth setting	140
6.10	The calculated peak density in our trap	141
6.11	Entropy density, entropy per particle, and linear atom density in a harmonic trap with our trap frequencies	142
6.12	Line profiles of the excitation densities through the center of the real non-harmonic trap	143
6.13	Two approaches to dealing with the lack of confinement in the anharmonic trap	144
6.14	The energy spectra and entropy density in the anharmonic trap	146
A.1	Winding the slower	150

List of Tables

3.1	POLI-500 DC settings	35
5.1	Entropy per particle in some reference systems	46
5.2	Optical trap parameters	79
6.1	Summary of different entropy per particle estimates	145

Acknowledgments

While most of the experiments described herein took place in a literal vacuum, this work was performed far away from any figurative one. I am indebted to so many people for so many reasons that it seems futile to attempt to enumerate them. Nonetheless, I would be remiss if I did not at least try.

First and foremost, I acknowledge my wife, partner, and friend, Kayte, and our loving family—Glenn, Jeanne, Bob, Kathleen, DJ, Will, and Kelly—for their unending support and encouragement, even as the years turned into nearly a decade. While our poodle, Lucy, cats, Linus and Charlie, and step-dog, Alabama, cannot read, write, or speak English, they would want to know meow, *stomp* *stomp* *bite*, woof, meow, Bank of America *low growl*, *long stare*.

I acknowledge our housemates over these years, who have shared both my triumphs and my struggles: Linda Strubbe, David Strubbe, Will Coulter, Andre Wenz, Todd Gingrich, Pat Shaffer, Grant Rotskoff, and Sydney Schreppler.

I acknowledge the considerable support, patience, confidence, and guidance of my adviser Dan Stamper-Kurn, from whom I have learned so much. The members of the Stamper-Kurn group, ever changing, have proven to be first-class colleagues, friends, and mentors, and I could not have made it through graduate school without their varied contributions. I therefore acknowledge Lorraine Sadler, Kevin Moore, Deep Gupta, Mukund Vengalattore, Tom Purdy, Kater Murch, Sabrina Leslie, Jennie Guzman, Dan Brooks, Thierry Botter, Nathan Brahms, Gyu-Boong Jo, Friedhelm Serwane, Vincent Klinkhamer, Claire Thomas, Nicolas Spethmann, Maryrose Barrios, Tom Barter, Jonathan Kohler, Zephy Leung, Justin Gerber, Severin Daiss, and Daniel Tatum.

I am especially indebted to everyone who worked on my experimental apparatus, called “E4” in Stamper-Kurn lab speak. For this, I acknowledge in particular Toni Öttl, Michael Solarz, Enrico Vogt, Jo Daniels, Gabe Dunn, Sean Lourette, Holger Kadau, Eric Copenhagen, Fang Fang, and especially Ed Marti, who many have referred to as my “lab wife” over the preceding many years. His inventiveness, work ethic, and passion for science were indispensable during my graduate career, and his willingness to put up with my quirks and (strange?) sense of humor were rivaled only my “real wife.”

I acknowledge members of the larger physics community at Berkeley, including administrators and support staff—especially Anne Takizawa, Katalin Markus, Eleanor Crump, Stephen Pride-Raffel, and Anthony Vitan—and faculty, especially the members and supporters of the AMO community—Hartmut Häffner, Holger Müller, and Hitoshi Muryama—as well as my committee member Alex Pines.

Finally, I acknowledge the members—undergrad, grad, and postdoc—of the Compass Project, who helped me better understand my place in the Scientific community, and how to make that community more inclusive and equitable.

Chapter 1

Introduction

1.1 Purpose of this dissertation

The modern dissertation serves many purposes. First and foremost, it documents the author's innovations and novel contributions to his or her field, fulfilling the requirements for a Ph.D. My first priority is to accomplish this to the satisfaction of my committee members.

Equally important, though, at least among my experimental colleagues, a thesis documents the process behind the most substantive and summary figures, which are typically published in journal articles. For projects that did not ultimately lead to results, the dissertation serves to mark important lessons learned and sets signposts for future work. Thus, when results or details can be readily found elsewhere, I will generally defer to the most suitable reference. The bulk of this dissertation focuses on process and intermediate results, some of which paved the way to publication and some of which did not.

Finally, a good dissertation documents important toy models and “tricks of the trade”—be they technical or practical—that enabled or facilitated the work being discussed. While not necessarily original work, these tricks and toys constitute a significant part of the productivity advantage a Nth-year student enjoys over a first-year. Thus, primarily with younger practitioners in mind, this dissertation will, at some points, take a pedagogical turn.

1.2 Prologue

I began my journey to this dissertation some years ago in the summer of 2006. I had graduated from Caltech in 2005 without much in the way of plans and moved to Berkeley with my wife Kayte, who would complete her Ph.D. here at Berkeley in 2009. I had enjoyed majoring in physics, but I had never really found my research groove and I wasn't sure that graduate school was for me. I took a year to try new things, and to step away from the insular world of elite academics.

Obviously, I ultimately made the questionable decision of pursuing a Ph.D. in physics, but the experiences of that year have proven to have enduring value. Apropos to this thesis,

I spent a fair amount of time at the [Exploratorium](#), eventually as a volunteer, and my hours there there have helped me to better understand (or at least frame) what makes science so interesting to me.

I had been to the Exploratorium, San Francisco’s “museum of science, art, and human perception,” many times as a kid, but I don’t recall ever feeling as much awe at the wonders of nature as when I returned after graduating from Caltech. Far from diminishing the experience, budding scientific expertise put myriad interactive experiments into perspective. Questions streamed through my mind, and I felt like I could have spent days wandering the museum, pen and paper in tow, weaving threads of curiosity into a satisfying blanket. Even today, it’s not infrequent that I notice something new, or perhaps something old, but in a new light, even with a familiar exhibit. With the right tools, deep science can be found in deceptively simple packages.

At some base level, the brain— mine, at least— does not distinguish between discoveries that are novel and discoveries that are merely new-to-me. Anyone with a frontier— and everyone has a frontier— can be an explorer. At the Exploratorium, as in real life, a handmade musical instrument, a lens, or a slinky can be someone’s Large Hadron Collider. However, even as the reward centers of the brain crave novelty of any stripe, the Ph.D. is not just about a human’s knowledge, it’s about human knowledge.

Keeping eyes on that prize—discovering something truly new—has been a huge challenge for me. Ultimately, it’s been a very rewarding one. In part, the reward has been goals achieved, uncovering some small bits of new knowledge and giving it to the world. But in part, the reward has been due to a surprise turn: an education that has been, like many of the most sought after quantum states, narrow in one sense, but incredibly broad in another.

1.3 Why ultracold AMO?

According to the United Nations, 44% of the world’s population lives within 150 km of the coast [60], far more than would be expected if population were distributed randomly, but far fewer than might be expected if proximity to ocean were paramount to the viability of human life. The coast is generally convenient to have nearby, but not necessarily vital. Especially astonishing to me, though, is that even if one nullifies the practical amenities that the nearby ocean provides—food and transport are widely available even inland these days—the coast still draws us. The coast provides a diversity of ecologies, climates, and topographies that feed our competing needs for variety and integration. Our own San Francisco Bay area is a paragon of the compelling dynamics of life on the edge, at the confluence of mountains, plains, sea, river, and estuary.

Like our University, ultracold atomic physics sits at a multitude of appealing, if less poetic, interfaces:

classical	quantum
wave	particle
statistical	dynamical
many-body	condensed matter
N -dimensions	$N + 1$ -dimensions
non-equilibrium	equilibrium or steady-state
$\text{SO}(3)$	$\text{S}^2 \times \text{U}(1)/\mathbf{Z}_2$
...	...

These junctions avail us access to an unusually wide array of canonical physical models. Moreover, by meeting these models at the edge of applicability, our community of scientists pushes these models to their breaking points and watches them wash up on the shores of their neighbors. Deep physical puzzles like emergence and universality are nowhere more apparent, and tantalizingly deconstructable, than in cold-atom systems, where one can observe an analogue of Hawking radiation [84], a Higgs mode [17], and topologically-driven quantum Hall conductivity [30] in apparatuses that, even to most physicists, are nearly indistinguishable.

These features make systems comprised of cold, trapped atoms an attractive place to look for scientific breakthroughs¹, and there has been ample funding to support ambitious research programs around the world, leading to the sort of consistent growth that makes the academic job market abnormally viable. Beyond academia, the practical side of AMO puts cold atoms and cold atom analogues to work in strange places—clocks, precision sensors, computers, and computer networks—to great success. The practical work of building and running a cold atom experiment relies on a wide range of technical skills, leaving its practitioners with a range of (hopefully) employable engineering skills, in addition to scientific ones.

A scientist’s tools can be both boon and burden: Abraham Maslow was speaking of scientists when he coined the aphorism, “it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail.” [51] We must not be too attached to our favorite models and tools. One of the privileges of pursuing a Ph.D. is the freedom to hammer in a screw a few times, and to deal with the fallout. The diversity of tools and models employed in ultracold AMO guarantees countless such opportunities for young practitioners. Seeing the need for new, better tools takes practice, and it’s something we get to practice a lot.

For the curious, technically minded young physicist, life in AMO can provide quite an education. Here, indeed, simple ingredients come together to yield remarkable depth.




Figure 1.1: Lab then and now. Over several years, the empty lab (left) was filled with equipment and experimental apparatus (right), much of it custom built and designed.

1.4 Outline: An empty room

The story of this dissertation begins with an empty room, Fig. 1.1, left², and an ambitious vision. While the details of our long-term experimental goals were not entirely clear, the immediate technical goals of our apparatus were sufficient to get us running. Building on previous work in our group [23], we wanted to design and build a new apparatus for trapping ultracold, possibly Bose-condensed, rubidium and lithium in a ring-shaped potential. The magnetic potential, and the magnetic environment, would be well controlled so that opposing objectives could be achieved: the ring would encircle a large area (we would shoot for roughly 1 mm diameter, but the exact size would be adjustable *in situ*) and the trapping potential would be level enough to allow a Bose-Einstein condensate to fill it.

The team that Dan assembled to lead the experimental effort initially consisted entirely of two inexperienced young students: myself and Ed Marti. The two of us would work together closely for nearly eight years. A year after Ed and I began work on the experiment, with apparatus design and construction underway, we were joined by a talented postdoc, Toni Öttl. Over several years, the three of us, with assistance from a few others, most notably

¹I estimate that there has been, on average, approximately one new research article based on trapped atoms in Science or Nature each week since I began watching these journals carefully in 2006. This would correspond to nearly 1 out of every 30 articles that these journals publish! Nobel prizes have been awarded to achievements related to trapped atoms at a similar rate. In the last 20 years, 4 years have yielded such prizes, corresponding to 1 in 15 science Nobels (or 1 in 20 if you count Economics as a science).

²A minor detail: the empty room shown is actually the second empty room we encountered. Construction on the experiment began in one empty room before relocating to another.

Enrico Vogt, a diploma student from Germany, built 95% of the apparatus that would be required to achieve our lofty ring trapping goals. Unfortunately, the missing 5% was spread over a number of sub-systems, and ultimately Dan, Ed, and I decided to abandon the dual-species magnetic ring trap effort and to focus our efforts elsewhere. Many of the details of our efforts during this period will be lost to history³, though a few of our more novel and substantive efforts are documented here.

I describe some of our progress towards our lofty dual-species ring-trapping goals in Chs. 2 and 3. In particular, filling the ring would require a lot of rubidium and lithium atoms and to achieve this we designed and built the optimized dual-species atom trapping apparatus that is described in Ch. 2 and the accompanying references. To make precise, well controlled magnetic fields, we turned to microfabrication. Pertinent details of the construction of the actual ring-potential-producing coils can be found in Ch. 3 and references therein. Documentation of some of our other ring-trapping efforts can be found in Ed's dissertation, Ref. [49].

We ultimately did build a working ring trap interferometer, but the trap was formed optically instead of magnetically, it was a 30 micron wide doughnut rather than a 1 mm wide hula hoop, and instead of interfering atoms we interfered phonons, collective acoustic excitations of a BEC. While I'm quite proud of my work on this project, it is well described in our publication [46] and in Ref. Marti2014a, and I won't be covering that work here. Even so, the ring trap work is important to this dissertation for two reasons.

First, on a conceptual level, our work with phonons got us thinking about collective excitations of the BEC in general, and collective excitations were a core theme of our later work on the dispersion of magnons [47], collective spin modes of the spinor BEC, and the use of magnons as a thermometer and coolant [62], the subject of Chs. 5 and 6. On a practical level, in playing with the interaction between our ring BEC and a spherical quadrupole trap, we accidentally made our first interesting magnetic excitation, a spin vortex, the subject of Ch. 4.

³Or, if not lost to history, are buried in obscure grant reports for our primary funding agencies during this period, DTRA and DARPA.

Chapter 2

Dual-species apparatus

2.1 Overview

All of the fundamental physics relating to the experiments reported in this dissertation take place in a relatively small volume—at most a cubic millimeter, for example in the time-of-flight expansion of a barely degenerate gas, but typically much less than that, less than 0.0001 cubic millimeters—but as is quite common in experimental science, the entire apparatus for allowing and observing that physics occupies a much larger space. Most of our apparatus is quite standard, with individual components well described in a variety of places. The basics of laser cooling and trapping are covered in detail in Metcalf and van der Straten [57]. Our dual-species oven design was based on Stan and Ketterle [81] and Stan [80]. For detailed guidance on assembling a BEC system, we like Lewandowski et al. [40]. Our optical trapping procedure was inspired heavily by Lin et al. [43]. Finally, several other theses from our group describe constructing apparatus very similar to our own, either because ours was built in their shadow [26], or because ours was built contemporaneously and benefited from many mutual exchanges of ideas [24].

Finally, my close collaborator in designing, building, characterizing, and using the present apparatus has covered many of its aspects in Marti [49]. Fig. 2.1 shows an overview of the vacuum system employed in this work. In this chapter, I will focus on two areas of the system: the dual-species oven and Zeeman slower. The accompanying optical setup is discussed in Marti [49] and will not be discussed here, except as necessary to understand the Zeeman slower. Details of the apparatus relevant to particular experiments will be covered in their respective chapters.

2.2 Dual species oven

The original design for our oven is based heavily on the work of Stan and Ketterle [81], with minor modifications pertaining to the substitution of rubidium (in our system) for sodium (in his). The nozzle design in our oven has seen several iterations and is quite distinct from




Figure 2.1: The vacuum chamber. Most aspects of its design and construction can be found in the references. In this chapter, we discuss a few of its novel elements, the oven and the Zeeman slower.

that of Stan and Ketterle [81], but I will not be covering that aspect of our oven as it is discussed sufficiently in the thesis of Marti [49].

The role of the dual species oven is to output a bright, roughly collimated atomic beam directed towards the region of the magneto-optical trap. The oven is separated from the main vacuum system by a nozzle that serves to collimate the atomic beam and whose low conductance allows for a large differential pressure. The flux of the atomic beam is generally controlled by varying the temperature, and thus the pressure, on the oven side of the nozzle. In general, turning up the pressure increases atomic flux. Like most cold atom experiments, we operate our nozzle in the molecular flow regime, wherein the mean free path of the atoms is much larger than the diameter of the tubes comprising the nozzle. These basic facts, and the reasons behind them, are discussed in numerous variations in a variety of references. Beyond those sources already mentioned in this chapter, I enjoy the thesis of Higbie [26] and the extremely useful text *Building Scientific Apparatus* [58]. In particular, the contents of the “vacuum technology” section of this later reference (Ch. 3 in my copy) should be familiar to the reader before continuing. O’Hanlon [61] provides more in-depth coverage of a range of topics in vacuum technology.

In designing the parameters of the oven, it is helpful to work backwards from the oven nozzle and the desired atomic beam, and to keep in mind several constraints. First, the nozzle should be the hottest part of the oven to avoid clogging. The coldest part of any contiguous bit of vacuum chamber should be the atomic reservoir. We want gas-phase atoms to condense primarily at the atom reservoir and nowhere else. In this way, the steady-state vapor pressure in each chamber of the oven will be primarily determined by the temperature of the atom reservoir. The nominal temperature limit for Conflat flanges is 450°C . In reality, Conflat flanges can be operated at higher temperatures, however to be safe, we will plan to operate our nozzle at a maximum of 450°C .

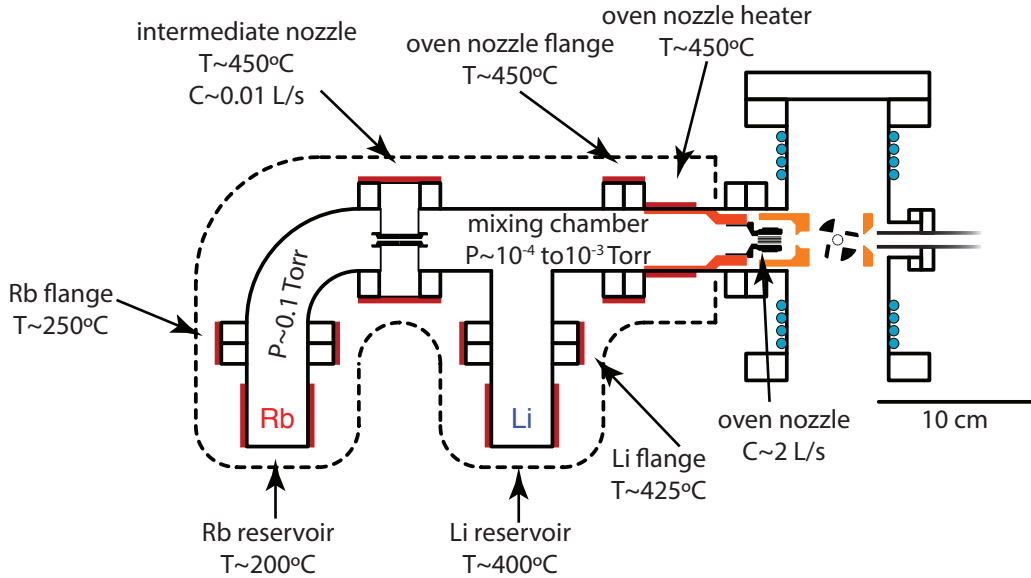


Figure 2.2: The dual-species oven and parameters. Approximate operating pressures, temperatures, and conductances are indicated. Band heaters are shown in dark red. A cross section of the cold surface is shown in orange and is maintained at around -15°C . The nozzle is attached to a copper mount (shown in red) that is press-fit into a steel nipple. The oven nozzle heater is placed around the nipple and heat is conducted to the nozzle through the copper. Water cooling coils, indicated in blue, keep the heat from the oven from heating the rest of the vacuum chamber. Oven insulation is schematically represented by the dotted black line.

2.2.1 Pressures and conductance

As covered in the thesis of Marti [49], we want to operate with a mean free path $\lambda_{\text{mfp}} \gtrsim 1$ cm in the mixing chamber, which is attached directly to the lithium reservoir. This implies partial pressures for both rubidium and lithium in the range of 10^{-4} – 10^{-3} Torr in the mixing chamber. The lithium partial pressure is controlled directly by the lithium reservoir temperature. The equilibrium rubidium partial pressure in the mixing chamber ($P_{\text{Rb}}^{(\text{mix})}$) depends on the conductance of the intermediate (C_{int}) and main (C_{main}) nozzles and the rubidium pressure in its reservoir ($P_{\text{Rb}}^{(\text{res})}$):

$$P_{\text{Rb}}^{(\text{res})} C_{\text{int}} = P_{\text{Rb}}^{(\text{mix})} C_{\text{main}}.$$

Here we have assumed that the pressure in the mixing chamber is negligible compared to the pressure in the rubidium reservoir, and similarly that the pressure outside the oven is negligible compared to the pressure in the mixing chamber.

Vapor pressure curves for the alkali elements can be found in many places. The CRC Handbook [41] has comprehensive data, while the particular properties of ${}^6\text{Li}$ and ${}^{87}\text{Rb}$ are conveniently cataloged online by Gehm [20] and Steck [82], respectively. The optical

properties of ${}^7\text{Li}$ can be found in Ref. [73]. At 400°C , the vapor pressure of lithium is just 2×10^{-4} torr, which handily sets the maximum lithium pressure at which we will operate. We have much more leeway with rubidium as the range of pressures achievable in the reservoir can reach several torr, if we desire.

In order to nail down the intermediate nozzle conductance, we'll consider one additional criterion: that lithium not back flow into the rubidium reservoir. This criterion is likely not critical, but it has some potential advantages. First, the vapor pressure of lithium in the rubidium reservoir will be vanishingly small, so allowing lithium to accumulate there is wasteful. That said, the lithium source should last for a long time, regardless, as five grams of lithium constitutes many more atoms than five grams of rubidium. Second, as we will see shortly, the temperature of the rubidium reservoir will be on the order of the lithium melting point. Lithium and rubidium can alloy and/or react in both solid and liquid forms, so it is possible that lithium in the rubidium reservoir could soak up rubidium, reduce its vapor pressure, and render it inaccessible. For more information on binary mixtures of alkali atoms, refer to, for example, Massalski et al. [52].

In order to avoid backflow, we would like the mean free path of lithium in the intermediate nozzle to be much less than the nozzle length. We accomplish this by using a relatively long intermediate nozzle (3–5 centimeters long) and small diameter (1.2 millimeters) with conductance a few orders of magnitude lower than the main nozzle, $C_{\text{int}} \approx 0.01$ L/s versus $C_{\text{main}} \approx 2$ L/s. The pressure in the rubidium reservoir would then need to be roughly 10^{-1} Torr, and the mean free path would be on the order of 10 microns with the temperature at 200°C . In this regime, the flow is viscous, not molecular, and conductance depends on the pressure and temperature both directly and via the viscosity. Many useful details on calculating mean free path, viscosity, and conductance in a wide range of regimes (including the presence of multiple species with different cross sections) can be found in the text of O'Hanlon [61].

In our case, when the gas mean free path is much less than the diameter of a tube, in the so called viscous regime, the viscosity can be calculated as

$$\eta = \frac{1}{2}nmv\lambda_{\text{mfp}},$$

with v the thermal velocity, m the mass, and n the number density. The conductance of a long tube can then be calculated for a gas in the viscous regime via the Hagen-Poiseuille equation,

$$C = \frac{\pi d^4}{128\eta l} \Delta P,$$

with d the tube inner diameter, l the tube length, and ΔP the pressure difference. The pressure and density in the intermediate nozzle are not determined by the nozzle temperature alone (they depend on the reservoir temperature as well), complicating the precise application of these formulas. However, using the mixing chamber density and nozzle temperature is sufficient to yield a ballpark estimate. Using these equations, for example, we calculate the conductance of the intermediate nozzle of Stan and Ketterle [81] to be within a factor of two

of the measured value. The viscosity of rubidium in our nozzle is expected to be roughly 2×10^{-5} Poise. In practice, we have a lot of latitude in the rubidium temperature and will optimize both rubidium and lithium atom flux by turning all the knobs at our disposal.

2.2.2 Temperature control

Once the nozzles are fixed and the physical design of the oven fixed, maintaining the temperature of the oven becomes a primary concern. We follow Stan and Ketterle [81] and use Mi-Plus band heaters with built-in thermocouples from TEMPCO to heat the oven chamber in multiple places, shown in red in Fig. 2.4. The Mi-Plus heaters can be ordered in a range of peak powers (the peak power is obtained by driving them with 120V AC), sizes, and geometries. Two relevant geometries that we consider are one-piece, two-piece, and expandable. Two-piece or expandable heaters should be used where a one-piece band heater cannot be slid in place. However, we have found that 1.5 inch 2.75 inch one-piece heaters can be stretched and wrapped directly onto 1.5 inch vacuum nipples and 2.75 inch flanges without breaking them, if necessary.

We found elementary order-of-magnitude calculations of heat flow to be sufficient to determine the required power of each heater. For the most part, any underestimate of the required heater power can be fixed with additional insulation. Overestimates can be corrected by reducing the voltage to the heater or by pulse-width-modulation, which in our case is automatically performed by PID controllers that stabilize the temperature as measured at the heaters. The main concern one should have in choosing heaters and insulation is that the requisite thermal gradients can be maintained in the oven.

For example, in our oven design, we may wish for the rubidium reservoir to be maintained at 150–175°C while the intermediate nozzle temperature is fixed at 450°C by the desire to prevent lithium accumulation. Heat will flow from the nozzle through the steel to the rubidium flange and reservoir. Maintaining the appropriate thermal gradient will require a sufficiently powerful heater on the intermediate nozzle and a sufficiently small amount of insulation around the rubidium reservoir and flange. A rough calculation using the parameters of steel and the heat flow equation

$$\dot{Q} = \frac{kA\Delta T}{L},$$

with \dot{Q} the heat flow rate, $k \approx 21$ W/m·K the thermal conductivity, ΔT the temperature difference, and L and A the steel tube length and cross sectional area, respectively, with equal amounts of radiation and convective heat loss from the rubidium flange and reservoir, convinced us that a 300 W heater on the nozzle flange would be more than sufficient. It's not necessarily bad to have too much power. Keep in mind, though, that the fuses and cables that power the heaters need to be sized according to the peak power in order to handle more peak current.

Based on our experience with the heaters and insulation, roughly one half watt per degree celsius desired, as well as good fiberglass insulation, seems to be sufficient for maintaining temperature. We went a little overkill and used 300 W and 400 W heaters for the nozzle

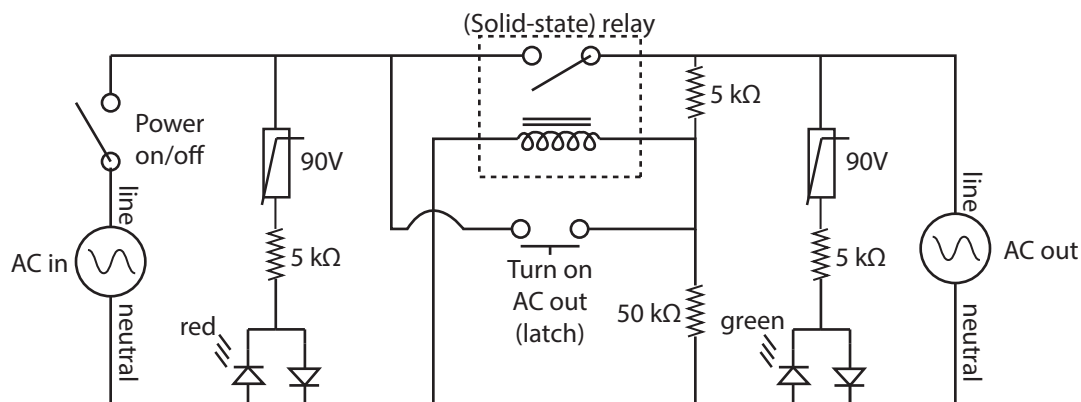


Figure 2.3: Self-latching relay circuit. The solid-state relay used is controlled by an AC voltage. Using a relay that requires DC control would require some sort of rectification.

flanges. By monitoring the temperature of the oven during operation, one can note areas that either are pulled above their setpoint by other heaters, or cannot reach their setpoint. Before swapping heaters, one should try adding and trimming insulation, in general keeping insulation the thinnest where the temperature should be minimal (the bottom of the reservoirs) and thickest where it should be maximal (the nozzle flanges).

One more word of caution: if one uses overpowered heaters, beware that should a heater be left running at full power, parts of the chamber may heat far above the desired temperature. If this happens, one may find the pressure rise to the point where the atom sources are reduced to nothing quickly, necessitating an oven change. Even worse, a vacuum leak may be created that contaminates the entire vacuum system. Excessive thermal gradients across flanges should be avoided (generally by ramping the temperature slowly, at most a few degrees per minute), and our 400 W heaters are more than capable of heating the chamber locally well over the 450°C maximum for Conflat. When our PID feedback is working, and the PID parameters set appropriately, such overheating or rapid heating should never happen. However, we have had the solid state relays that modulate power to the heaters break in the closed position. Additionally, after a total power loss, our PID controllers recall their old (possibly hot) setpoint and may attempt to heat the chamber up quickly once power is restored.

We have adopted a few strategies to avoid such situations. First, the power to the heater runs through percentage controllers that set a maximum duty cycle for the pulse width modulation. This reduces the maximum average power output of the heaters to a safe level while still allowing our temperatures to ramp as desired. Running the heater power through an autotransformer (variac) would be better, reducing both the average and the peak power. Second, the heaters and their PID controllers are powered via a self-latching relay, shown in Fig. 2.3. When the system loses power, the self-latching relay opens and thus when power is restored, the PID controller and heaters will not automatically regain power. The self-latching relay must then be closed by activating a momentary switch, at which point the




Figure 2.4: Original oven design (left) and improved oven design incorporating lessons learned (right). (a) Fewer flanges are used, and the flanges present are oriented horizontally to avoid alkali accumulation. The flanges are farther from the reservoirs, allowing them to be heated more. (b) The intermediate nozzle is welded in place, avoiding the manifold issues associated with the original double-sided flange. (c) The nozzle assembly includes a larger gap between the nozzle heater and the high vacuum chamber, reducing the thermal gradient across the flange that affixes the oven to the oven chamber.

relay will close and remain closed. In practice, the self-latching relay that I built only opens upon power failure when a load is attached, seemingly because of some quirks of the solid state relay. This is fine for turning off power to the PID controller and heaters, but in general is an unwanted bug. Using a spring-loaded mechanical relay in the same circuit would be an obvious improvement.

2.2.3 Iterating on the design

Several drawbacks to the initial oven design were revealed over a few years of operation, which included several catastrophic leaks. While we eventually switched to a less fraught single-species rubidium oven for the bulk of the work described in this dissertation, we designed and built an updated dual-species oven that incorporates several important improvements. A cross sectional schematic of the updated design is shown in Fig. 2.4. The problems we encountered, and proposed solutions incorporated, are discussed below.

First, lithium is well known to react slowly with copper gaskets. Thus, it is commonplace to use nickel gaskets at flanges with substantial lithium exposure. Nickel is not an ideal gasket material, however, as it is much less pliable than copper. This makes nickel gaskets less forgiving of knife-edge imperfections and, in our experience, repeated cycles of thermal expansion and contraction. The updated oven design minimizes the number of flanges and their exposure to lithium. In this design, only one flange is exposed to lithium at a modest pressure, and this flange does not have to transmit any torque. A nickel gasket can safely be used for this flange. Importantly, the flange that affixes the oven to the vacuum system and transmits a torque is “outside” the oven. A normal copper gasket, or better, a silver coated copper gasket (so that it is less likely to stick to the steel when it is time to replace it), can




Figure 2.5: Schematic of zeeman slower. Atoms have a Maxwell-Boltzmann velocity distribution in the oven. The atoms that emerge from the oven are collimated and those with a longitudinal velocity $v_z < v_{\max}$ are slowed and cooled by scattering light in the Zeeman slower. As the atoms travel towards the MOT, their transverse phase space distribution evolves owing to both the passage of time and light scattering. Some fraction of the atoms will have a longitudinal velocity and transverse position such that they can be captured by the MOT.

be used there.

Our original design included a double-sided flange in the construction of the intermediate nozzle. This required using long threaded bolts to affix the adjoining oven parts, and these bolts had to apply pressure to two different knife edges, one of which required a nickel gasket. Long bolts expand more than short ones when thermally cycled, and the requirement that strong force be applied to two knife edges made this junction particularly troublesome. When the nickel gasket originally used at the lithium side of the double-sided flange failed, we replaced it with a softer silver-coated copper gasket. However, the orientation of the flange at the intermediate nozzle allowed liquid lithium to accumulate in a cold spot and the copper gasket was eaten away slowly in spite of its silver coating. The new oven design avoids the double-sided flange. Instead, an intermediate nozzle is welded in place.

Finally, the nozzle assembly incorporates a longer gap between where the nozzle is heated and the rest of the vacuum system. This allows the large thermal gradient between the oven nozzle heater and the vacuum system to be more gradual. As a result, the flange affixing the oven should experience less thermal stress.

2.3 Dual species Zeeman slower

The basics a Zeeman slower are well described in a number of the aforementioned sources [57, 26]. Please refer to one of these resources for an overview of the basic operating principles. In this section, I will only briefly mention the basic elements of a Zeeman slower that are most relevant to the design and construction of our novel dual-species design.

A schematic of a single-species Zeeman slower is shown in Fig. 2.5, which illustrates some of the essential trade offs in Zeeman-slower design. An ideal slower decelerates an atom at some fraction η of the maximum deceleration a_{\max} , which is given for a saturated atom by

$$a_{\max} = \hbar k \Gamma / 2m = v_r \Gamma / 2,$$

for an atom of mass m slowed by a counter-propagating laser of wavenumber k on a transition with linewidth Γ . The recoil velocity is v_r . Increasing the length of the slowing region increases the length over which deceleration can occur, and thus the maximum velocity that will be entrained in the slower. For a slowing length L_z the maximum velocity entrained is

$$v_{\max} = \sqrt{v_f^2 + 2\eta a_{\max} L_z},$$

where v_f is the velocity of the slowed atoms at the exit of the slower. Since the number flux of atoms in the atomic beam rises initially like the velocity cubed, there is generally some advantage to be had by increasing the slower length above zero. On the other hand, as the slower is lengthened, the solid angle subtended by the MOT decreases, reducing the number of atoms in the beam with the appropriate trajectory for capture. Transverse heating of the atomic beam owing to spontaneous emission may also reduce the number atoms with a trajectory suitable for capture.

2.3.1 Optimizing the slower length

For heavy atoms like rubidium, transverse heating can more or less be ignored in slower design, making the resulting calculation of ideal slower length essentially a problem of geometry, as in Higbie [26]. For lighter atoms, transverse heating cannot be ignored, as we will see, and it can be important to account for its effects. We optimize the slower length by considering the slow (capturable) atom flux at the MOT. To calculate the slow flux, we consider the effect of the slower on two main factors: the transverse phase space and the longitudinal velocity distributions. We begin by considering a single-species slower.

To model the transverse phase space, illustrated in Fig. 2.5, we will derive equations for its first and second moments. The first moments, which correspond to the mean transverse velocity, are trivially zero, so we focus on the second moments only. Using the definition of the expected value $\langle \rangle$ and the relationship between the phase space variables (in one dimension), $v_x = \partial_t x$, with ∂_t a short hand for the partial time derivative, we can derive a

simple differential relationship between the second moments,

$$\begin{aligned}\partial_t \langle x^2 \rangle (t) &= 2 \langle xv_x \rangle (t) \\ \partial_t \langle xv_x \rangle (t) &= \langle v_x^2 \rangle (t) + \langle x \partial_t v_x \rangle (t) \\ \partial_t \langle v_x^2 \rangle (t) &= r(t) (\Delta v)^2,\end{aligned}\tag{2.1}$$

where $r(t)$ is the rate of photon scattering and Δv is the root mean square transverse velocity imparted by each scattering event. The term $\langle x \partial_t v_x \rangle$ can be non-zero, for example when the transverse beam profile is inhomogeneous and the scattering rate depends on transverse position. For simplicity we will ignore such effects for now, effectively modeling the cooling beam as translation invariant in the transverse plane. When $r(t)$ is constant, Eqs. (2.1) are easy to solve analytically. One can account for initial atomic beam distributions that have nonzero higher moments (i.e. are not Gaussian) by the appropriate convolution of the initial beam and a narrow Gaussian kernel propagated via Eqs. (2.1); however, for simplicity we will consider beams that are initially Gaussian and can be described by the second moments directly. The estimated capturable flux is sensitive to assumptions about the initial beam and the slowing process, a fact that has some implications for the manner in which we operate the slower.

To solve Eqs. (2.1) for the transverse phase space distribution at the MOT of atoms with a particular longitudinal velocity $v_f < v_z < v_{\max}$, we break the trajectory into three phases: first, the atoms travel ballistically from the oven nozzle into the slower, with $r(t) = 0$, for a time t_1 such that

$$v_z t_1 = L_d + L_z - (v_z^2 - v_f^2)/(2a_{\max}\mu).$$

Here, L_d is the distance between the oven nozzle and start of the slower, and the later terms account for the fact that an atom with initial velocity $v_z < v_{\max}$ will travel in the slower for some distance before it is Zeeman shifted into resonance with the cooling light. Second, once the atoms are brought into resonance with the cooling light, they scatter light at a constant rate as they are slowed, with $r(t) = \eta\Gamma/2$, for a time

$$t_2 = (v_z - v_f)/(v_r r(t)).$$

Finally, the atoms once again travel with $r(t) = 0$ from the end of the slower to the position of the MOT during a time $t_3 = L_m/v_f$. Atoms with $v_z < v_f$ travel the whole length of the slower ballistically. The capturable fraction of atoms with initial longitudinal velocity v_z is then computed by comparing the marginal probability of a slowed atom having a position within the capture range of the MOT to that of an unslowed ballistic atom from the same beam. Assuming a Gaussian atomic beam, integrating the phase space distribution is trivial and the capturable fraction f takes a trivial form:

$$f(v_z) = \left(\text{erf} \left(\frac{\text{MOT size}}{\sqrt{2} \langle x^2 \rangle (v_z)} \right) / \text{erf} \left(\frac{\text{MOT size}}{\sqrt{2} \langle x^2 \rangle_0 (v_z)} \right) \right)^2.\tag{2.2}$$




Figure 2.6: Capturable fraction without transverse cooling (dotted/dashed lines) and with transverse cooling and with variable initial beam sizes (narrow beam: violet, broad beam: red, in reverse rainbow order). Dotted curve neglects transverse heating. Without transverse cooling, the reduction in the capturable fraction is dominated by the reduction of the longitudinal velocity without a commensurate reduction in the transverse velocity. This “blooming” effect can be minimized by focusing the slower beam on the oven aperture, effectively cooling the transverse motion in proportion to the longitudinal motion. The solid rainbow curves have such transverse cooling applied. The presumed distance between the slower and the MOT is 10 cm, each atom is presumed to slow at $0.7a_{\max}$, and the final velocities are $v_f = 30$ m/s for rubidium and $v_f = 90$ m/s for lithium.

Here, “MOT size” refers to the transverse extent over which atoms with longitudinal velocity less than v_f will be entrained in the MOT. Typically, this size is taken to be roughly the size of the beams that form the MOT. An assumption inherent in this formula is that there is a sharp cutoff in the transverse direction between 100% and 0% of atoms with speed v_f being captured. In reality, there is not a sudden cut-off, but we have found this simplifying assumption to yield actionable results nonetheless. More nuanced expressions for the capturable fraction that consider a more detailed model of the MOT can certainly be produced, but considering the level of approximation employed in our treatment thus far, they are unlikely to contribute much additional actionable information.

Fig. 2.6 shows some sample capturable fractions and highlights the importance of some amount of transverse cooling of the atomic beam. With longitudinal cooling only, the beam blooms dramatically because the ratio $\sqrt{\langle v_x^2 \rangle}/v_z$ increases as the slower reduces v_z . As shown by comparing the dotted and dashed curves, even for the lighter lithium, this sort of beam blooming dominates transverse heating in the reduction of capturable flux. Zeeman slowers are typically operated with a slowing laser that is focused near the oven aperture, applying, in the ideal case, transverse and longitudinal cooling in (at least) equal proportion.

The effect of slowing in the presence of such transverse cooling, which we model by assigning our atom beam zero initial transverse velocity, is indicated by the rainbow colored curves

in Fig. 2.6. The different colors correspond to different initial (spatial) beam sizes. Here, the reduction of the capturable flux is due to transverse heating. Under the assumptions of our simple model, the effects of transverse heating exhibit a strong dependence on the model parameters. For beams with initial Gaussian width much larger than the MOT (not physically realistic), transverse heating has little effect in our model, as atoms are equally likely to be redirected toward the MOT as away from it. On the other hand, for very small, focused atomic beams (such as those emerging from a very small aperture), a substantial amount of transverse heating must take place before any atoms are lost. The true physical situation inside our slower lies between these extremes, as the extent of the atomic beam and the cooling laser are both finite. The most physically relevant atomic beam size that works with our model is thus initially slightly smaller than the MOT, giving the “worst case” dark blue curves. Numerical simulations of the slowing process can allow one to implement more elaborate models, allowing, for example, atoms to exit the cooling beam and cease scattering, or to have spatially dependent scattering rates. For estimating the ideal slower length, however, such simulations are almost certainly overkill, as we have found that including more sophisticated physical models beyond what we have already presented changes the ideal lengths by relatively small amounts, even if it does shed better light on the overall atom flux that can be expected.

We are now ready to estimate the impact of the slower on the capturable flux by combining our model slower with a model oven. We take the model oven to operate with a pressure of $P = 2 \times 10^{-4}$ Torr at $T = 400$ K and to have a thin round aperture with 5 mm diameter. The flux emitted into solid angle $\pi\theta^2$ at velocity v_z is then, using the Maxwell-Boltzmann velocity distribution [35],

$$I(v_z) = \pi\theta^2 n A v_z \left(\frac{m}{2\pi k_B T} \right)^{3/2} v_z^2 \exp\left(\frac{-m v_z^2}{2k_B T} \right), \quad (2.3)$$

where A is the aperture area and n is the atom density in the oven at pressure P and temperature T . With $\pi\theta^2$ set to the solid angle subtended by the MOT at the oven aperture, and therefore a function of the slower length, the estimated capturable flux \dot{N} can be calculated as

$$\dot{N} = \int_0^\infty I(v_z) f(v_z) dv_z. \quad (2.4)$$

Note that when the lithium and rubidium partial pressures are equal, as we are assuming, the total lithium flux emerging from the oven far exceeds the rubidium flux emerging from the oven owing to the higher average velocity of the lithium atoms at the same temperature. As noted in Sec. 2.2, the lithium and rubidium partial pressure, and thus the corresponding number fluxes, can be varied independently.

Estimated slow fluxes for a single species slower are shown in Fig. 2.7 for particular design parameters. Notably, we consider a slower with initial dead length of only 10 cm. In general, some amount of dead space, wherein the atoms are not being slowed, is required to allow




Figure 2.7: (Left) Estimated slow flux for single species slower as a function of slower length L_z , with slowing efficiency $\eta = 0.7$, $L_d = 10$ cm, $L_m = 10$ cm, final velocities v_f for rubidium and lithium of 30 m/s and 90 m/s, respectively, and oven parameters described in the text. The dotted curve indicates the slow flux of lithium in a slower optimized for rubidium. (Right) The slow fluxes from the left panel (solid lines) are compared to the slow fluxes that would be estimated without accounting for transverse heating (dashed lines). In our simple model, transverse heating reduces the optimal lithium slower length considerably.

room for vacuum hardware in between the atomic beam nozzle and the Zeeman slower¹. Our initial chamber design did not account for a second differential pumping stage, which increases the dead length. Notably, operating a slower designed for rubidium does increase the capturable lithium flux by some amount, but is far inferior to a dedicated lithium slower and strongly favors rubidium flux over lithium flux². In the next section we talk about improving the lithium performance of a primarily rubidium-focused slower.

2.3.2 Dual-species slower design

Our main innovation in slower design is to take advantage of the large differential in a_{\max} between rubidium and lithium by adding slowing sections that (ideally) impact lithium without impacting rubidium, as shown in the annotated magnetic field profile in Fig. 2.8. Because of more pronounced transverse heating of lithium, and because it can be slowed a great deal in a small amount of space, additional slowing of lithium at the end of the rubidium slower provides the greatest returns in terms of slow lithium flux with the least cost in slow rubidium flux. With the parameters modeled in Fig. 2.9, the model predicts that slow lithium flux jumps 3.5 times by the addition of an optimized 3 cm section at the end of the rubidium slower. The rubidium flux drops by a mere 15% with the same addition.

While the stage III slower provides a generous return, its length is limited by the requirement that rubidium cease its slowing and that lithium be slowed effectively by the stage II

¹Recently, a permanent magnet Zeeman slower has been developed [44] that can be placed in vacuum very close to the atomic beam nozzle, allowing a very short dead space and thus a very short optimal length.

²Recall that owing to its smaller a_{\max} , a slower optimized for lithium cannot slow rubidium.

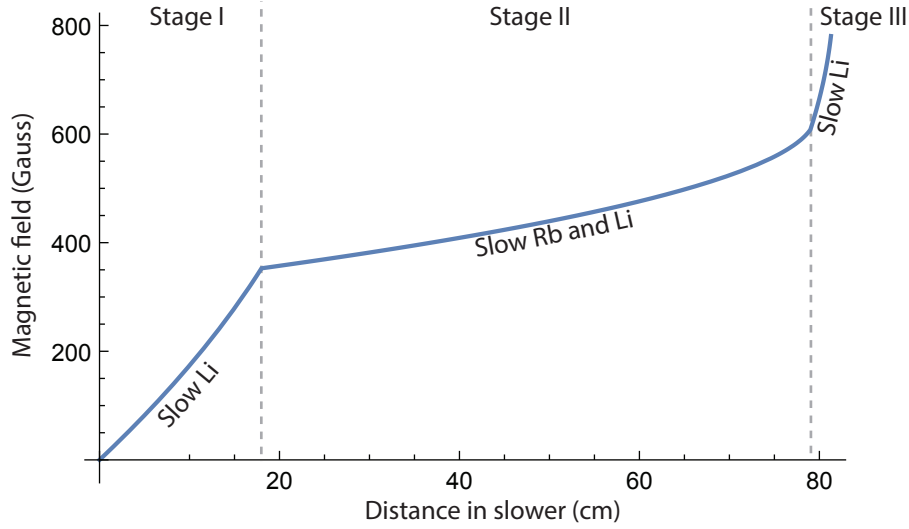


Figure 2.8: Dual-species slower design with three sections. Section I (shown as 18 cm long) slows lithium only. Section II (here 61 cm long) slows rubidium primarily, but lithium comes along for the ride. Section III (merely 2 cm long) slows lithium only. The kink in the field going from section II to III is important. In order for rubidium to exit the slower with the desired final velocity, its velocity must be unable to track the change in the magnetic field at that spot. This places limits on the length of the final lithium slower section, as discussed in the text.

slower. This condition may seem at first glance to be trivially satisfied owing to the fact that $a_{\max}^{(\text{Rb})} < a_{\max}^{(\text{Li})}$, but it is not, because if the velocities of lithium and rubidium are different, the acceleration required for the velocity to follow a particular magnetic field will be different. Rather than thinking about dv/dt we need to consider $dv/dz = (1/v) \times dv/dt$.

The requirement that rubidium not be slowed in the stage III slower can be expressed as

$$\eta a_{\max}^{(\text{Li})}/v_{\text{III}}^{(\text{Li})} > a_{\max}^{(\text{Rb})}/(\eta v_{\text{III}}^{(\text{Rb})}), \quad (2.5)$$

where we have written v_{III} for the velocity of rubidium/lithium at the end of stage II/beginning of stage III. We have used the fact that the stage III magnetic field is designed to slow lithium at a rate $\eta a_{\max}^{(\text{Li})}$ and, for simplicity, we have arbitrarily imposed a requirement that the acceleration required to slow rubidium exceed its maximum value by a factor of $1/\eta$. In deriving Eq. (2.5), note that the magnetic moments of lithium and rubidium in their optically-pumped states are the same. Likewise, the requirement that lithium be slowed in stage II takes the similar form

$$\eta a_{\max}^{(\text{Li})}/v_{\text{III}}^{(\text{Li})} > \eta a_{\max}^{(\text{Rb})}/v_{\text{III}}^{(\text{Rb})}. \quad (2.6)$$

For $0 < \eta < 1$, both Eq. (2.6) and Eq. (2.5) are satisfied if Eq. (2.5) is.

Using the fact that $v_{\text{III}}^{(\text{Rb})} = v_f^{(\text{Rb})}$ and $v_{\text{III}}^{(\text{Li})} = \sqrt{v_f^{(\text{Li})2} + 2\eta a_{\max}^{(\text{Li})} L_{\text{III}}}$, the restriction on the




Figure 2.9: Dual-species slower trade-offs. (Left) Slow flux of rubidium and lithium is plotted as a function of total lithium slower length. Lithium length is applied first to the stage III slower, up to its maximum length, indicated by the vertical dashed gray line. Additional length is allocated to the stage I slower. Rubidium slower (stage II) length is set to maximize rubidium flux at the particular slower length. The stage II length is plotted at the right. Adding lithium sections to the rubidium slower decreases the rubidium flux, but over a wide range of parameters the fractional gain in slow lithium flux far exceeds the cost in rubidium. The lithium performance of the dual-species slower does not approach that of a slower optimized for lithium alone. Parameters used here are $\eta = 0.7$, $L_d = 10$ cm, $L_m = 10$ cm, and final velocities v_f for rubidium and lithium of 30 m/s and 90 m/s, respectively, and oven parameters described in the text.

length of stage III can be written

$$L_{\text{III}} < \frac{\eta^3 a_{\text{max}}^{(\text{Li})} v_f^{(\text{Rb})}}{2a_{\text{max}}^{(\text{Rb})2}} - \frac{v_f^{(\text{Li})2}}{2\eta a_{\text{max}}^{(\text{Li})}}. \quad (2.7)$$

In essence, this condition is meant to guarantee that the kink in the desired magnetic field at the interface of section II and III is sufficiently large. In reality, the slope of the magnetic field cannot be made to increase abruptly. In the physical design of our slower, we will have to make sure it changes fast enough. Like the factors η in the designed rate of deceleration, the margin $1/\eta$ in Eq. (2.5) is meant to help accommodate the differences between an ideal and an actual slower.

Even with a limit on the length of section III, increasing the total length of the lithium slower, by including section I, may be advantageous, as shown in Fig. 2.9. Section I offers an additional benefit that will be apparent in Sec. 2.3.3: some of the coils that produce the field for section I can be made large, such that their field extends into the oven chamber. To some extent, this allows additional slowing of lithium without increasing the total length of the slower; rather, dead space at the beginning of the slower is reduced. In addition, it is advantageous for a rubidium Zeeman slower to operate on top of a large, at least 250 G bias field, at which point the excited state hyperfine structure is in the Paschen-Back regime, in order to reduce the likelihood that an atom is pumped to a dark state [57]. Lithium, which

reaches the Paschen-Back regime at much lower field, might as well be slowed as the bias field is ramped on.

2.3.3 Designing the slower windings

Now that we have a decent model that allows us to weigh the effect of various choices of slower lengths and parameters η , which determine the magnetic field profile $B_{\text{target}}(z)$, we consider what sort of slower we can actually build. Our general approach to this problem is to try certain parameters, producing a target magnetic field profile, and then see if we can expect to build an approximation of that profile that is “good enough” by playing with a model of the magnetic field produced by an achievable configuration of coils and currents.

We found it helpful to break down what constitutes “good enough” into a manageable list of quantitative metrics. The key ingredients are as follows.

Correct v_{max} and v_f The parameters that produced the target magnetic field imply or assume a maximum velocity and final velocity for each rubidium and lithium. Even if the actual field deviates from the target field, we will insist on the difference between its minimum (the magnetic field at which slowing begins for a particular species, which is determined by the slower laser detuning) and maximum (the field at which slowing ceases) maintain the desired velocity limits.

Achievable deceleration The target field is meant to produce uniform deceleration. In reality, it’s fine if the deceleration is not uniform, as long as the rate of deceleration expected by the actual field is achievable given the local laser power. Recall that we are assuming a focused beam, so the intensity is highest at the beginning of the slower. Given an expected field $B(z)$, the condition of achievable deceleration is

$$\frac{\mu}{m} \frac{v_z(z)}{v_r} \frac{dB(z)}{dz} < \frac{s(z)}{1 + s(z)} a_{\text{max}}, \quad (2.8)$$

where μ is the magnetic moment, m the mass, and $s(z)$ the local intensity in units of the saturation intensity. Keep in mind that this equation does not consider the effects of finite (perhaps even power broadened) linewidth. However, we will seek to satisfy it completely and hope that linewidth will help us overcome imperfections in the field produced by imperfections in the actual winding.

Rubidium must stop being slowed at the right point Related to the first point, but worth explicit mention, in order for rubidium to exit the slower with the right velocity, the acceleration required for rubidium to remain in the slower must substantially exceed the acceleration achievable at a particular value of the magnetic field. Thus, we look for violation of Eq. (2.8) for rubidium the appropriate location. Here, the finite linewidth is a potential liability. We can be hopeful, though, as the beam is weakest at the end where rubidium exits the slower. Nonetheless, it will be prudent to look for a large and sustained violation of the achievability condition.




Figure 2.10: Slower winding and target field on same linear position scale. The main section consists of two layers of double-diameter hollow core wire with up to sixteen layers of narrow gauge solid wire on top. The large diameter windings shown at the left are the “stretch” coils that extend the section I field into the oven chamber. The section III field is aided by a high-current “boost” winding, shown in blue to the right. The field peak is sharpened by the high current “anti-gradient” winding, shown in green to the right.

We modeled our physical slower by considering a configuration of closed loops of wire in various positions and with various currents. In reality, we constructed our slower from helical windings; the closed loop approximation should be good as long as the coil spacing is much less than the coil diameter. To assist in finding a configuration of coils that is good enough, it helps to have additional constraints. These constraints should incorporate physical restrictions in the slower winding. To get a sense of the constraints we employed, take a look at the winding pattern we ultimately used in our slower, shown with the target field in Fig. 2.10.

We chose to have five distinct sets of coils. The main section of the slower would comprise two sets of coils: one set composed of 1/8 inch wire and another of 1/16 inch wire. Up to two layers of the 1/8 inch diameter hollow wire would be spaced at multiples of 1/7 inch and covered by up to sixteen layers of 1/16 inch wire, which would be spaced at multiples of 1/14 inch. The 1/8 inch base layer of coils were to be wound on a one inch outer-diameter tube. Each subsequent layer would rest on the layer below, giving the layers 1/8 or 1/16 inch spacing in the radial direction. A key constraint imposed on these windings was that every coil be placed on top of another coil. Gaps in the winding were allowed only on exposed layers. In addition, the average winding pitch was constrained to increase from left to right, as shown in the zoomed in regions of Fig. 2.10. All of the wire in the main section was designed to operate at the same current, but in optimizing the winding pattern the current required was allowed to vary between five and ten amperes. Water can be run through the 1/8 inch hollow layers to prevent some of the heat generated by the slower windings from conducting through to the vacuum chamber.

A large diameter coil was included in the coil optimization routine. The first layer of this “stretch” coil was fixed at an inner diameter of five inches in order to allow the magnetic field to stretch into the oven or differential pumping region. Up to six layers of winding with 1/8 inch wire were allowed in the stretch section. The axial position of the stretch coils were constrained to limit the overlap with the main section. The designed stretch coil current, like the inner diameter, was constrained to be five times the current in the main section, reducing the number of free parameters and increasing the speed of finding the optimized winding pattern.

The final two sets of coils are meant to facilitate the fast rise in magnetic field at the end of the slower. The “boost” coil consisted of up to two layers of the 1/8 inch wire wound on top of the two layers of 1/8 inch wire from the main section. In retrospect, winding the boost directly on the 1 inch tube would have been a better choice, allowing the field slope to change more abruptly and making it easier to stop slowing rubidium. The boost coil current was allowed to reach several hundred amperes. The final coil is the “anti-gradient” coil, which is wrapped on 1.5 inch tube that can fit over the mini-flange that connects the slower vacuum tube to the MOT chamber. The anti-gradient coil consisted of three turns of 1/8 inch wire running up to 200 amperes.

The particular pattern of winding for the main and stretch coils, along with the currents of all of the coils, was optimized using an iterative algorithm that involved adding and removing turns of simulated wire in various places, subject to constraints already mentioned, and varying the currents in order to minimize a target function. The target function weighed five criteria:

1. The squared difference between the target field and calculated field,
2. The condition that the acceleration required, Eq. (2.8), was achievable as necessary in the three stages,
3. The condition that rubidium cease accelerating at the appropriate place,
4. The condition that the gradient field at the MOT be minimal, and
5. The condition that the field reaches the right maximum value so that the final velocity of lithium is correct.

Condition 4, in retrospect, was not important as the gradient due to the slower can be made to add constructively with the gradient of the MOT coil, and does not hurt the operation of the MOT. We would not include it if we were to perform the optimization again. The bias field of the slower is canceled by compensation coils.

Each of the terms in the target function was weighted empirically to assist the algorithm in finding a suitable local minimum. Once a winding and current set were found, tweaks were made by hand to further optimize the field (e.g. increase the likelihood of rubidium ceasing deceleration at the expense of larger gradient field at the MOT) and to make the winding pattern easier to produce (e.g. make the spacing more even so that fewer changes




Figure 2.11: Design field (left) and acceleration required (right). (Left) The target field is shown in solid blue, and the calculated field from the coil winding is shown in dashed orange. (Right) The acceleration required is expressed in terms of a_{\max} for each species. Divisions between the sections, based on the target values of the magnetic field, rather than the design lengths L , are shown as vertical dashed lines. The maximum acceleration achievable given the laser power, assuming a minimum intensity $I/I_{\text{sat}} = 2.5$ 10 cm past the end of the slower, at the presumed position of the MOT, is shown in solid dark gray. As desired, the acceleration required of rubidium exceeds the maximum possible under the simple model that does not account for finite linewidth or transverse variation in the laser power.

in spacing would be required). Throughout the optimization process, we kept tabs on a plot of the acceleration required and a comparison of the actual and target field, as shown in Fig. 2.11. The target field parameters ultimately employed included a 5 cm transition region between stage I and stage II in which the target field did not change. Additionally, the design parameter η varied from section to section. The parameters were $L_I = 18$ cm, $L_{II} = 61$ cm, $L_{III} = 2.6$ cm, $\eta_I = 0.78$, $\eta_{II} = 0.67$, and $\eta_{III} = 0.68$. Lithium and rubidium final velocities v_f were 40 m/s and 30 m/s, respectively. We later realized that higher final velocities up to 100 m/s for lithium could have been used.

2.3.4 Winding the slower

The ambitious nature of the dual-species slower requires precise winding. The main innovation that we employed to reach the desired precision was the use of the threading functionality of a lathe in laying each turn of wire. In the process of winding the main and boost sections of the slower, which are wrapped on a single tube, the threading gears were always engaged and the turns of the chuck (performed by hand rather than by motor) were carefully counted. In this manner, each layer of wire could be absolutely referenced to a particular “home” position of the compound tool holder.

The use of the lathe’s threading facility was inspired by the fortuitous fact that both seven




Figure 2.12: An overview of the first half of the wound slower. One-quarter, one-half, three-quarter, and full density turns are clear.

and fourteen threads-per-inch are standard pitches for 1-1/4" and 7/16" bolts, respectively, and are thus among the limited options built-in to larger imperial lathes. The standard square cross section magnet wire used in our group (hollow conductor kapton-wrapped magnet wire from S&W Wire) happens to be about 1/8" wide. In considering ways of making more precise magnetic fields, I noted the compatibility of 1/8" wire with a seven-per-inch pitch and, furthermore, was able to find square magnet wire with approximately 1/16" width (14 AWG square conductor wire with polyimide enamel from MWS Wire Industries), compatible with 14-per-inch pitch. Furthermore, half density windings of the thinner wire could be achieved. Using the thinner wire allows more layers with different numbers and pitches of turns in the same space, and allows finer control of the turn density. The procedure for winding the slower is detailed in Appendix A. A snapshot of the precise winding is shown in Fig. 2.12

2.3.5 Testing the slower

Before installing the slower, it is very helpful to run a few tests. Resistance measurements between unconnected conductors can reveal shorts (oops!) and four-point resistance measurements of the various contiguous layers will serve as a baseline for future reference. Drops in resistance over time could indicate shorts between layers. It's also important to make an accurate measurement of the magnetic fields generated by the various coils, both as a test for shorts between layers and also to aid in simulating the function of the slower, if necessary.

We used F.W. Bell transverse and axial hall probes to measure the magnetic field at centimeter intervals along the center of the coils. As expected, the field is mostly axial. The magnetic field measurements were performed in a differential manner with DC current switched on and off at each position. The main set of coils, including the 1/8" base layer, is composed of seven different pieces of conductor. The field due to each conductor was measured separately to allow for the possibility of optimizing the operation of the slower by running slightly different currents through the different conductors that comprise the main section. The boost, stretch, and anti-gradient coils were measured separately as well.




Figure 2.13: Measured slower field. The agreement with the expected field is very good, but the measured field curvature at the transition from stage II to stage III is less than we had hoped for. Details on the performance of our slower, and suggestions for further improvements, can be found in our publication, Marti et al. [48].

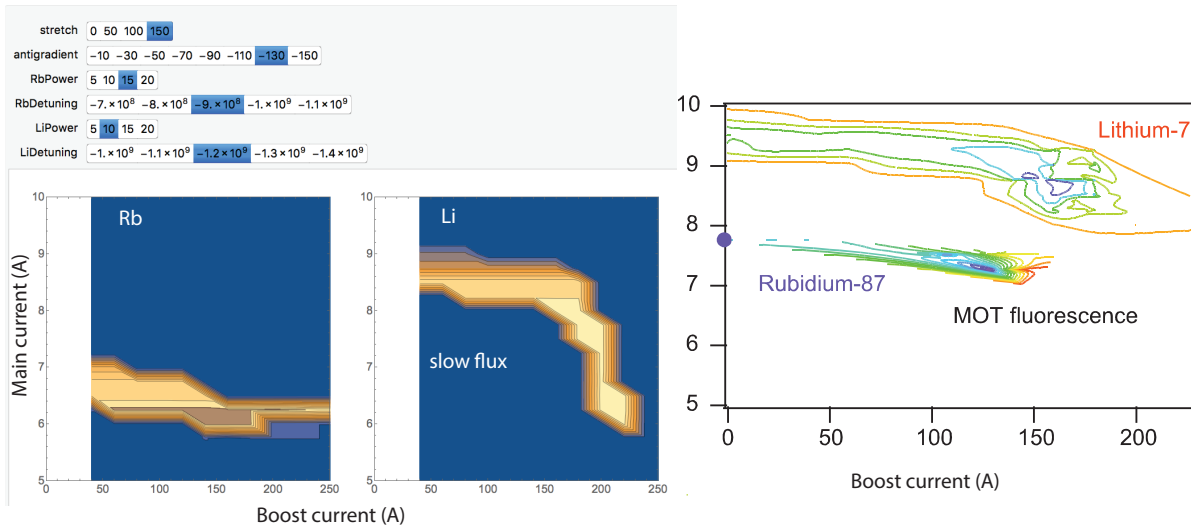


Figure 2.14: A screenshot of our numerical optimization using the simulated slower and fields (left) allowed us to vary multiple parameters and estimate the effect on the capturable flux. Data from our actual optimization (right) showed similar trends, but aligning slow lithium and rubidium flux with a large boost current proved more difficult in reality than in theory. The large purple dot in the right-hand plot represents large rubidium flux with zero boost current. An island of high flux with high boost is seen around 125 amps. Parameters aside from the boost and main currents are not the same in the left and right figures.

With the field well measured, the performance of the slower can be simulated numerically in great detail—accounting for beam inhomogeneity and finite linewidth, for example—allowing one to vary a wide range of parameters to assist in finding suitable settings. We found that including the magnetic fields produced by the MOT in the simulations of the slower improved their ability to predict our observed loading rates. Ultimately, the proof of the slower is in its performance in loading a real MOT. The measured fields of our slower, its real-world performance, and suggestions for future improvements based on lessons learned since its original design are detailed in our publication, Marti et al. [48]. A screenshot of some of our numerical optimizations, in which many simulated parameters are varied, and MOT fluorescence data tracking a similar variation in parameters, is shown in Fig. 2.14.

Our goal of producing a stage III that slows lithium but not rubidium was only partly achieved. Operating in dual species mode as intended, we were unable to find settings that produced a large rubidium flux with exit velocities low enough to be captured by the MOT, owing to the limited field curvature at the transition from stage II to III. New slower designs based on this concept would do well to place more emphasis on increasing the curvature in this region, perhaps with smaller boost coils or by placing the boost coils further from the main coils, as suggested in our publication.

The slower works very well as a single-species slower—no stretch, boost, or anti-gradient currents required—as very few compromises were made in its rubidium performance in the effort to accommodate lithium, and the smooth field and aggressive deceleration payoff in the rubidium flux. Other than producing the data for Marti et al. [48], the slower has been operated in rubidium-only mode for all of the experiments described in this work. In this mode, only the main sections of the slower are powered. The stretch, boost, and anti-gradient coils remain, but are idle.

Chapter 3

Microfabricating a magnetic ring trap

The initial goal for our team was to trap and condense atoms in a magnetic ring trap, building upon work that had been done along these lines in our group on a different apparatus [23]. In contrast to that work, which employed hand-wound millimeter-scale coils with sufficient magnetic field inhomogeneity to preclude filling a ring with a BEC, our experiment would build upon the micro-fabrication expertise developed by members of our group's atom chip experiment [70] to produce coils capable of trapping atoms in an extremely flat potential with a circular minimum.




Figure 3.1: Overview of the microfabricated chip design, including labeled crosssection of the assembled stack of four chips, magnetic field profiles, and schematic of the overall assembly.

The basic requirements and design of the microfabricated chips were worked out by our postdoc Anton Öttl, and are shown in Fig. 3.1. Prior to my assuming a role in the fabrication process, two members of our research group, Tom Purdy and Dan Brooks, had developed and demonstrated a process for micro-machining deep (100s of microns), high-aspect ratio ($>10:1$) and sub-micron positioned through-silicon features on double-sided silicon wafers. Preliminary work had been done on copper plating and planarizing these microfabricated chips, though these aspects of the process were largely undeveloped. Tom had been able to show very promising proof-of-concept electroplating of the larger, lower-aspect-ratio (2.5:1) antibias coils. The metalization process we employ leaves metal on the surfaces of the chips, which must be removed. Tom had demonstrated the ability to remove this metal in antibias chips with copper on one side.

An overview of the project goals and status as it stood when I assumed the lead, produced as part of a request for proposal in an attempt to outsource the chip fabrication project to Teledyne Scientific, can be found in Appendix B. Several aspects of the chip fabrication had yet to be developed completely. In particular:

1. Consistent void-free copper filling of the high-aspect-ratio “wires” (6:1) on the curvature chip had not been achieved.
2. A compatible planarization process for removal of excess copper had not been developed for the curvature chips.
3. Planarization of double-sided chips had not been achieved.
4. Post-processing procedures (e.g. annealing of plated copper) were largely unexplored.
5. The methods and procedure for the assembly and precise alignment of chips into a coil stack was required.
6. Destructive and non-destructive methods of testing the electrical and mechanical robustness of the chips required refinement and further development.
7. Operational capabilities and limitations (i.e. current carrying, heat sinking and distribution, handling) needed to be determined.

The unique nature of our chip coils (large, high-aspect-ratio, high density features distributed over a large area of both sides of the silicon) place them in a regime which has been sparsely explored by industry and academia. The development of our process has therefore been guided by a combination of prior research on thin (nanometer scale) films and on sparse 100 nanometer to 10 micron-scale high-aspect ratio features.

The development of an appropriate process was further complicated by the fact that we had a very limited number of quality samples from the first (micro-machining) stages of production, necessitating excessive care on ensuring a high yield on these samples. Tom was available to introduce me to many aspects of the chip fabrication procedure, but the

knowledge and skill required to micro-machine quality double-sided chips remained his alone. Thus, most process development was done on stand-in chips differing in key ways from the actual production samples. In most cases, single sided chips were used at first. The process was then extrapolated to double-sided chips and tested on discarded and early-stage (often with shallower features) or defective (with broken sidewalls, for example) double-sided chips. A “best effort” would then be performed on the production samples. In an ideal scenario, many identical production chips would have been available for testing and process development, allowing for more parallel testing and more accurate indication of the process efficacy.

The details of the chip micro-machining procedure, as well as the specific parameters of the electroplating process that we developed for both the curvature and antibias wafers can be found in Appendix A of Purdy [70]. In this chapter, I will touch on a few of the challenges that we overcame in developing the electroplating recipe, including some suggestions for further improvement. Mainly, though, I will report on subsequent stages of development, including our planarization, annealing, chip verification and testing, and bonding procedures. Preliminary measurements of the thermal capabilities of the completed chip stack are also reported.

3.1 Electroplating

Before the electroplating process is to begin, a metal layer (or two) is deposited onto the silicon-oxide substrate. In our case, a thin layer of titanium serves as an intermediary between the silicon-oxide and a layer of copper. The metal layers serve two main purposes. First, the copper serves as a seed on which more copper can be deposited by electrochemical means. Second, the metal serves to provide electrical connectivity between all areas where we desire copper to grow.

There are two basic strategies for electroplating deep features. The most straight-forward way is to deposit copper (or electrically connect) only the bottom of the future wires and grow the copper in a bottom-up fashion. This approach works when all of the wires can be joined with minimal electrical resistance such that the growth rate in all of the wires is nearly the same. The other approach is to coat all of the surfaces with a seed layer and grow copper everywhere, even where it will not be wanted eventually. This approach has the advantage of great electrical conductivity, and is useful in cases like ours where the wires are very long (so that the resistance of a seed layer lying within them would be large) or where many features that require copper are electrically isolated. The challenge of this approach is filling deep features completely, without voids. In order for this to be done, the copper in the trenches needs to grow from the side in, with a preference for faster growth deep in the trenches.

The electroplating procedure that we developed was based on the work of Sun et al. [85]. There, 10 micron square by 70 micron deep vias were filled with copper using a commercial plating apparatus, a homemade plating solution, and special current sequence involving a

forward plating pulse, a quick reverse de-plating pulse, and a period of no current flow. Electroplating dynamics are very complex, and *a priori* it was not clear how to best apply their work to ours. There were several key differences between their application and ours. Notably,

1. Their features are square and sparse, resulting in a much lower (several orders of magnitude) density of “deep” features. Our features are not square, but are distributed, and are densely packed. Our features have 3 distinct dimensions rather than two.
2. Their features are evenly distributed over the entire surface, whereas ours have very high density in some areas and low density in others.
3. Their commercial plating apparatus agitates the solution in a different way than our homebuilt one.

The plating process is, by necessity, very sensitive to the local concentration of different molecular constituents and chemical intermediaries present in the plating solution. The local electric field and rate of various chemical reactions on the chip surface interact with (and maintain) these chemical gradients, allowing the copper deposition to occur in the non-uniform, bottom-up manner that we require to completely fill the deep wire and via features present on the chips. Thus, we were not surprised that the parameters that worked for Sun et al. [85] did not apply directly, in an obvious way, to our needs.

Through much iteration of plating parameters (solution additive concentration and maintenance procedures, plating current values and pulse lengths) we found a process that produces adequate results. Limitations of the process as it stands are that plating quality is not uniform across the whole sample being plated, though it may be adequate for our needs. Improvements to the plating apparatus that would create more symmetric and/or uniform electrical and chemical environment, such as more uniform agitation, would likely improve this, though the degree to which the inherent non-uniformity of the chip (rather than the plating apparatus) is responsible for plating non-uniformity we do not know.

Another possible improvement may be the use of a larger amount of plating solution, either by performing the plating in a larger bath or using some form of recirculation with an external reservoir. The overall amount of copper deposited in our process is quite large, and the plating solution likely changed a lot in constitution throughout the process. We found that using fresh solution in each plating run produced the best results; it’s possible that keeping the solution more fresh would have further improved things.

Photos of electroplated samples are shown in Fig. 3.2. Such polished cross-sections, along with surface inspection, are the tools with which we can evaluate the process at this stage. Resistance measurements of the final polished and annealed sample would likely provide a better “figure of merit,” however other variations in the chip (etch depth, most importantly) can make comparisons between chips difficult. Most importantly, we did not have the time or resources to optimize the process using this better figure of merit (resistance) and each step of the process was tackled individually by necessity. Details of our electroplating apparatus, and our adapted plating solution, parameters, and process, are detailed in Purdy [70].




Figure 3.2: Photos of the curvature coil electroplating. (Left) A polished cross section of the wires in the curvature coil shows voids. (Center) Using our best known parameters, void-free filling has been achieved, at least in the section of the chip shown here. Extreme copper overfilling in this section of the chip was deemed necessary to insure adequate copper filling in the large pads of the chip. In retrospect, overfilling with so much copper may not have been prudent, as removing such a thick layer of copper poses a challenge. (Right) The surface of the chip as it appears during electroplating.

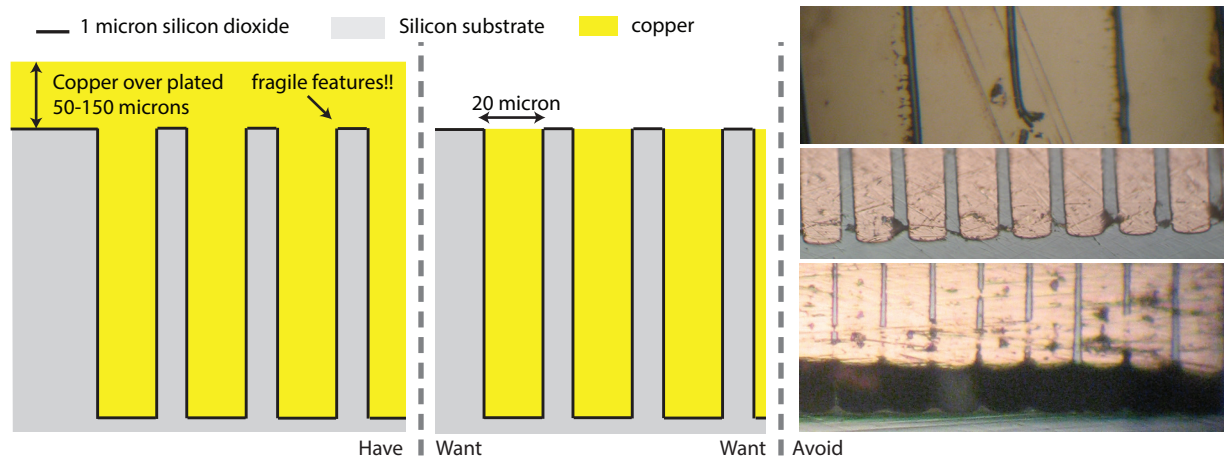


Figure 3.3: Planarization scheme and photos of some common pitfalls of purely mechanical polishing: (top to bottom) surface shorts, cracking of wire sidewalls, and wholesale delamination of copper.

3.2 Planarization

The plating procedure we developed leaves 50–100 microns of copper on the surfaces of the chip (really the “chip” at this stage is a quarter wafer containing several individual coils). This excess copper must be removed while preserving delicate silicon features and copper below the surface of the chip. This process is complicated by several factors:

1. Chips with copper on both sides do not have a flat reference surface and copper is not evenly distributed on the surfaces. Thus, purely mechanical polishing that does not distinguish between Cu and Si or Si-oxide will tend to remove a substantial amount of material below the chip surface.

2. Each side of the wafer may have a different amount of copper because of differences in the duration of plating. Ideally, this factor could be minimized by more consistent plating.
3. The wafer must be mounted to a chuck to be polished, and standard (reversible) mounting procedures involve hot wax bonding. Thermal expansion and bowing of the piece during temperature cycling may result due to different thermal expansion properties of Si and Cu, and due to different amounts of Cu on different sides of the chip. Thermal stress on the Cu-Si surface can cause delamination of copper from this surface and/or damage to delicate silicon features.
4. Traditional mechanical lapping procedures are rough on the thin Si walls.

Photos illustrating some of the effects of these complications are shown in Fig. 3.3. These sort of poor polishing results lead us towards developing a chemical-mechanical polishing (CMP) process for our chips.

CMP, when done right, can allow for highly selective and highly planar material removal. In the ideal case for our application, the CMP process would be unable to remove any Si or Si-oxide while gently removing only copper that is in contact with the polishing surface. The way in which this ideal is approximately achieved is an active area of research [88, 75], but it can be schematically thought of as a competition between several processes. First, a chemical reaction dissolves copper, removing small amounts of copper from the surfaces of the chip and quickly dissolving any copper particulates that have been removed by mechanical means. Second, a slow reaction passivates the copper surface, inhibiting its chemical removal. Third, slight mechanical agitation removes the passivating layer, along with small amounts of copper, allowing the chemical etch to continue in abraded areas. The resulting effect is that copper areas that are not in contact with the polishing pad are protected by the passivation reaction, but that copper in contact with the pad is removed quickly by the combined chemical and mechanical process. Ideally, the Si and Si-oxide parts of the chip are affected only by the mechanical aspects of polishing, which can be minimized when the primary means of copper removal is chemical.

After a lot of trial and error and in consultation with CMP experts in industry and in other departments at Berkeley—notably Professor Fiona Doyle and former students of hers, Shantanu Tripathi, now at Intel, and graduate student Seungchoun Choi, along with Professor David Dornfield’s graduate student Joshua Chien—we found an effective CMP slurry and process for our double sided chips. This process is detailed in the rest of this section.

First, excess copper on the surface of the chips is reduced from 100–200 microns to 10–50 microns by a purely chemical wet copper etch process. Reducing the thickness of the copper layers not only reduces the CMP time (the CMP process removes material slowly), but it reduces the stresses on the surface of the chip during CMP and during the thermal cycling employed in mounting. The primary reason the chips have so much additional copper on the surface in the first place is that the narrow wire features accumulate copper at a much



Figure 3.4: Chip after wet etch. The area around the pads, which were masked by super glue, is etched at an accelerated rate.




Figure 3.5: CMP chuck and machine in action. The chuck (left) consists of an aluminum mount with inset region for the handle wafer. The handle wafer has space in the center for the piece to be polished (shown) and consists of a 4 inch wafer with slightly thinned wafer pieces epoxied around the edges. The chip is affixed to the center of the handle wafer with bonding wax.

higher rate than the large-area pads used for connecting the leads to the surface. In order to bring the level of copper in the pads to the level of the silicon substrate, the rest of the chip must accumulate an excess. The pads, though, have only a tiny excess. Thus, during the initial chemical removal of excess copper, the large interconnection pads are protected by covering them with super glue, so they are not etched.

The wet procedure employs a commercial etchant, Aluminum Etchant Type A from Transene Co., Inc., heated to 50°C. We etch the chips one side at a time by covering the non-etched side with Kapton tape. The etch speed is roughly 3 microns/minute. The etch time is varied in order to leave 10–50 microns of copper on the surface of the chips and, most importantly, above the chip wires. Interestingly, we observe the etch rate is accelerated around the masked (super glue-covered) regions, as shown in Fig. 3.4.

Next, we employ our CMP procedure to remove the remaining excess copper. We use a G&P Technology POLI-500 DC Chemical-Mechanical Polisher with a custom chuck, shown in Fig. 3.5. The general procedure for using the polisher can now be found in chapter 10.3

table on	90 rpm	oscil on	DI off
head1 on	60 rpm	slurry 1 on	slurry 2 off
head2 off	0 rpm	head1 down	head2 up
pressure	100 g/cm ²		
slurry flow	100 mL/min		

Table 3.1: POLI-500 DC settings. The piece to be polished is affixed to head1.

of the Marvell Nanolab manual [50] (though I don’t believe this document existed when our work was underway), and include our procedure for conditioning the IC1000 polishing pad. The POLI-500 settings we settled on are shown in Tab. 3.1.

The slurry that ended up working best for us consisted of (per liter deionized water) 10–20 grams 0.3 micron alumina abrasive powder, 7.5 grams glycine, 0.15 grams benzotriazole (BTA), 50 milliliter of 30% hydrogen peroxide, with smaller amounts of abrasive used for the final stages of planarization. In this slurry, BTA plays the role of an inhibitor [34], reducing the chemical etching of copper by the hydrogen peroxide and glycine, and allowing the copper to be polished with minimal “dishing” of the copper beyond the Si or Si-oxide surface, as shown in the profiles in Fig. 3.6.

In order to polish both sides of the chips, we first run the CMP on one side using our optimized slurry. The chip is mounted using a low-temperature bonding wax with only moderate hardness (MWM070 from South Bay Technologies) to reduce thermal expansion stresses. Not all surface copper is removed from this side. We cease CMP when polishing appears to be uneven and Si-oxide is exposed on part of the wafer. Si-oxide can be identified by its deep blue or Newton ring rainbow coloration. We then continue CMP on the second side. When polishing appears uneven, the wafer/wax/chuck assembly is temperature cycled to relieve strain on the wafer caused by the thermal expansion. This periodic re-seating of the polished piece by temperature cycling is critical to obtaining uniform polishing of double-sided samples. Polishing continues on the second side until all surface copper is removed. Finally, CMP is resumed on the other side, temperature cycling as necessary, until all of the surface copper is removed.

3.3 Singularization, Annealing, and Testing

The polished chips were inspected and photographed under a microscope for defects acquired during the polishing and plating steps. An ASIQ profilometer can produce a surface profile of features $\gtrsim 10$ microns in size with sub-nanometer precision, allowing us to evaluate “dishing” and surface quality. Sample microscope images and ASIQ profiles are shown in Fig. 3.6.

Next, the chips were “singularized,” that is, the individual coils were cut out from the quarter-wafer of similar coils that have been processed thus far. The singularization process was performed on an Esec dicing saw. The dicing procedure is well described in Purdy [70], Appendix A, section 1.5.

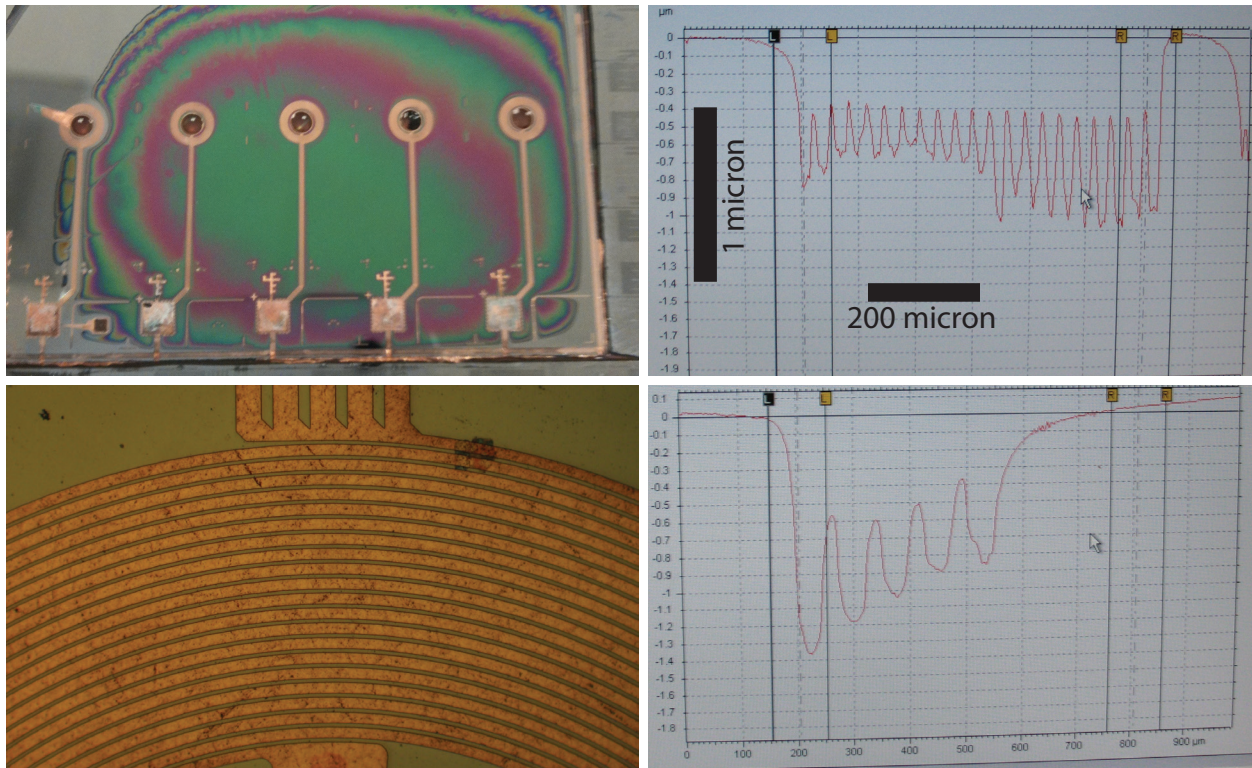


Figure 3.6: CMP results. Top left, even polishing of one side of the curvature chips is indicated by the mostly intact oxide layer. The oxide is originally one micron in thickness. Right, elevation profiles of the annulus (top) and leads (bottom) show sub-micron dishing of the copper and sidewall features. Bottom left, a close up of one of the polished coils.

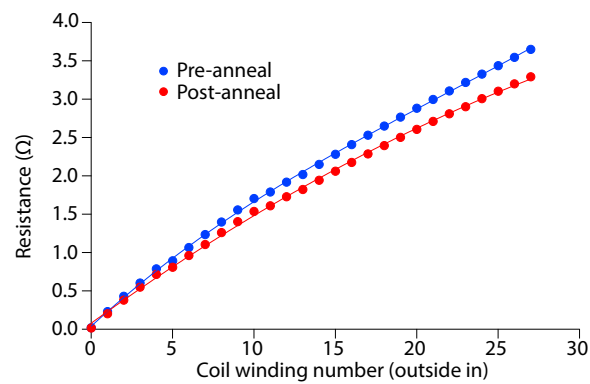


Figure 3.7: Chip resistance pre- and post-anneal.




Figure 3.8: Resistance curves without and with a short. Fits are to a quadratic polynomial.

Annealing of the plated copper can reduce stresses within the copper and on the silicon as well as decrease the resistance of the copper. For thin films, at least, the reduced resistance arises more quickly than the bulk restructuring of the copper crystal, and can actually be achieved with a relatively short low temperature anneal [10]. This is good because heating our chips too much can cause the copper to pop out of the silicon owing to the differences in thermal expansion. Therefore, we do a 250–300°C anneal of the polished chips for several hours in a nitrogen atmosphere. We have seen this reduce the resistivity of the copper by approximately 10%, as shown in Fig. 3.7. The resistance can be accurately measured by four-point probe, where probes are placed with the aid of a micro-motion stage under a microscope.

Such four point resistance measurements, taken at various (and numerous) positions along the turns of the polished chips, provide a non-destructive window into the quality of the chips at this point in the process. If the depth of the etching on the chips is known and approximately uniform, such measurements provide an indication of local electroplating quality and can indicate the presence and location of shorts between turns of the coils. Shorts, while undesired, need not pose too much of a problem if they are sufficiently large that they completely “cut out” a single turn of the coil, and can for the most part be compensated for with a small amount of extra current. A four point resistance measurement as a function of current across a short could determine the robustness of the short by measuring, essentially, its heating via resistance rise.

3.4 Chip stack bonding

We bonded the stack of four chips together using a low-viscosity, heat-curing, non-conductive epoxy on a home-built temperature-controlled micro-manipulation stage. The chip stack was back-illuminated with infrared light. Each of the chips had alignment markings that could be




Figure 3.9: Chip stack bonding. (Left) Alignment marks of two chips, as seen through the chips, are nearly in alignment. (Right) A cross section of the bonded chips reveals a sub-10 micron wide interface between the two chips.




Figure 3.10: Our best, fully-assembled ring trap coils.

viewed through the stack on a microscope-mounted camera. First we bonded the outer pairs of coils, one curvature and one anti-bias coil. During the bonding and alignment process, etalon fringes were generated by the narrow gap between chips and served as an indicator of parallelism.

Then, as shown in Fig. 3.1 and in Appendix B, the two sides of the stack were bonded and aligned with a 2 mm spacer in between. The spacer was made of highly-parallel pieces of silicon. We did not precisely measure the alignment precision of our bonding setup, but I estimate it is the range of 10 microns. The parallelism across the chips was on the order of 5-10 microns, which was the thickness of the epoxy layer. Images of alignment and of the bonded interface, with annotated scale bars, are shown in Fig. 3.9.

We measured the thermal conductivity through a stack of chips and found that a thin layer of non-conductive epoxy more effectively transfers heat than a thicker thermally conductive epoxy, by about a factor of two. This was likely because all surfaces are very smooth and thus the 50 micron-scale silver particles present in many thermally conductive epoxies could not do their typical job of filling in the cracks.




Figure 3.11: Mounting and heat sinking the chip stack. (Right) A CAD model of the chip stack mounting and heat sinking hardware, as described in the text, with inner ceramic mount (top) in contact with outer ceramic mount (middle) in contact with copper cooling tubes (middle, bottom). (Left) An assembled (test) chip stack mounted in the holding arm, which is formed by the copper cooling tubes, ready to go into vacuum.

Our final best-effort chip stack is shown in Fig. 3.10.

3.5 Mounting, thermal testing, and *in situ* monitoring

The cooling and mounting apparatus we employed was designed by our postdoc Anton Öttl and is shown in Fig. 3.11. The mounting setup could be considered a sandwich, with the chip stack as the meat and two symmetric form-fitting ceramic mounts as the bread. The ceramic bread interfaces with complementary ceramic pieces that are epoxied to copper cooling tubes, through which we could flow coolant. The ceramic employed was a machinable aluminum nitrate called shapal.

By applying a known heat load to the chip stack and measuring the equilibrium temperature at multiple places along the path of heat flow—on the chips, on the inner ceramic mount, on the outer ceramic mount, and on the coolant-filled copper tubes—we determined

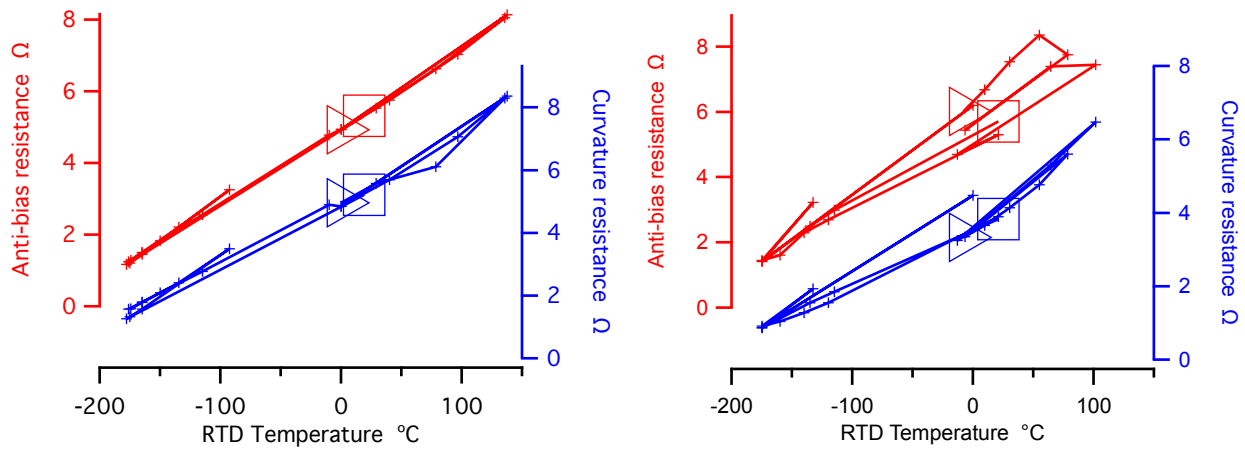


Figure 3.12: Thermal cycling curves for two different pairs of chips, each with a bias and anti-bias coil. The chips tested on the left appear more robust than the ones on the right, as indicated by the lack of hysteresis. The first points in each series of measurements are indicated with the large right-pointing triangle, while the final point is indicated by the large square.

the heat sinking capabilities of our chip stack and the cooling/mounting apparatus in which it sits. In particular, we determined the thermal bottlenecks.

Overall, we measured a temperature rise at the chips of 1°C per Watt. The thermal bottleneck was the interface between our outer ceramic mount and the copper tubing. Heat transfer at the interfaces between the chips and the inner ceramic mount, as well as between the inner and outer ceramic mounts, was facilitated by thin graphite sheets. Graphite has excellent thermal conduction only in-plane, so the graphite sheets were kept thin. These thermal joints were overall very effective at transmitting heat. On the other hand, the interface between the outer ceramic mounts and the copper tubing relied upon thermally conductive epoxy.

Regardless, with the setup as-is, we determined that we should be able to constantly run our target current, four amps, through each chip with a temperature rise of 150°C . This temperature rise would be acceptable if we were to cool the chips to -60°C with methanol, which was within our capability. With a modest two-fold improvement of the ceramic-copper junction the estimated temperature rise would be reduced to 75°C , which would be both more manageable and leave a greater margin for error. Still, we had a lot of doubt as to whether the chips could survive repeated cycles of temperature variation.

Thus, we investigated the question of what sort of temperature cycling our chip stack can stand, at least in the short term. To do this, we temperature cycled the chips by flowing various amounts of current while monitoring the resistance and temperature of the coils. Changes of the resistance with temperature are normal, and can be predicted if assumptions about the distribution of resistivity in the copper wire are made, but hysteresis or changes in resistance at a given temperature after cycling can indicate structural changes and failure of

the stack. Preliminary data tracking the resistance of two different pairs of chip coils as the temperature was cycled are shown in Fig. 3.12. The data suggested that chip resistance over time may be a good indicator of robustness and possible failure. Preliminary measurements of the radio-frequency transfer function of early versions of the chips were also performed. Such transfer function measurements could also be tracked over time and may serve as additional indicators of chip quality and reveal shorts or other material changes over time.

3.6 Moving on

Because of issues with other aspects of the apparatus, most notably the magnetic transport of cold atoms into the area where the chips were to be placed, and because of the fact that we only managed to assemble one chip stack that we deemed suitable for use in our apparatus, we abandoned efforts to use the microfabricated chips in our actual apparatus. We deemed the substantial effort required to fully realize the magnetic transport through the tight apertures imposed by the holding arm and the chip stack not to be worthwhile, as the likelihood that the chip stack would fail in some way upon installation and repeated temperature cycling was substantial. Without a backup stack, such a failure would make all of that additional effort for naught.

As true believers in the sunk cost fallacy, we decided to move on. The “best-effort” chip stack shown in Fig. 3.10 remains in a drawer in the lab, and we decided to pursue ring trap and spinor physics in the MOT chamber, in spite of the fact that this chamber lacked some important features that might have been desired of a top-tier apparatus. Desirable, but missing features included plentiful optical access, an RF antenna, and magnetic field control. Nonetheless, we made it work, as evidenced in Marti [49] and the rest of this dissertation.

Were we to continue towards our initial goal of trapping atoms in the magnetic ring trap, our next step would have been to measure the magnetic fields produced by the assembled stack of chips. We began development of a system for accurately measuring the fields of the chip stack that I will briefly describe.

First, we obtained a sensitive magnetic field probe from Micro Magnetics with a spatial resolution of must a few microns. We mounted this probe on a 3D micro-translation stage from Newport so that we could sample the magnetic field in the region of the chip stack with high resolution and precision. The magnetic probe was linear and sensitive only at fields much less than one Gauss, which meant we could not simply map the field at anything approaching our process currents. Our plan, which we never implemented, was to put the magnetic sensor in loop with a compensation coil. This compensation coil would be large enough to provide a very uniform magnetic field in the region of the chip stack and would be driven with a precision current source such that the magnetic field at the position of the sensor be precisely canceled. The relevant magnetic field signal would be the current through the compensation coil.

In order to mitigate the effects of large near-DC magnetic field noise, our plan was to drive both the chip stack and the compensation coils with AC currents. Owing to differences

in complex impedance, the phases of the driving currents of the chips and the compensation coil might well have needed to vary. We never got around to deciding on a concrete implementation for the phase-locked AC driving currents.

Chapter 4

Spin vortices

Though we abandoned our efforts to place atoms in an ultra-smooth magnetic storage ring, we did not abandon ring shaped traps all together. Our explorations of Bose condensed ^{87}Rb in a $30\ \mu\text{m}$ diameter *optical* ring trap are detailed in our publication, Ref. [46] and in Ch. 3 of Marti [49]. In this work, we characterized the collective modes of the scalar BEC order parameter— phonons—and showed how they could potentially be used to perform accurate and bias-free rotation measurements, even in an imperfect trap. While we did not have sufficient signal-to-noise to measure a rotation signal, we performed a detailed experimental and theoretical analysis of the noise in our system and demonstrated that it operated at the atom shot noise limit. Our signal was limited primarily by unexpectedly large damping rates, which we now attribute to finite-temperature effects.

The bias-free nature of our phonon gyroscope has its origin in the fact that the trapped BEC is irrotational at low energy: The order parameter, the phase of the condensate, must be single-valued, and thus rotation, which corresponds to an accumulated phase around a closed contour within the condensate, must be quantized.

The order parameter(s) of spinor condensates admit a wide range of interesting structures [31], some of which are topologically non-trivial. In future versions of this chapter we will summarize our work creating, imaging, and characterizing interesting spinor structures. Our main focus will be on the spin vortex in a ferromagnetic condensate, which is topologically non-trivial in the magnetization space but trivial in the larger order parameter space.

While not officially part of my dissertation, my plan is to release an updated version of this chapter¹ wherein I discuss the following topics:

- Making the first spin vortex.
- Varying the vortex core size.
- Analyzing the topological character of the spin vortex.

¹Try the Stamper-Kurn group website <http://ultracold.physics.berkeley.edu> or the ArXiv pre-print server <http://www.arxiv.org>.

- Comparing the stability of the spin vortex with similar looking spin structures without rotation.
- Accounting for missing transverse magnetization in the imaged spin vortices
- Improvements to the experimental apparatus that would be necessary to improve upon this work.

At the time of this writing, a discussion of our imaging techniques and images of some of our spin vortices can be found in Ch. 4 of [49].

Ultimately, our work on spin vortices was put on hold to make the aforementioned improvements to the apparatus. While many of these improvements were made (some are discussed in Ch. 5), we never returned to the spin vortices. Instead, we initiated a productive investigation of an even simpler excitation of the spinor condensate: the magnon, a Goldstone mode associated with symmetry breaking of the ferromagnetic order parameter. Eventually, a return to the study of spin vortices using the techniques that we developed might prove quite fruitful.

The basics of magnons are discussed and experimentally validated in Ch. 5 of [49]. In the following chapter, I use these properties of magnons to extend accurate thermometry of the BEC into new regimes of low entropy. I also show how magnons can be used to cool a spin-polarized BEC.

Chapter 5

Thermometry and cooling with magnons

The low [56, 38], even (gasp!) negative [9] temperatures reached in ultracold atomic and molecular [66] systems are deservedly praised, sometimes even in the popular press (*cf.* Ref. [16]), for allowing experimenters to “access new regimes of physical phenomena.” However, the emphasis in this archetypal phrase is almost always misplaced. The key phrase should really be *new access* rather than *new physical phenomena*. Most of the strange quantum phenomena the cold atom community explores are quite common. A paper titled “Cooling to 0.002 times the Fermi temperature” could represent a huge technological and scientific breakthrough in cold quantum gases, but would be rather unremarkable if the cooling in question was performed by placing a block of aluminum in a typical household refrigerator¹.

New access to phenomena that exist in a wide range of materials is very appealing. In exchange for a massive effort in cooling technology, ultracold atomic gases have provided experimenters with a remarkable window into the quantum world, allowing time- and space-resolved readout and control of archetypal many-body and condensed matter systems. Still, in spite of their name, relative to relevant physical energies, many ultracold gases might be considered quite hot relative to solid-state systems of interest. The entropy per particle (S/N), rather than the temperature, is the pertinent parameter when comparing physically similar systems with different length and energy scales. See Table 5.1 for a sample of S/N in various systems.

In some simple cases, the entropy per particle can be calculated in a straightforward way. For non-interacting harmonically trapped quantum degenerate Bose and Fermi gases, for

¹In solid aluminum, the electron Fermi temperature $T_F \sim 135,000$ K.

System	Temperature	$S/N (k_B)$
Nitrogen gas	300 K	22
Rubidium (metal)	300 K	8.8
Water	300 K	8.4
Sodium (metal)	300 K	6.1
Ideal Fermi gas, harmonic	T_F	4.2
Ideal Bose gas, harmonic	T_c	3.6
Aluminum	300 K	3.0
Ideal Bose gas, box	T_c	1.3
Two-level systems	∞	0.7
Diamond	300K	0.28
Ideal Fermi gas, harmonic	$0.05 \times T_F$	0.25
Ideal Bose gas, box	$0.3 \times T_c$	0.2
Ideal Bose gas, harmonic	$0.3 \times T_c$	0.1
Aluminum	77 K	0.06
Aluminum	2 K	0.0001

Table 5.1: Table of entropy per particle in some reference systems. Even when entropy per particle cannot be calculated from first principles, the measured entropy per mole, or “standard molar entropy,” can be looked up at for certain substances and a specified temperature in chemical tables.

example,

$$(S/N)_{\text{Bose}} = k_B \times 4 \frac{\zeta(4)}{\zeta(3)} \left(\frac{T}{T_c} \right)^3 \quad (5.1)$$

$$(S/N)_{\text{Fermi}} = k_B \times \frac{\pi^2}{2} \frac{T}{T_F} - \mathcal{O}(T/T_F)^3, \quad (5.2)$$

where ζ is the Riemann zeta function. For strongly interacting Bose gases deep in the Bogoliubov phonon regime, and, in general, for a medium that supports phonon excitations, such as a crystal lattice well below the Debye temperature²,

$$(S/N)_{\text{phonon}} = k_B \times \gamma \frac{2\pi^2}{45n\hbar^3 c^3} (k_B T)^3, \quad (5.3)$$

where c is the speed of sound, n is the number density of the medium and γ the number of phonon polarizations (one for BEC, three for a crystal). In the general case where the equation of state is not known, or must be calculated numerically, as in Ch. 6, estimating the entropy per particle is either not quite as straightforward or simply impossible. Regardless, using these simple relations, the 500 pK sodium gas of Ref. [38] has an entropy per particle $S/N \sim 1.5 k_B$, about that of the same sodium cooled as a solid to liquid nitrogen temperature.

²Note that the Debye temperatures of the alkali are quite low. For rubidium it is 56 K.

At first glance, (weakly interacting) gaseous and (strongly interacting) solid sodium may seem too dissimilar for the entropy per particle to provide a meaningful comparison, but because of the manner in which experiments on ultracold atoms are often performed, this is not necessarily the case. Weakly interacting gases are routinely transformed, ideally adiabatically, into strongly interacting lattice-trapped systems [22, 3, 65, 86, 78] where the physics of solids are at the forefront. Without powerful techniques for reducing the entropy in optical lattices, which have largely, though not entirely [4], remained elusive [53], the entropy of the initially prepared gas sets a lower bound on the entropy that can subsequently be achieved in the transformed system.

Entropies per particle in the 0.5–few k_B range commonly reached in cold atom experiments have been sufficient to implement quantum simulators exploring a range of fundamental effects, including the superfluid to Mott insulator transition of the Bose-Hubbard model and the universal dynamics of Fermi gases at the BEC-BCS crossover [8]. However, the entropy required to reach the regimes of non-trivial (e.g. antiferromagnetic or frustrated) magnetic ordering in multi-component Bose-Hubbard [1, 11] and Fermi-Hubbard [64] systems is expected to be just beyond what is commonly achieved. At even lower entropy, perhaps a few orders of magnitude lower than has been achieved with cold atoms thus far, with fractional filling of the lattice, such systems may be important to understanding the d -wave superconductivity known to exist in the cuprates [29]. Nobody knows the full richness (or not) of the multi-component Hubbard models at partial filling, and whether exotic superfluid/superconducting phases exist somewhere in them, because they remain theoretically intractable. To be able to simulate them, then, would be a Very Big Deal.

Why does a small amount of entropy spoil some quantum phases and not others? I am not aware of a good discussion of this question, but I can offer a few thoughts. Some systems display quantum behavior at high entropy owing to entropy compartmentalization. For example, in a BEC, the non-condensed fraction can carry a large amount of entropy over a wide range of low energy modes, leaving a zero entropy condensate. The non-condensed fraction interacts weakly with the condensate, making the distinction between condensed and non-condensed fractions quite stark. In other systems, such as Mott insulating states, strong localization allows entropy to be compartmentalized spatially, for example, at the edge of the Mott lobes. The center of a Mott lobe may appear to have quite low entropy.

Interesting spin models often present different situation, as they frequently rely on the fact that spins are not localized, and interactions between spins necessarily constitute a dominant energy scale. In “strongly correlated” systems, interesting states may require correlations that build up over long distances, and perhaps it doesn’t take many thermal fluctuations to disrupt the long-range quantum ones. In Fig. 5.1, the “spin” pattern is revealed at $S/N = 0.1 k_B$. If, however, “Berkeley” were more spread out, or had larger quantum fluctuations, it might still be obscured. In addition, the number of degrees of freedom in spin systems is typically quite small. For a two-component spin system with one spin per site, the maximum entropy is $S/N = k_B \times \ln 2 \approx 0.7 k_B$. Clearly, to reach an ordered phase, the entropy needs to be less than the maximum. Little is known about the microscopic physics behind the d -wave superconducting states, but in solid state systems they appear with




Figure 5.1: What does entropy look like? In general, it’s not easy to determine, let alone “see”, the entropy in a system. Here is a special case. Each of these images is 200 by 200 pixels, and each pixel can independently take on one of 16 grayscale values. When all values are equally likely, the entropy per pixel is $\log(16) \approx 2.77 k_B$. In digital images, entropy is more frequently measured in bits: $\log_2(16) = 4$ bits per pixel. If not all colors are equally likely—here white is favored as the entropy is reduced—the expected entropy per pixel is the average $S/N = \sum_i p_i \log(1/p_i)$, where p_i is the probability of a pixel having color i . At lower entropy, a message may emerge from the “thermal” noise if the contrast of this figure is high enough. This message has some noise of its own, however not all noise is thermal! Incidentally, the GIF format does a good job of compressing these images, with the overall file size tracking closely to the true average entropy per pixel.

less doping (i.e. less filling) and at lower temperatures than the antiferromagnetic phases.

The greatest promise of quantum simulation using cold atoms hinges on accessing models about which theory has been able to say very little, and in order to do this, experimenters need to be able to reliably create and characterize samples with lower entropy than has been done before. In this chapter, I describe our experimental efforts exploring cooling and thermometry of a highly degenerate bulk ferromagnetic $F=1$ Bose gas using collective spin excitations known as magnons [27]. Our study reveals that forced evaporative cooling can produce gases an order of magnitude colder than had been commonly thought, as low as $T/T_c \approx 0.02$ in our experiments, which we estimate (see Ch. 6) corresponds to $S/N \sim 0.001 k_B$ at equilibrium. Our thermometry is remarkably precise, operating near fundamental noise limits, and could pave the way for precision studies of thermodynamics at low entropy. In addition, we demonstrate a novel cooling method, cycled decoherence cooling, a form of demagnetization cooling, that does not rely on lowering the trap depth, and that we find can reach $S/N \sim 0.3 k_B$. With technical improvements, it may be able to reduce the entropy even further. Producing low entropy bulk gases does not solve the problem of producing low entropy gases in optical lattices, however, the ability to create and characterize such samples is a substantial step forward.

5.1 The challenge of low entropy

Both thermometry and cooling rely on distinguishing the excitations that differentiate a system from its ground state. At very low entropy per particle the system is very near its ground state and detecting (or removing) these minor differences becomes difficult, as doing so requires very high sensitivity to small and rare fluctuations in the state of the system as well as detailed knowledge of the ground state. For example, the temperature of an ideal gas is manifest entirely in the variance in its kinetic degrees of freedom. At high temperatures, crudely sampling the velocity of the constituent particles is sufficient to establish the temperature. At sufficiently low temperature (and entropy), however, the variations in velocities can get arbitrarily small, and so arbitrarily precise measurements would be required. Likewise, reducing the entropy of such a gas in isolation requires a velocity dependent force or the selective removal of its most energetic constituents.

In a weakly interacting Bose gas with harmonic confinement, the situation is not improved, as the number of non-condensed atoms, N_{th} drops very quickly below the condensation temperature, $N_{\text{th}}/N \approx (T/T_c)^3$, so that not only does the variation in the kinetic motion decrease, so does the number of atoms with varying motion. For weakly interacting gases at temperatures comparable to the chemical potential, $k_B T \lesssim \mu$, the kinetic degrees of freedom acquire a collective phonon-like character, and can no longer be characterized by their velocity. In this regime, it becomes increasingly difficult to distinguish thermal fluctuations from quantum ones [74, 84], and it is not obvious on theoretical grounds [63, 45, 33] that evaporative cooling should continue to function.

While direct measurements of the system itself may not readily yield its temperature, thermometry may still be possible by placing the system in thermal contact with a system that would exhibit considerably higher entropy at the same temperature. Additionally, unless the thermometer is prepared at the same temperature as the system in question before being placed in thermal contact, the thermometer can be expected to change the temperature of the system it is measuring. For a highly degenerate ferromagnetic condensate, spin excitations can constitute such a thermometer.

Spin excitations can constitute an excellent thermometer for several key reasons. First, in most of the commonly used ultracold gases, collisions that change the projection of the gas magnetization along the magnetic quantization axis essentially do not occur. Thus, the number of spin excitations created out of an initially polarized gas can be controlled precisely. In particular, the number of spin excitations created can be made to be as large as possible while still keeping the number of spin excitations below the critical number for their condensation at equilibrium. In addition, in the presence of a weak magnetic field, spin excitations that differ in magnetic moment from the majority spin component can be addressed spectroscopically or separated from the majority component by a magnetic field gradient, allowing them to be imaged with high signal to noise independent of the majority gas by common techniques, as in Ch. 4.

Any spin or pseudo-spin excitation that thermalizes at constant (pseudo-)magnetization might make an excellent thermometer, but magnons have some advantages over generic spin




Figure 5.2: Schematic of effective trapping potentials for different hypothetical spin excitations. (a) In the Thomas-Fermi limit, the majority spin condensate (red) fills the harmonic trap $V(x)$ up to the level of the chemical potential μ . (b,c,d) The interaction energy between minority spins and the majority spins is represented in gray. (b) For magnons, the peak mean field interaction energy $n(0)g$, with $n(0)$ the majority density at the trap center, is equal to μ . The effective potential $V_{\text{eff}}(x) = V(x) + n(x)g$ has a flat bottom. (c) A minority spin with $n(0)g > \mu$ interacts with the majority spins more strongly than the majority spins interact with each other and is repelled from the trap center. (d) Minority spins with $n(0)g < \mu$ experience a weakened harmonic trap in the region occupied by the minority condensate.

excitations. Neglecting small magnetic dipolar interactions, the magnon is a gapless Goldstone excitation corresponding to the rotational symmetry of real spins [47]. Accordingly, the s -wave scattering lengths that characterize the interactions between two majority-spin atoms, $a_{-1,-1}$, and between a majority and minority-spin atom, $a_{-1,0}$, are identical. Thus, the mean-field interaction energy ng between magnons and the majority-spin atoms, where n is the majority density and $g = 4\pi\hbar^2 a_{-1,0}^2$ is the interaction strength parameter, is exactly the same as the interaction energy between majority-spin atoms themselves. Because of this, in the Thomas-Fermi limit, the magnons see an effective box-like potential, shown schematically in Fig. 5.2(b), as variations in the majority density compensate energetically for variations in the confining potential. The consequences of this are that magnons disperse as free particles down to $T = 0$ and have a relatively high density of states, and thus larger critical number for condensation, compared to confined particles of the same mass. Additionally, magnons can be created with zero energy, less the Zeeman energy. Apart from the presence of magnetic field gradients, the Zeeman energy is constant and conserved, and thus does not contribute to the gas dynamics [79]. Mathematically, the Zeeman energy can be eliminated by boosting the system to a frame rotating at the Larmor frequency.

In contrast, spin excitations involving a different hyperfine state are gapped, in general, and creating them involves depositing (or removing) interaction energy inhomogeneously, as illustrated schematically in Fig. 5.2(c,d). In addition to providing a heating mechanism, the low energy density of states of the spin excitations is reduced by the variation in interaction energy, and with it the critical number of condensation, reducing the potential temperature signal. In practice, in ^{87}Rb , all of the inter- and intra-species hyperfine ground-state scat-

tering lengths are nearly equal [89], and the temperature at which magnons should have a distinct advantage over hyperfine pseudo-spin excitations is quite low, on the order of a few percent of the chemical potential. Thermometry using co-trapped pseudo-spin impurities is a common technique in Fermi gas systems [72, 59, 77, 54].

5.2 Cooling at low entropy

The same characteristics that make spin excitations a good thermometer—high relative energy and entropy per particle—enable them to cool the majority gas in two distinct ways. First, spin excitations created in a degenerate Bose gas can decrease its temperature by a process known as decoherence [39] or demagnetization [19, 56] cooling. Immediately after the RF pulse that creates the magnons is applied, the minority spin population has the same energy and momentum distribution as the initially polarized degenerate Bose gas, with a large condensed fraction, which carries no energy above the chemical potential, and a small normal fraction, which carries on the order of $k_B T$ of energy per particle. Upon thermalization, the normal fraction within the minority spin gas increases, bounded from above by its critical number for condensation, and the energy and entropy of the minority spins increases. The energy gained by the minority is energy lost by the majority, as, for magnons, the entire process occurs at constant energy, neglecting from the effects of magnetic field inhomogeneities and dipolar interactions. After thermalization, the magnons can be expelled from the sample and the process of decoherence cooling can be repeated. We call this “cycled decoherence cooling.”

This cooling process can also be understood as a form of demagnetization cooling. The majority gas and the magnon gas are the longitudinal and transverse magnetization, respectively, of a magnetized gas. Transverse magnetization is created by slightly tipping the magnetization of the longitudinally polarized gas. While the longitudinal magnetization is conserved owing to the symmetry breaking applied magnetic field, the transverse magnetization is allowed to dissipate at constant energy.

In Fermi gas systems, a similar process would lead to heating as the minority spins would be created with relatively high average energy (some fraction of the majority Fermi energy) and equilibrate to a lower average energy (some fraction of the much smaller minority Fermi energy). Energy lost by the minority spins is, more or less, energy gained by the majority, leading to heating overall. Thus, to perform thermometry using spin excitations in Fermi gases, the spins are typically allowed to equilibrate at higher temperature before evaporating to the final trap depth. It turns out that this sort of evaporation protocol can also lead to enhanced cooling in bosons by increasing the efficacy of evaporative cooling.

Forced evaporative cooling from a trap with an effective trap depth $U_{\text{eff}} \gg k_B T$ has a cooling power proportional to the number of thermal excitations with excitation energies above U_{eff} . In a weakly interacting single-component degenerate Bose gas, the number of thermal excitations is determined by the temperature and is independent of the total particle number, fixing the evaporative cooling rate. By seeding the gas with additional

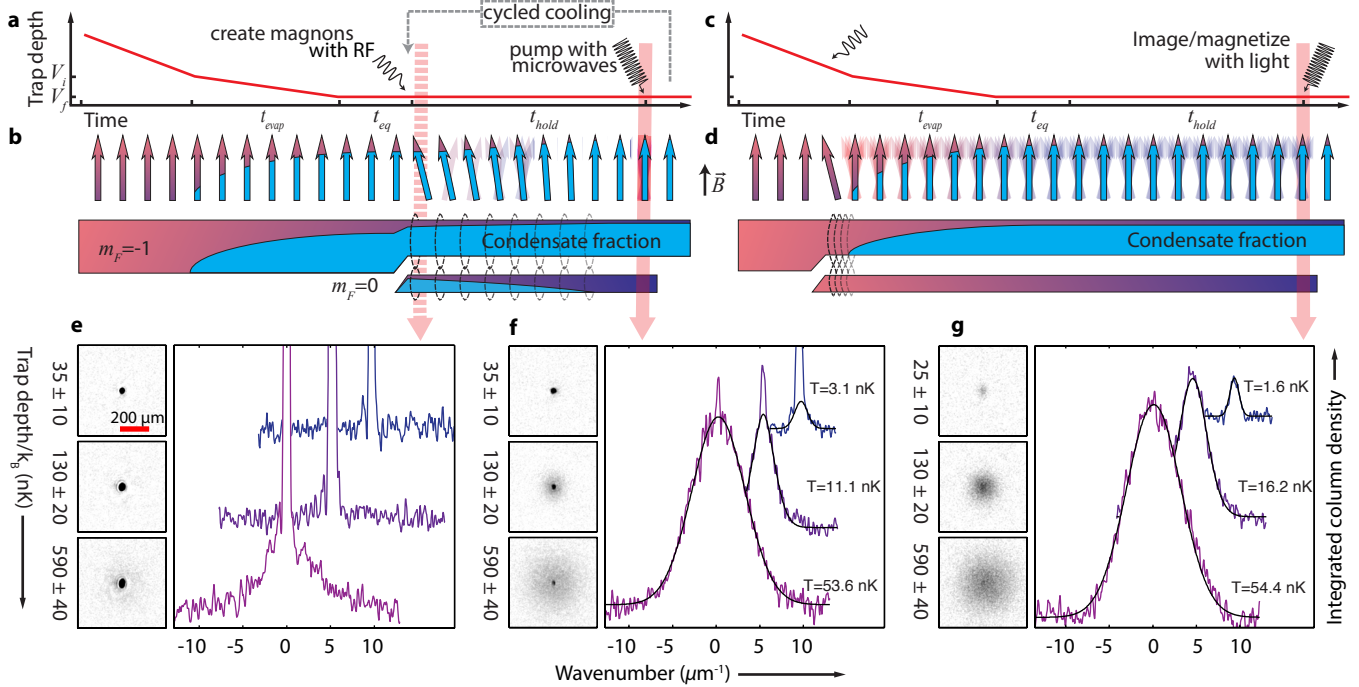


Figure 5.3: Schematic illustration of magnon cooling and thermometry protocols with sample thermometry data. Magnons are created via a RF pulse in a spin-polarized gas at the final trap depth (a, b) or at an intermediate trap depth (c, d). The magnetized gas can be viewed schematically as an arrow, initially pointed in the direction of an applied magnetic field. Introducing magnons corresponds to tilting the gas magnetization. Unlike the longitudinal magnetization, the transverse magnetization is not conserved, and as the magnon population loses coherence, the average transverse magnetization dissipates. This is indicated schematically by the magnetization arrow losing its tilt. The incoherent transverse magnetization is indicated by the fuzzy background arrows. Equivalently, the magnons can be viewed as an admixture of $m_F=0$ atoms whose coherence dissipates quickly at the intermediate trap depth and slowly at the final trap depth. Magnons created at the final trap depth cool the gas via magnon decoherence cooling, which can be repeated in the same sample as cycled magnon decoherence cooling. Magnons created at the intermediate trap depth improve the efficiency of forced evaporation. In both cases, the magnons serve as a sensitive thermometer. As the gas is cooled, the majority gas (or, equivalently, longitudinal spin component) becomes highly condensed and the temperature signal is lost (e). In contrast, the magnons ($m_F=0$ component or minority gas) provide a relatively large temperature signal even at low temperatures (f,g).




Figure 5.4: Sample preliminary cycled decoherence cooling data. Top row images show the majority condensate and thermal cloud atoms in time-of-flight. The bottom row of images shows the corresponding distribution of minority atoms. The graph shows the absolute and relative (to the critical temperature T_c) temperatures extracted from the majority atoms. The estimates of temperature are not well calibrated and the critical temperature is unreliable owing to large uncertainty in the trap frequencies and trap anharmonicity. Regardless, cooling is clearly working and shows real promise.

spin excitations at constant total particle number, the total number of thermal excitations, and thus the evaporative cooling power, increases. We call this sort of cooling “magnon assisted evaporation.” Once again, in Fermi gases, a different situation presents itself, as the excitations of the minority spins will in general have much lower energy than the excitations of the majority, and will not contribute to evaporation. Our cooling and thermometry schemes are illustrated schematically in Fig. 5.3.

In this chapter, we discuss the details of our implementation of sensitive and accurate low-entropy thermometry using magnons, as well as our efforts characterizing the two different magnon cooling schemes. As most of the most salient scientific results are in our publication, Ref. [62], the focus of this chapter will be the experimental and behind-the-scenes details and evolution of that work. It will probably help the reader to be familiar with that reference before proceeding.

5.3 First cooling: a benchmark for improvements

Our first attempts at doing cycled decoherence cooling went remarkably smoothly. For reasons of history and convenience (we had microwaves set up and not RF) we chose to cool by mixing in a small amount of $F=2$ into the polarized $F=1$ condensate using a brief

microwave pulse. That cooling was occurring was evident immediately from the *in situ* images of the majority gas. Even at this early stage, it seemed that a magnetic field gradient was important in achieving thermalization in a finite amount of time. An example of our preliminary cooling data are shown in Fig. 5.4.

The preliminary data were very promising, but indicated a need for substantial improvement to be quantitatively convincing. For example, we had little confidence in the estimated temperature and critical temperature. To do accurate thermometry, lots of things need to work right and be well calibrated. There are many things to cross-check to assure us that we're moving towards reliable measurements, or not.

The BEC critical temperature for an ideal non-interacting Bose gas in a harmonic trap depends on several quantities that we need to calibrate accurately:

$$k_B T_c = \hbar \bar{\omega} \left(\frac{N}{\zeta(3)} \right)^{1/3}. \quad (5.4)$$

Here, $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean of the trapping frequencies, N is the number of atoms, and ζ is the Riemann zeta function. We can experimentally produce a cloud with temperature $T = T_c$ by finding settings at which a small condensate barely forms. If everything is well calibrated, we should find that the T_c calculated using Eq. 5.4 and our measurements of N and $\bar{\omega}$ agrees well with our estimate of the temperature of the cloud. The equality should hold over a range of T_c and N .

In the Thomas-Fermi regime, the chemical potential, μ , is another quantity that can allow us to find consistency between our parameters, relating the trap frequencies, condensate number, and the *in situ* size of the condensate:

$$\mu = \frac{1}{2} m \omega_i^2 R_i^2 = \frac{\bar{\omega}}{2} \left(15 \hbar^2 N_c a (m \bar{\omega})^{1/2} \right)^{2/5}, \quad (5.5)$$

with R_i the condensate radius along the direction with trap frequency ω_i , and N_c the condensate number.

Finally, we can measure temperature by different methods, and these methods should agree in regimes where they are valid. For example, near T_c we can measure the temperature of our gas by imaging its size as a function of the time following its release from the optical trap (“time-of-flight” or TOF) along two different axes, with two independently calibrated imaging systems.

Our first attempts at cooling with magnons clearly showed a gas getting colder, but exactly how much we could not be sure, as the thermometry failed several of these consistency checks. This spurred us to make a range of changes to our thermometry, imaging, and trapping setups. In the next sections I will describe many of the calibrations and improvements to our apparatus that ultimately enabled our final study.

5.4 Momentum space focusing

In a non-interacting, non-degenerate classical gas (obeying Boltzmann statistics), the temperature is manifest in both spatial and momentum degrees of freedom *separately*. At every point in position space, the momentum distribution is identical, apart from a multiplicative constant, and likewise for points in momentum space, as is easily seen from the separability of the Hamiltonian entering the exponential of the classical phase-space probability distribution function:

$$P(\mathbf{x}, \mathbf{p}) = e^{(\mu/k_B T - V(\mathbf{x})/k_B T - \mathbf{p}^2/2mk_B T)} / Z. \quad (5.6)$$

Unwanted degrees of freedom can essentially be ignored, as integrating over them yields trivial multiplicative factors.

For a gas at or near quantum degeneracy, the Boltzmann distribution in Eq. (5.6) is replaced by the appropriate Bose or Fermi distribution, and the separability of real-space and momentum-space distributions is no longer obvious. However, the dependence of the momentum distribution on the trapping potential turns out to take a fairly simple form, at least for potentials that are well described by a single power law (e.g. box and harmonic). In the case of bosons, integrating over j quadratic degrees of freedom yields a probability for the remaining quadratic degrees of freedom \mathbf{s} of

$$P(\mathbf{s}) \propto g_{j/2} \left(e^{(\mu/k_B T - (\mathbf{c} \cdot \mathbf{s})^2/k_B T)} \right), \quad (5.7)$$

where \mathbf{c} is a vector of coefficients for the \mathbf{s} degrees of freedom, e.g. $\sqrt{1/2m}$ for momentum degrees of freedom, and $g_{j/2}$ is the Bose function or polylogarithm of order $j/2$. The signature of the integrated degrees of freedom (e.g. the trapping potential if we are observing only momentum degrees of freedom) is in the order of the polylogarithm and, indirectly, the chemical potential μ . μ appears in Eq. (5.6) but is irrelevant to the shape of the probability distribution; it's a multiplicative factor. In Eq. (5.7), however, μ appears inside the polylogarithm, and can affect the overall shape of the probability distribution. Note that $g_j(z) \approx z$ for $z \ll 1$ and thus far from quantum degeneracy, or at sufficiently large $\mathbf{c} \cdot \mathbf{s}$, Boltzmann-like statistics prevail.

In-trap (position-space) images of non-condensed gases thus provide direct access to the temperature, provided one knows the potential V . However, *in situ* images are rarely used for thermometry for several reasons. First, in the presence of a condensate, the non-condensed part of the gas may be largely obscured, and high optical density makes imaging such distributions challenging. Even in the absence of a condensate, tightly confining traps push the relevant temperature signal to small length scales where imaging may be less sensitive. Finally, the temperature signal is sensitive to the details of the trapping potential, which may be poorly constrained.

Time-of-flight measurements are the thermometric workhorse of the cold-atom community and overcome many of the limitations of *in situ* thermometry. In most imaging setups,

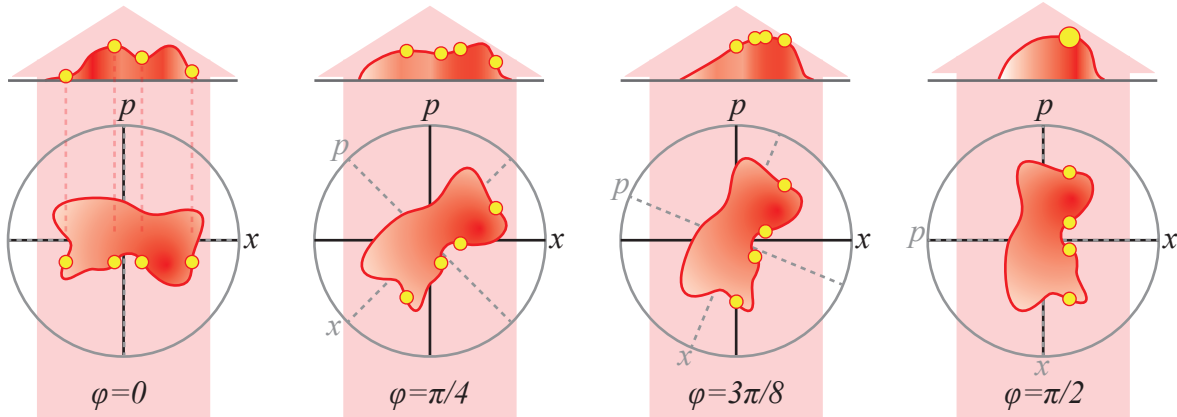


Figure 5.5: Principle of momentum space focusing. The action of a harmonic oscillator is to rotate a system in phase space, shown here with one space and one momentum dimension versus the harmonic oscillator phase $\phi = \omega t$. A few points with the same *in situ* ($\phi = 0$) momentum are marked with yellow dots. The imaged distribution, at the top of each column, shows only position space. The dashed axes show the positions of the initial axes rotated through angle ϕ . At $\phi = \pi/2$, points with the same momenta converge to the same real-space coordinate.

the signal-to-noise for temperature measurements increases with free expansion time (for some time, at least) and the expanding cloud asymptotically approaches a shape that reflects its momentum distribution.

Momentum space focusing [87, 76] (MSF) takes time-of-flight one step further in yielding pure momentum-space distributions in an adjustable (and perhaps signal-to-noise optimized) time, removing most of the dependence on the trap geometry from temperature calculations. For magnons, which can experience an effective trapping potential that varies with their temperature (which is related to how much of the trap they explore beyond the edge of the majority condensate) [18], this feature is quite attractive.

The key insight involved in MSF is that harmonic motion manifests as rotation in phase space (with coordinates appropriately scaled), and thus a distribution of momenta manifests as a distribution of positions after a quarter cycle (rotation in phase space by $\pi/2$), and vice versa, as shown in Fig. 5.5. Atoms with initial momentum \mathbf{p} will converge on real-space point $\mathbf{x}_p = \mathbf{p}/m\omega$, where ω is the harmonic trap frequency. By releasing our atoms from their trap into a two-dimensional isotropic harmonic magnetic potential, with the symmetry axis coincident with the imaging axis, the atoms' *in situ* in-plane momentum distribution will appear after a quarter cycle, neglecting interactions after the trapping light is extinguished.

In the following subsections, I will describe how we implement and characterize MSF in our system.

5.4.1 Experimental procedure

Two-dimensional confinement is provided by the same coils that produce the spherical quadrupole³ and bias fields for the MOT. Near the center of the quadrupole field, the field produced can be written in terms of the offset and gradients as

$$\mathbf{B}(\mathbf{r}) = \mathbf{B}_0 + \mathbf{J}_B \mathbf{r}, \quad (5.8)$$

where \mathbf{B}_0 is the bias field, and \mathbf{J}_B is the Jacobian matrix of partial derivatives of \mathbf{B} . For the quadrupole, in a basis with the z -axis aligned with the vertical symmetry axis, this Jacobian has a simple diagonal form:

$$\mathbf{J}_B = \begin{pmatrix} \partial_x B_x & 0 & 0 \\ 0 & \partial_y B_y & 0 \\ 0 & 0 & \partial_z B_z \end{pmatrix} = \begin{pmatrix} -B'/2 & 0 & 0 \\ 0 & -B'/2 & 0 \\ 0 & 0 & B' \end{pmatrix}, \quad (5.9)$$

where we write B' for the vertical gradient $\partial_z B_z$. The trapping potential for a spin-polarized atom with magnetic moment μ and bias field $\mathbf{B}_0 = B_0 \hat{z}$ oriented along the z -axis is

$$V(\mathbf{r}) = \mu |\mathbf{B}| \approx \mu \left(B_0 + zB' + \frac{x^2 + y^2}{2} \frac{B'^2}{4B_0} \right), \quad (5.10)$$

where the approximation assumes B_0 dominates the other terms in parenthesis. The in-plane curvature of V implies a harmonic trapping frequency $\omega^2 = \mu B'^2 / 4mB_0$. The quadrupole and bias field are configured such that atoms in the $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = 1, 2\rangle$ states are confined by the field curvature and supported against gravity by the field gradient.

The MSF sequence begins 5 ms after the optical trapping light is extinguished. During this time, the purging and state preparation (Sec. 5.13) sequences are completed and the atoms we wish to image are in the $|F = 2, m_F = 1\rangle$ state and in free fall under the influence of gravity. At this point we quickly ramp on the bias field $B \approx 36$ G over 500 μ s. B' rises more slowly from zero to $B' \approx 42$ G/cm with an exponential time constant of roughly 5 ms. This gradient overcompensates for gravity ($mg = \mu \times 30.4$ G/cm) in order to slow the falling atoms to a stop. After approximately 45 ms, the atoms are stopped (Fig. 5.6b) and B' is set to 30.4 G/cm in concert with B in order to maintain a constant transverse curvature of $B'' \approx 12$ G/cm², which corresponds to a transverse trapping frequency of $\omega \approx 2\pi \times 3$ rad/s. Calibrations, discussed below in Sec. 5.4.3, will reveal a more precise version of this number.

With the magnetic field so configured, the vertical field gradient opposes gravity, levitating the center of mass of the cloud at a constant height. Nonetheless, the atoms are free to continue their expansion in the vertical direction, as motion along the vertical axis is unconfined. I will have more to say about this in Sec. 5.4.4.

³Aside: Etymologically, *multipoles* are assembled from a hodgepodge of prefixes of Greek and Latin origin. *Pole* can swing both ways. *Multi* comes from Latin; the Greek would be *poly*. *Mono* and *di* are Greek, while *quad* is Latin and *Octo* can go either way. For higher-order moments, people tend to steer towards the Greek, as in *tetrakaidecapole* for the fourteenth moment.

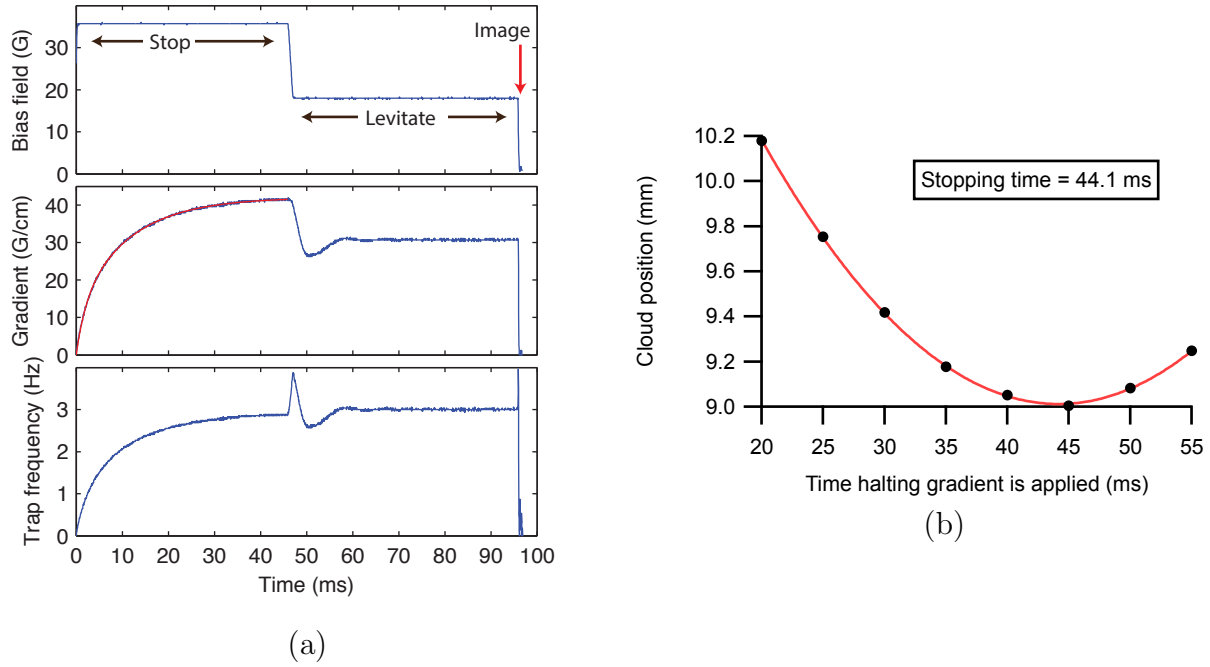


Figure 5.6: Experimental sequence for MSF. (a) Gradient and bias fields are first configured to slow and stop the falling cloud. Blue traces show measured coil currents plotted as magnetic fields. The gradient coil ramps up exponentially with a weak 2 ms time constant and dominant 11 ms time constant (red fit). As soon as the atoms are stopped (b), the fields are reconfigured to levitate the cloud until it is focused. The MSF trap frequency is held roughly constant.

One complication that we need to consider carefully is that the ramps of the gradient field are slow, limited by the slew rate and bandwidth of the driving power supply, as shown in Fig. 5.6a and Fig. 5.7a. At one point early in developing MSF on our apparatus we optimized the gradient field turn-on using the resonant switch in order to have appropriate field curvature established within a few milliseconds, however we realized later that these settings were never used for subsequent calibrations and data. To figure out the impact, or lack thereof, of non-constant field curvature during MSF, we simulate the process using the experimentally measured field ramps. These simulations will verify that MSF behaves as we expect, and will guide our calibrations.

5.4.2 Simulating the magnetic focusing lens

The results of simulating the magnetic focusing lens with our measured gradient and bias coil ramps are summarized in Fig. 5.7. The only free parameter in these simulations is the equilibrium magnetic trap frequency, $\omega_0 = 2\pi \times 3.05$ rad/s, which I adjust to achieve focusing at the observed time $t_{\text{foc}} = 97$ ms. The focusing trap is assumed to be harmonic

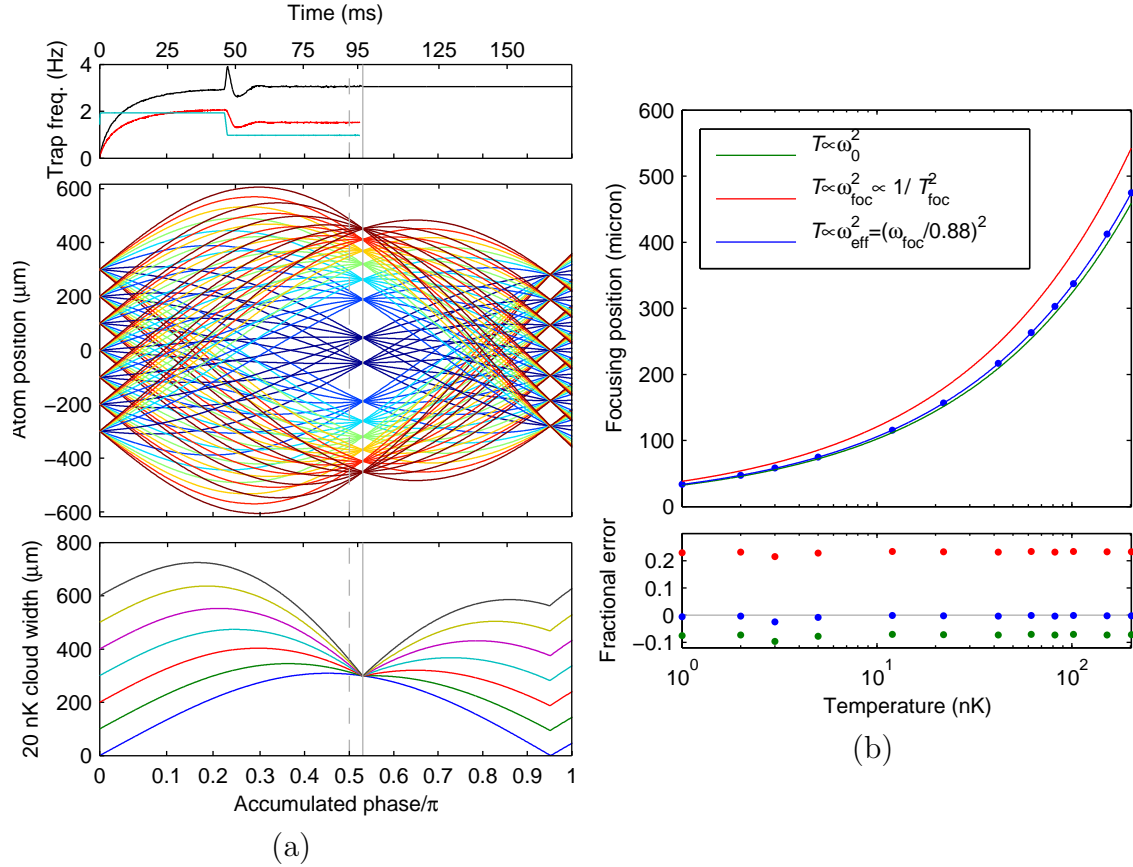


Figure 5.7: Simulation results for MSF with our measured field ramps and focusing time of 94 ms. (a) The gradient (red) and bias coil currents (teal) in arbitrary units, along with inferred magnetic trap frequency (black), are shown in the top panel. The middle panel shows the real-space trajectories of atoms with different momenta (different colors) originating from different locations in the cloud. The bottom panel shows width of a 20 nK cloud with different initial sizes (different colors). (b) The simulated real-space focusing position of atoms with various energies (blue dots) are plotted in the top panel along with the focusing positions that would be expected in a SHO with various frequencies. Fractional errors of the different SHO models, shown in the bottom panel, are constant. In these plots, for simplicity, representative momenta and temperatures are related by the relation $p^2/2m = 2k_B T$.

and isotropic with time-varying trap frequency $\omega(t)$ in both transverse directions, and the longitudinal direction is ignored.

The simulations show several important features of the 2d-harmonic focusing lens. First, the magnetic lens is free of aberration in spite of the time varying frequency. Individual momenta focus to a point (lines with the same color in Fig. 5.7a cross at a point), and different momenta focus at the same time (all colors cross themselves at the same time).

Further, the time at which the momenta focus is close to, but not quite, the time at which the oscillator phase $\theta(t) = \int_0^t \omega(\tau) d\tau = \pi/2$, as it would be for a perfect turn-on. Importantly, the focusing time t_{foc} corresponds to the time at which the cloud width achieves a local minimum. For our cloud, with one very long axis, this local minimum is quite sharp. t_{foc} can thus be readily measured, as done in Fig. 5.8. Note that the focusing time does not coincide with the global maximum of the width of small clouds, as it does in the constant ω case.

Finally, the simulation shows in Fig. 5.7b that the real-space positions at which different momenta focus are well described by the simple harmonic oscillator (SHO) model with a single effective momentum-independent frequency $\omega_{\text{eff}} = \omega_{\text{foc}}/0.88$ such that $\mathbf{x}_p = \mathbf{p}/m\omega_{\text{eff}}$. Thus, in spite of the varying curvature of the magnetic lens, the momentum space images it produces are faithful to the simple model presented above, with one simple adjustment. Computing temperatures using $\omega_{\text{foc}} = 2\pi/4t_{\text{foc}}$, instead of ω_{eff} would lead one to underestimate temperatures by 25%.

5.4.3 Calibrating MSF

The key parameters in MSF are the magnification, $1/m\omega_{\text{eff}}$, which allows us to infer momenta from real-space data, and the focusing time, t_{foc} , at which we should take images in order to find the momentum distribution in focus. Because these two parameters are related, knowing one of them is sufficient.⁴ Nonetheless, we can boost our confidence, and precision, by measuring both parameters quasi-independently.

The focusing time t_{foc} is perhaps the easiest parameter to pin down. The long axis of the condensate focuses very tightly, making the focusing point easy to spot to within a few milliseconds by eye. The focusing time can also be easily extracted from non-condensed gases, so long as the focused size is smaller than the initial spatial extent of the gas. In this limit, the cloud size has a minimum that can easily be fit. The colder the gas—the smaller the focused size—the sharper the minimum. Furthermore, neglecting effects of the expanding condensate and misalignment of the focusing plane (see Sec. 5.4.4), the size of the non-condensed gas should be equal in all directions at the focusing time. Sample data for these calibrations are shown in Fig. 5.8a. With $t_{\text{foc}} = 97 \pm 3$ ms, and including 0.5% uncertainty in the ratio $\omega_{\text{foc}}/\omega_{\text{eff}}$, we find $\omega_{\text{eff}} = 2\pi \times 2.9 \pm 0.1$ rad/s.

Measuring ω_{eff} directly, without the help of a simulation, is easy if one knows *a priori* the momentum distribution and an approximate value of t_{foc} : simply look at the spatial distribution around the focusing time and use the relation $\mathbf{x}_p = \mathbf{p}/m\omega_{\text{eff}}$. Thus, accurately measuring the temperature via a traditional method, such as time-of-flight, can serve to calibrate MSF directly. We would prefer, however, to measure ω_{eff} without reference to temperatures, as this allows us to gain confidence in the temperatures measured by MSF and time-of-flight by their mutual agreement.

⁴In principle, ω_0 would also be sufficient information to fully calibrate MSF. However, in our system, observing more than a single oscillation in the MSF trap proved difficult, owing to persistent vertical expansion and residual vertical motion of the gas.




Figure 5.8: Calibrating MSF. (a) The time at which the vertical size of the condensate reaches a minimum can be discerned easily by eye (top). The focusing time t_{foc} of the non-condensed atoms shows a distinct minimum as well, even if it is hard to see by eye. Coupled to simulations, these measurements allow us to determine the magnification of MSF. (b) The magnification of MSF is measured directly by fitting in-trap center-of-mass oscillations *in situ* and in MSF to an exponentially decaying sinusoid. The decay times τ and oscillation frequencies ω_i are the same, but the oscillations differ in phase by $\pi/2$, the expected phase difference between real-space and momentum-space quadratures of the same oscillation, and they have different amplitudes. The frequency associated with the MSF magnification is $\omega_{\text{eff}} = \omega_i A_i / A_m$. Vertical axes are shown on the same scale.

Rather, in Fig. 5.8b we illustrate another (related) way of measuring ω_{eff} by comparing the amplitude and phase of center-of-mass oscillations measured *in situ* and in MSF. In effect, we are measuring the in-trap velocity and mapping it to a position in MSF, but by considering the amplitude of oscillation we have a straightforward way of averaging many measurements. Given the amplitude A_i and frequency ω_i of in-trap motion, the maximum center of mass momentum is $p_i = m A_i \omega_i$. This maximum momentum maps via MSF to the maximal position $x_p - x_0 = A_m = p_i / m \omega_{\text{eff}}$, with A_m the measured amplitude of the oscillations as imaged in momentum space. Thus, $\omega_{\text{eff}} = \omega_i A_i / A_m$. Our calibration yields $\omega_{\text{eff}} = 2\pi \times 2.8 \pm 0.2$ rad/s.

This calibration depends only weakly on the chosen value of t_{foc} . The center-of-mass motion is quite small, and the MSF trap is nearly centered on the atoms during all phases of the in-trap oscillation. Thus, even with the field ramps as slow as they are, as can be seen from the trajectories that begin near zero in Fig. 5.7a, the center-of-mass velocity reaches a quadratic turning point near t_{foc} . Therefore, we can consider this direct measurement of ω_{eff} independent of the one obtained from t_{foc} and average the two values, weighting by uncertainty.

As it stands, the measurement based on t_{foc} dominates the weighted average, but we gain confidence knowing that the two measurements agree with each other. The direct measurement of ω_{eff} could be repeated, or even better, more points could be sampled at early times of the *in situ* oscillation to reduce the uncertainty in our determination of A_i . The 3% uncertainty in our measurement of ω_{eff} is the primary source of systematic uncertainty in the temperature measurements.

5.4.4 Practical considerations and limits of MSF

So far, we have been relying on two assumptions that merit further scrutiny. First, we have assumed that the momenta of the atoms are fixed when we release them from the optical trap, and that the momenta we observe in MSF correspond to the momenta *in situ*. For weakly interacting atoms, for example non-degenerate alkali atoms, this assumption is typically very good. Clouds composed exclusively of such atoms expand nearly ballistically and isotropically. However, the interactions between atoms in the condensate are not negligible, and interactions between the condensed and non-condensed parts of the gas may not be either. In addition, we have mostly ignored the vertical axes of MSF, imaging, and the optical trap, and we have tacitly assumed they are well aligned to each other. Poor alignment of these axes, and the expansion of the condensate along one or more of them—aggravated by interactions—can broaden the momentum signal anisotropically, as discussed below.

For a large condensate deep in the Thomas-Fermi regime, the change in the velocity of the condensate atoms is well described by a hydrodynamic theory [14]. In particular, upon being released from confinement, a condensate expands most rapidly in the direction in which it had been most tightly confined, with radii growing as $R_i(t) = R_i(0)b_i(t)$. The scaling parameters b_i , with $i \in \{x, y, z\}$, satisfy the equation

$$\ddot{b}_i = \frac{\omega_i^2}{b_i b_x b_y b_z}, \quad (5.11)$$

where ω_i are the (extinguished) trap frequencies. To understand intuitively why the condensate expands rapidly along the tight axis, recall that the phase of the trapped Thomas-Fermi condensate in equilibrium is uniform, evolving at a rate $\dot{\theta} = \mu(0) = \mu(\rho) + V(\rho)$, where ρ is a radial coordinate and the minimum of the trapping potential is $V(0)$. The interaction energy varies across the cloud, but the sum of the interaction and potential energy is constant⁵, and

⁵Reminder: in the Thomas-Fermi approximation one ignores the kinetic energy of the ground state.




Figure 5.9: The expansion of the condensate. (a) Interaction energy turns into kinetic energy after the optical trap is turned off, shown here for the case of optical trap setpoint $V_{\text{ODT}} = 2$. The magnons might inherit this velocity, which, via the relation $mv_{\text{RMS}}^2/2 = k_B T$, corresponds to a temperature, shown in (b), where v_{RMS} is the RMS velocity of the condensate after 100 ms.

hence the phase evolves at the same rate everywhere. When the trapping potential is suddenly turned off, this is no longer the case, and the phase accumulates most rapidly where the interaction energy $\mu(x)$ is largest. A phase gradient develops and is most steep where $\mu(x)$ changes most rapidly: along the direction of (formerly) tightest confinement. In a superfluid, a gradient in phase corresponds to a velocity. An equivalent semi-classical picture is that the interaction energy constitutes a type of potential energy, and thus the expansion force is proportional to its gradient, allowing the interaction energy to be reduced as quickly as possible.

Our optical trap is highly anisotropic, and the condensate expansion is predominantly directed along the vertical axis. This can be a nice feature, as it means the image of the condensate in momentum space is quite small, allowing narrow momentum distributions to be seen unobscured. However, the vertical expansion is not without side effects. First, the optical density of the condensate increases dramatically as its shadow on our camera gets smaller; we make no attempt to accurately count condensate atoms in MSF. Second, a condensate expanding at 5 $\mu\text{m/ms}$ (see Fig. 5.9a) will be nearly 500 μm long at the focusing time. The Rayleigh range of our approximately 8 μm -resolution imaging system

is around $250 \mu\text{m}$ and thus it is unsurprising that MSF images with a condensate include aberration at and around the location of the condensate. Finally, if the tight z-axis of the optical trap does not coincide precisely with the unconfined axis of MSF, a component of the condensate’s “vertical” expansion may be present in the transverse momentum signal, causing the condensate to appear larger or more anisotropic than it really is.

The simple Thomas-Fermi hydrodynamic theory applies only to the condensate; in general, it’s not particularly useful in quantitatively understanding the interactions between condensed and normal gases⁶. Nonetheless, to the extent that the condensate can be considered a medium for excitations—magnons or non-condensed atoms—we can imagine that, in the worst case, the expanding medium imparts its own additional expansion velocity onto its excitations. Thus, the additional velocity acquired by the condensate during MSF serves as an *upper bound* on the additional velocity obtained by the normal atoms, neglecting interactions between only normal atoms. In service as a thermometer, this extra velocity imparted by the condensate has the potential to set limits on the temperatures we might be able to measure, and so it is worth estimating this effect, as in Fig. 5.9b.

In our system, all three axes—MSF, imaging, and optical trap—would ideally be aligned to gravity and thus closely coincide. The spherical quadrupole coils used in MSF were roughly aligned to gravity upon their installation, but there is likely some residual tilt. A small tilt around an in-plane axis by an angle θ would rotate the plane in which focusing occurs around that same axis by 3θ , but would not lead to any other more onerous aberrations. To calculate this, consider a variation on Eq. (5.8) with the quadrupole coils rotated around the \hat{x} -axis,

$$\mathbf{B}(\mathbf{r}) = B_0 \hat{z} + R_x(\theta) \mathbf{J}_B R_x^T(\theta) \mathbf{r}, \quad (5.12)$$

where \mathbf{B} is written in the gravity-referenced coordinates, and \mathbf{J}_B is diagonal in the rotated coil-referenced coordinates. Once again, write the trapping potential V as in Eq. (5.10) and compute the Hessian (curvature) matrix \mathbf{H} of V ,

$$\mathbf{H}_{ij} = \frac{\partial^2 V}{\partial x_i \partial x_j},$$

evaluated at the origin and keeping terms up to $\mathcal{O}(\theta)$. Unlike in the case of perfect alignment, \mathbf{H} is not diagonal in the gravity-referenced basis. However, the trapping frequencies, the eigenvalues of \mathbf{H} , are unaltered, while the axes of oscillation/no confinement, the eigenvectors of \mathbf{H} , are rotated from the gravity-referenced basis by an angle 3θ .

The imaging axis can easily be well aligned to gravity: when viewed from above, falling atoms should not vary in their transverse position. The optical trap can also be aligned well to gravity, especially in the direction of the long axis of the condensate, where our procedure for leveling is most sensitive (see Sec. 5.8). However, at the lowest temperatures explored in this work, the anisotropy of the momentum signal seems to be the most sensitive probe of the

⁶Ref. [32] has a well written section on the ways in which normal and condensed fluids can interact. Though much work has surely advanced the state of the art in this area, this reference suffices as a first pass for many experimenters.

optical trap alignment in the tight in-plane direction. Thus for two reasons—broadening of the momentum distribution from in-plane condensate expansion and additional broadening from out-of-plane expansion projected into the imaging plane—we will use the temperature measured along the long axis of the optical trap as our primary thermometer. Regardless, misalignment should never cause us to underestimate the temperature.

5.5 Imaging

To extract quantitative, accurate temperatures from our images, we need to make sure our images accurately represent the distribution of the atoms we wish to image. Taking and interpreting absorption images of atoms is so integral to most ultracold atom experiments that most of the salient points have been detailed countless times. On the other hand, because imaging systems are so common, it is easy to take many important aspects of their operation for granted. Regardless, in the course of conducting this research we found that understanding our imaging process was both essential and non-trivial.

Our apparatus has two imaging systems, shown in Fig. 5.10, oriented at right angles to each other. The side imaging system has a wide field of view, roughly 1 cm across, and is used primarily for atom counting and calibrating the atoms' vertical position. The top imaging system has a narrower field of view, roughly 2 mm across, and is the source of most of our quantitative data. All of the imaging relevant to this study was performed with light resonant with the $F = 2 \rightarrow F' = 3$ cycling transition of the D₂-line. The camera used for imaging from the top is an Andor Ixon 885 EMCCD camera (with EM, electron multiplier, turned off), while the camera used for the side is a Stingray CCD camera by Allied Vision Technologies.

One salient feature of our top imaging system is that the primary objective is set back from the vacuum chamber to allow a MOT mirror to slide in under it. This reduces the resolution of the imaging along this axis to about 8–10 μm , but has some advantages. Notably, with the MOT mirror out of the way, there is space to place a removable pick-off mirror that places the focus of the imaging system outside the chamber. The pick-off mirror is mounted on a Thorlabs KB1X1 kinematic base, allowing it be replaced with a high degree of accuracy. This has proved very useful in aligning optical paths that enter or exit the chamber from top.

If the imaging system is well aligned to the atoms, optical targets or cameras can be accurately placed in the corresponding image plane outside of the chamber by moving them into the focus of the imaging system. Once a target is so positioned, the target can be used as a proxy for the atoms in future alignments. This allows us to change aspects of our imaging system quite quickly. For example, once a target is placed, and an image of it is saved with the position of the atoms noted, lenses could be reconfigured to change the magnification and subsequently realigned to the target in *real time*. Aligning to the atoms directly might involve a lengthy search through the alignment parameter space with one observation every




Figure 5.10: Optical trap and imaging system. The ODT axis, side imaging axis, and top imaging axis are oriented at right angles to each other. The focal position of the top imaging system can be scanned by translating a lens on a translation stage, though the change in magnification with focal position must be accounted for. A beam pick off point is a useful feature of this system. Several features of the imaging system are not used in most of this work described in this chapter, including an intermediate focal plane for a razor blade and a cube for overlaying additional trapping or manipulation beams from the top.

twenty seconds. Alternatively, if a camera is placed at the secondary image plane, top-down beams can be readily aligned and focused on the “atoms.”

It is not always trivial to find the focal plane in which the target should be placed, though; one needs something sharp to focus on, and whatever this thing is should be illuminated with light near the imaging wavelength of 780 nm. Optical test patterns with features at many different length scales, back lit by a diffused 780 nm laser, can be roughly placed by hand, and brought into fine alignment with a 2D or 3D translation stage. Cameras are more difficult to align, as the sensor surface is typically very smooth and reflective, leaving few features to focus on. In this case, we roughly find the imaging focus by illuminating a piece of paper with small type with diffuse laser light. Type on the paper should reveal itself when the paper is near the appropriate position, and the combination of letters seen can serve to indicate the portion of the paper that is being imaged. With this location noted, the camera sensor is positioned to roughly this spot and likewise illuminated. At this point, one hopes that there is dust on the sensor that can be discerned and brought into focus. Some sensors

also have lines that can be discerned if illuminated properly.

5.5.1 Focusing the imaging system

We focus our imaging system *in situ* using the procedure described in Ch. 4 of Guzman [24]. In short, this procedure involves imprinting a spin winding of known wave vector k onto the BEC by applying a magnetic field gradient to a condensate with a coherent transverse magnetization (or transverse pseudo-spin). We then use our spin-selective imaging (described in Refs. [49, 47]) to image a single component of the spin winding, which should have features with 100% contrast at the specified k . By noting the contrast of these features in our images for different values of k , we determine the modulation transfer function of our imaging system. We can then adjust lenses in our imaging system to optimize the modulation transfer function.

It is important to note that it is not sufficient to bring a single sharp feature, such as a narrow BEC, or a single wave vector into focus. In a real imaging system with aberrations, one can bring some length scales into greater focus at the expense of others. In addition, the Talbot effect will cause features with a narrow range of wave vectors to appear with high contrast when the imaging system is focused on any of several planes aside from the plane of the condensate. A rough focus can be found by looking at long length scale features where the Talbot effect can be mostly ignored owing to the finite extent of the condensate.

In many of our experiments we image the atoms after they have been released from the trap. In many cases, the atoms fall out of the focus of our imaging system, such as after momentum space focusing or after several milliseconds of time-of-flight. Thus, we run the MTF calibration procedure at a few times of flight. The resulting focal positions are consistent with a simple model that adjusts the lens position in accord with the expected atom position:

$$\text{lens position at time } t = \textit{in situ} \text{ lens position} - ct^2.$$

The calibrated constant $c \approx 0.064 \text{ mm/ms}^2$ incorporates the acceleration due to gravity as well as the imaging system magnification.

Once the focus is calibrated in time-of-flight, we can focus any cloud by observing its position from the side, relating this position to that of a condensate dropped in time-of-flight, and setting the focusing lens accordingly. We checked the vertical position of our atoms after MSF and adjusted the focal position employed during MSF every day (or more) to assure that the images we used for thermometry were in focus.

5.5.2 Calibrating the imaging system

The images we want to end up with are three-dimensional: two in-plane spatial dimensions, along with an number coordinate. We derive the number of atoms imaged by pixel i , N_i , by combining data from three exposures.⁷ An exposure of the atoms back-lit by a light pulse

⁷ Exhaustive treatments of quantitative absorption imaging are available in Refs. [32] and [49].

of duration t yields a number of “counts,” arbitrary units delivered by the camera’s ADC, n'_a (“a” for absorption). An exposure of empty space back-lit by a light pulse of the same duration yields n'_b counts (“b” for bright field). An image without imaging light and without atoms yields n_d (“d” for dark field). n_d is non-zero primarily because of background light in our lab, but also because of noise on the CCD and offset in the camera readout. This background is subtracted to yield $n_{b/a} = n'_{b/a} - n_d$, and the number of atoms is

$$N_i = \frac{A}{\sigma_0} \log \frac{n_b}{n_a} + \frac{2}{\Gamma t q g} (n_b - n_a), \quad (5.13)$$

where $A = (\text{pixel size/magnification})^2$ is the real area imaged by the pixel, σ_0 is the resonant cross-section for the imaging light on the atomic transition, Γ is the line-width of the transition, q is the quantum efficiency of the CCD (photon to photo-electron conversion factor), and g is the “gain” of the camera readout (photo-electron to ADC count conversion factor). The first term dominates at low intensities, $I \ll I_{\text{sat}}$, and represents a regime in which the atoms scatter a certain *fraction* of photons that pass through them. The second term dominates at high intensities and represents a regime in which the atoms are saturated by the imaging light and scatter a fixed *number* of photons at a rate $\Gamma/2$.

Our side imaging system operates in the regime $I \ll I_{\text{sat}}$ (Sec. 5.5.6), and thus only the first term is relevant. Improper calibration will lead to errors in overall atom counting. Our top imaging system operates with $I \sim I_{\text{sat}}$, and thus both terms need to be calibrated. In this case, because the intensity of the imaging light varies across our field of view, failure to properly calibrate both terms in Eq. (5.13) can lead to errors in the column density that *vary with position*, which can bias the temperature measurements. Fortunately, in the end, we have a powerful way to check that our system is working: we can verify that our atom counting is independent of imaging intensity, and that our atom number counting is consistent between the two independent imaging paths.

5.5.3 Magnification

Calibrating the magnification from the side is quite straightforward. We release a cloud of atoms from the trap and watch it fall. Since gravity acts in the plane of our imaging, we can be sure that, neglecting magnetic or electric effects, the position of the cloud’s center-of-mass will change in accordance with a well-known uniform acceleration. Vertical magnetic gradients can skew this measurement and are difficult for us to measure in our setup owing to our trap geometry. By dropping clouds composed of $m_F = -1$, $m_F = 0$, and $m_F = 1$ atoms, we can deduce the presence of gradients that act with force greater than about 1% that of gravity. In our system, no such gradients were detected.

To calibrate the magnification from the top, we employ an optical micrometer with 25 μm divisions placed in the pick-off path. The magnification changes slightly with focal position, so the apparent size of the micrometer needs to be measured in a range of focal positions, as shown in Fig. 5.11b. The fractional change in length scale over the range of focal positions

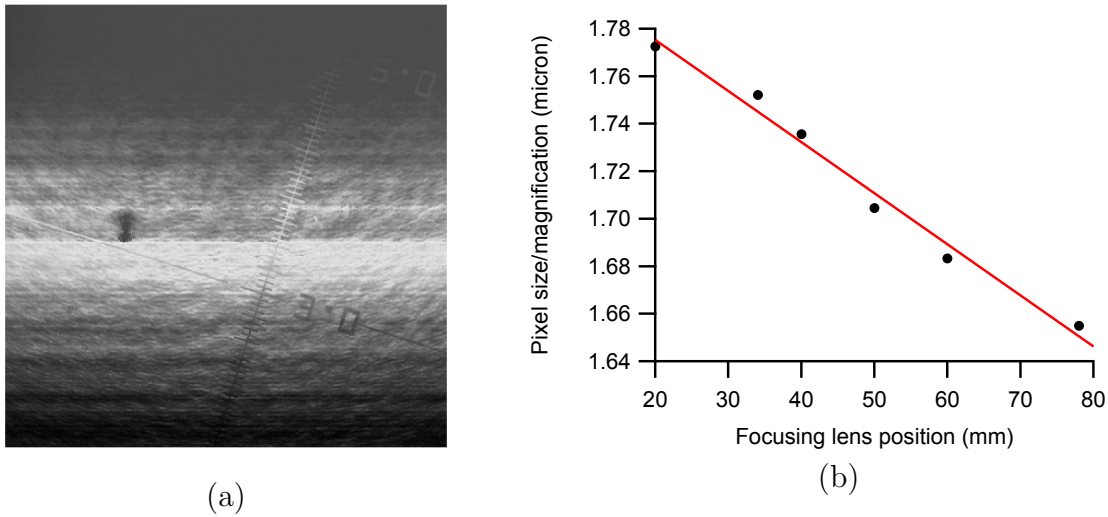


Figure 5.11: Calibrating the top imaging magnification. At several positions of the focusing lens we bring the micrometer into focus by translating it and measure the apparent size of the pixels on the CCD.

is around 7%, which, if unaccounted for, would lead to errors in inferred temperatures and atom numbers of up to 14%.

5.5.4 Gain and efficiency

The gain-efficiency product qg in Eq. (5.13) represents the conversion rate of photons in the image plane to A/D counts in the digital output of the camera. Gain and efficiency have different effects on the noise of the imaging system—a fact employed by camera manufacturers to determine the gain of the ADC for individually calibrated cameras—but for the purposes of atom counting the product is important.

The procedure we use is straightforward. First, we focus or iris the imaging beam so that it is small enough to fit on the measurement surface of an optical power meter near the atoms—before and/or after the vacuum viewports—and on the camera sensor. The inferred optical power at the atoms can be turned into a photon flux and compared to the rate at which counts accumulate on the camera. The resulting value of qg should be consistent with the value estimated from the camera’s specified gain and approximate quantum efficiency, combined with reasonable optical losses between the atoms and the camera, as it is in our case.

5.5.5 Polarization, cross-section, and optical pumping

The D_1 and D_2 lines in alkali atoms are very well characterized. In the case of ^{87}Rb , the relevant numbers are all contained in Ref. [82]. The main thing for us to do is to make

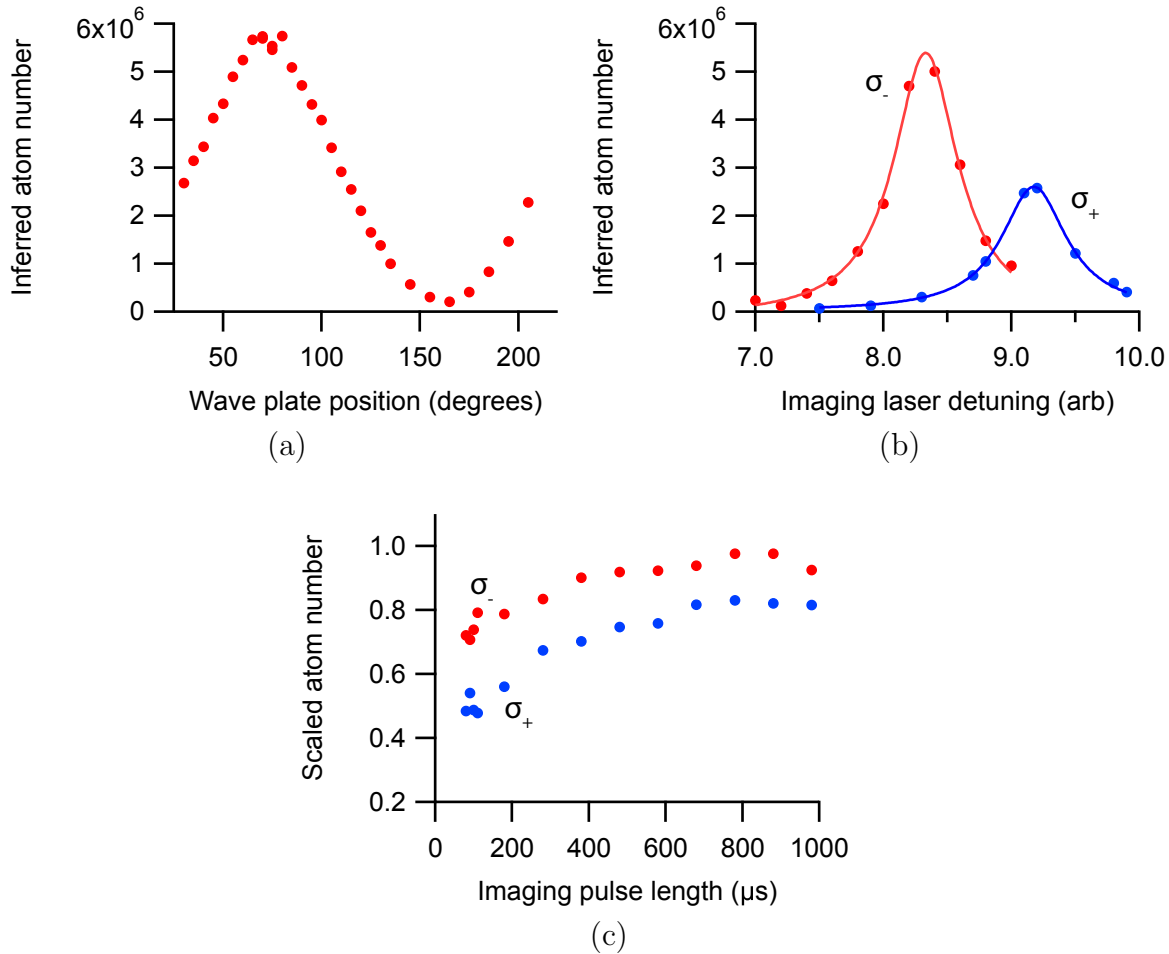


Figure 5.12: Side imaging light polarization and pulse duration. (a) Having the right imaging light polarization is important to accurate atom counting. Fortunately, calibrating the polarization is straight forward. (b) The maximum atom signal for opposite circular polarizations differ as a consequence of slow optical pumping by the low-intensity imaging light. (c) In this case, accurate atom numbers require long imaging pulses. Pulses can be made shorter by increasing the pulse intensity as described in the text. Atom number is scaled by the maximum atom number.

sure that our atoms and light are polarized in accordance with the assumptions implicit in the cross-section and line-width we are using. In our case, we assume that the atoms are optically pumped into the stretched state (with respect to the quantization axis set by the imaging light, to which we also align a small applied magnetic field to simplify optical pumping), and that the light is circularly polarized.

In alkali atoms, the strongest transition (largest cross-section) is the stretch-state transition with the appropriate circularly polarized light. Thus, as shown in Fig. 5.12a for the side

imaging system, we can verify that the imaging light is polarized properly by maximizing the atom signal with respect to the optical polarization, controlled in our case by the orientation of a $\lambda/4$ -plate following a polarizing cube.

The difference in the maximum signal for atoms initially (before any repump is applied) in the $|F = 1, m_F = -1\rangle$ and $|F = 1, m_F = +1\rangle$ states is due a breakdown in our assumption that the atoms are quickly pumped into the stretch state and scatter most of the photons that contribute to our signal from this state. In the case of the side imaging system, with light intensity far below the saturation intensity, each atom only scatters a handful of photons, and the rate at which atoms scatter on the weaker optical pumping transitions, as well as their initial Zeeman sublevel, is important. We can reduce the impact of the uncontrolled scattering rate during optical pumping by making the atoms scatter more photons either by using longer imaging pulses or more intense imaging light. In the case of the side imaging system, we place a ND filter in the imaging path *after* the atoms to avoid saturating the camera. We do not observe any problems with optical pumping in the top imaging setup, which can use more intense light without saturating the camera sensor, owing to the higher magnification.

5.5.6 Pulse duration and intensity

There are both practical and fundamental considerations involved in choosing an imaging pulse duration and intensity. For our purposes, the following factors were considered most important.

Signal to noise Probably the most fundamental consideration is signal to noise. Appendix A of Marti [49] includes a discussion of signal and noise in resonant absorption imaging, ultimately concluding that, everything else being constant and in the limit of low optical density ($n_a \sim n_b$), the noise is minimized at intensity $I = I_{\text{sat}}$, with $\delta N \geq \sqrt{16A/(q\sigma_0\Gamma\tau)}$. Equality in this relation is achieved in a system where photon shot noise is the dominant source of noise. By this consideration alone, the pulse length τ should be maximized. The condition $n_a \sim n_b$ may seem to make this analysis moot—don't we care most about noise when and where atoms are present?—but it doesn't. In optically dense regions, the dominant noise source is likely to be *atom* shot noise, rather than photon shot noise.

Laser detuning drifts Imaging near saturation, as already mentioned in Sec. 5.5.2, means that the calibration of the imaging system is important to avoid errors that are correlated with light intensity. The validity of the calibration, however, depends on the imaging light being on resonance. We found that in our system, our laser lock set point could drift on the order of 1–2 MHz, enough to noticeably change the apparent number of atoms and, more importantly, their spatial distribution. Imaging at lower powers, perhaps with longer pulses to increase the signal to noise, can reduce sensitivity to




Figure 5.13: The inferred atom column density (left) and the repump beam profile (right). Red circles are for spatial reference and mark the same area. The wings of the atom column density should be smooth. Instead, they have structure that correlates with the repump intensity. Using microwaves to repump removes these features. We note that the repump and imaging beam have the same spatial structure, and in some cases the imaging beam can be the source of such features.

these drifts⁸. Regardless, during data acquisition, we had to take care to check the lock drift regularly and recenter it around zero as necessary.

Camera dynamic range Longer laser pulses with more photons increase the signal to noise, but the camera can only count so many photons per pixel. To have reliable images, it's important the none of the camera pixels in the region of interest are saturated by any of the imaging pulses.

Atom motion Each time an atom scatters a photon, its momentum changes. On average, scattering N photons leaves an atom with momentum $p_{\parallel} = N\hbar k$ momentum along the direction of the imaging laser, and $p_{\perp} = \sqrt{N}\hbar k$ transverse to that direction. The longitudinal motion and Doppler shift limit the time the atoms remain in focus and the number of photons that can be scattered, respectively. The transverse motion causes the image to blur. In imaging small sharp features, such as the edge of a Thomas-Fermi profile, it's important to verify that imaging pulses are sufficiently short to capture the desired level of detail. In our system, we found that changing from imaging pulses of $200 \mu\text{s}$ at $I \sim I_{\text{sat}}/4$ to $40 \mu\text{s}$ at $I \sim I_{\text{sat}}$ was important to bring the calculated Thomas-Fermi radii in accord with the one observed in *in situ* images of the condensed gas, per Eq. (5.5).

⁸At very high powers, above saturation intensity, power broadening may also help alleviate the effect of drifts, though we have not explored this option.

5.5.7 Dealing with high OD

In our work with phonons, described in our publication (Ref. [46]) and Ch. 3 of Marti [49], we imaged optically dense samples by repumping a fraction of the total number of atoms using off-resonant repumping light. Off-resonant, not merely low intensity, light was key to this scheme, because it guaranteed that the gas appear nearly transparent to the repumping light, allowing for the gas to be uniformly excited. The excitation, however, can only be made as uniform as the excitation beam. In the phonon work, the gas was sufficiently small that this condition was well met. In the present work, wherein we image thermal magnons in momentum space, the clouds can cover a substantial portion of the repumping beam, and unless the cloud is fully excited to the $F = 2$ state, which in the case of high OD may not be desired, the spatial profile of the repump impacts the apparent spatial profile of the atoms. On longer length scales, this can bias temperature measurements. At shorter length scales, fringes on the repumping beam add unwanted noise, as in Fig. 5.13.

In this work, we employed a microwave repump instead of an optical repump, even when imaging spin-polarized samples for which the state selectivity of the microwaves is not needed. In free space, the microwave field is uniform at scales smaller than its wavelength ($\lambda \sim 4.5$ cm) and even on resonance interacts weakly with the gas, such that even the largest condensates have negligible optical depth and are excited uniformly. There are some downsides to using microwaves for repumping the atoms before imaging that should be considered. Chief among these are:

- The microwave excitation is sensitive to the magnetic field direction and strength, so inhomogeneous magnetic fields can cause inhomogeneous excitations. Thus, magnetic fields need to be well characterized and controlled.
- The absence of a dark state makes fully repumping a sample less automatic and less robust.
- The time it takes to repump is typically longer with microwaves (100-400 μ s in our system) than with resonant light (< 50 μ s in our top imaging setup).

5.5.8 Tracking photons

Eq. (5.13) relies on an assumption that all of the counts present at a pixel correspond to photons that were present in the appropriate area of the image plane, converted via a stochastic process with the quantum efficiency of the camera and nearly perfectly by the gain of the ADC. This assumption can be violated in a number of ways. Relevant to this work were variations in dark counts, vignetting, and light scattering between the atoms and the camera.

Dark counts

Dark counts are generally well dealt with in cold atom experiments, in which frames typically have thousands of photo-electrons, by subtracting the dark image. This process is not perfect though. The dark levels in each frame have fundamental variations from shot-to-shot which means that dark counts cannot be perfectly subtracted. This is a very important consideration in applications with long exposures and low light levels. In our system, photon shot noise should dominate this noise source. However, beyond the fundamental issues, we discovered that our camera seemed to have considerably more dark counts in the first frame acquired (the one with atoms) than in subsequent frames, and that the dark counts were not uniformly distributed, leading to a non-uniform bias in the distribution of atoms. We're not sure where these dark counts came from—they did not seem to be related to the room lights—but we were able to get rid of them by reading out a throwaway frame before the typical imaging sequence. When it is not acquiring, our camera is supposed to be continually flushing the CCD so that dark counts don't accumulate. Evidently in our camera this process was not ironclad. Fortunately, the solution was straightforward.

Vignetting

In vignetting, well described in Ref. [40] as it pertains to fluorescence imaging, the effective aperture varies with position in the image plane. In fluorescence imaging the problem here is obvious, as the solid angle over which photons are collected, which relates to the aperture size and position, is a key parameter in estimating the number of scattered photons and hence the number of scatterers. If the aperture varies spatially, getting smaller at the edge of the image, the inferred spatial distribution of scatterers will be skewed.

This effect can also be present in absorption imaging. Recall that according to the optical theorem, the total cross section of a scatterer is related to the forward scattering amplitude, $\sigma_{\text{tot}} \propto \Im(f(0))$. The $k = 0$ component of the atoms' shadow has all of the information on their total cross section, or number. Non-zero k components of the shadow, or forward scattered light, contain the spatial information and are key to allowing the imaging system to reconstruct how atoms are distributed. Young students often wonder why the imaging lenses need to be so much larger than the imaging beam; this is why. The solid-angle over which forward-scattered light is captured, or equivalently the numerical aperture, is an important parameter in the resolution of the imaging system regardless of whether images are captured with absorption or fluorescence. If vignetting is present, even though the total atom number information may be intact, the inferred spatial distribution of atoms can be skewed.

Light scattering after atoms

Light scattering between the atoms and the camera has two main effects. If the scattered photons leave the imaging system, the scattering effectively reduces the quantum efficiency of the camera, perhaps in a spatially dependent way. At low light levels, where the calibration of the camera gain and efficiency is not important, this has only minor impact on the noise. At




Figure 5.14: Saturation of the OD. Scattered light after atoms limits the maximum OD that can be inferred and can lead to under counting atoms in dense areas. In this image (left frame), the condensate OD is given a “flat top.” (Right frame) A slice through the image (red) is compared to a harmonic guide to the eye (dashed gray line).

higher light levels this could lead to spatially dependent calibration errors and, in principle, bias quantitative analysis if the errors are correlated in some way.

More onerous in our system was light that scattered after the atoms and remained within the imaging system, leading to counts at a pixel that cannot be blocked by the atoms imaged by that pixel. This effect is contained within the description of image analysis in the appendix A. to Ref. [32], in which the first term of our Eq. (5.13) is recast as, with our labeling,

$$\frac{n_a}{n_b} = e^{-N_i\sigma_0/A} + s_i (1 - e^{-N_i\sigma_0/A}), \quad (5.14)$$

where s_i is the fraction of photons at pixel i that are present owing to scattering after the image plane, and is presumed to be small. The quantity $N_i\sigma_0/A$ is the optical density. With $s_i = 0$ and $I \ll I_{\text{sat}}$, this equation is equivalent to Eq. (5.13). If one calculates the optical density in the usual way as $-\log(n_a/n_b)$, the effect of non-zero s_i is manifest in a few important ways. First, the maximum optical depth that can be inferred is $-\log(s_i)$, reducing the dynamic range. Moreover, the number of atoms in each pixel will be underestimated with the degree of error getting more severe when there are more atoms and also when light intensity is lower. This can lead, for example, to peaked atom distributions getting flattened, as shown in Fig. 5.14, which is a problem if one is trying to measure temperatures accurately.

In principle, one could measure the s_i and take them into account in calculating N_i . We were able to reduce s_i to a suitable level by cleaning the optics in the imaging path. In particular, dust accumulates on horizontal optics, such as the chamber’s top viewport, and needs to be removed from time to time. We use a shop vac for this purpose. At the time the data of Fig. 5.14 were taken, these optics had not been cleaned in several years.




Figure 5.15: A condensate in the new, more harmonic, trap. Condensate profile before had clear non-harmonic structure, as seen in Fig. 5.14 where the condensate outline is not round. The condensate in the new trap, shown here, has a harmonic profile, but narrow directions are broadened slightly by our imaging resolution and, if we are not careful, imaging pulse duration.

5.6 Characterizing the optical trap

With imaging and MSF well characterized and calibrated, we begin to have confidence in the momentum distributions and atom numbers we measure. However, for the pieces to fit together in Eqs. (5.4) and (5.5), it's also important to have a well characterized trap. For our trap, this means measuring the harmonic trap frequencies as a function of the optical power and verifying that the trap is, in fact, sufficiently harmonic.

For a sufficiently cold condensate to which the Thomas-Fermi approximation should apply, an easy first pass at characterizing the trap is to image the atom density *in situ* and to compare its profile with the one expected for a harmonic trap,

$$n(x, y) \propto \left(\mu - \frac{\omega_x^2 x^2}{2} - \frac{\omega_y^2 y^2}{2} \right)^{3/2},$$

where the column density n is the integral of the full density along the z -axis. A discrepancy between the measured and expected size of the condensate could indicate anharmonicity, though it is also important to note that such a discrepancy could also be indicative of other issues, such as imaging pulses that are too long, poor imaging resolution, or an out-of-focus imaging system. Other indications of anharmonicity include heavily damped center-of-mass oscillations and oscillation frequencies that vary with amplitude.

The optical trap we used for many of the preliminary cooling and thermometry experiments had significant anharmonicity, probably owing to the exceptionally weak focus along the x -axis coupled to the exceptionally strong focus along the z -axis. The anharmonicity can even be seen by eye in the shape of the condensate shown in Fig. 5.14. For such a trap, the trap minimum is determined by the position of the z -focus, and the position of the beam focus along x is difficult to determine. Ideally, the lens responsible for forming the ODT transforms short length-scale structure in the collimated trapping beam into long length-scale structure at the beam focus, and vice versa. Thus, ideally, the atoms, confined near the center of the focused beam, are not exposed to the short length-scale crud that is present on the collimated laser beam. However, if the focus is poorly determined or too weak, the desired short and long length-scale decomposition might not occur. Additionally, the ratio

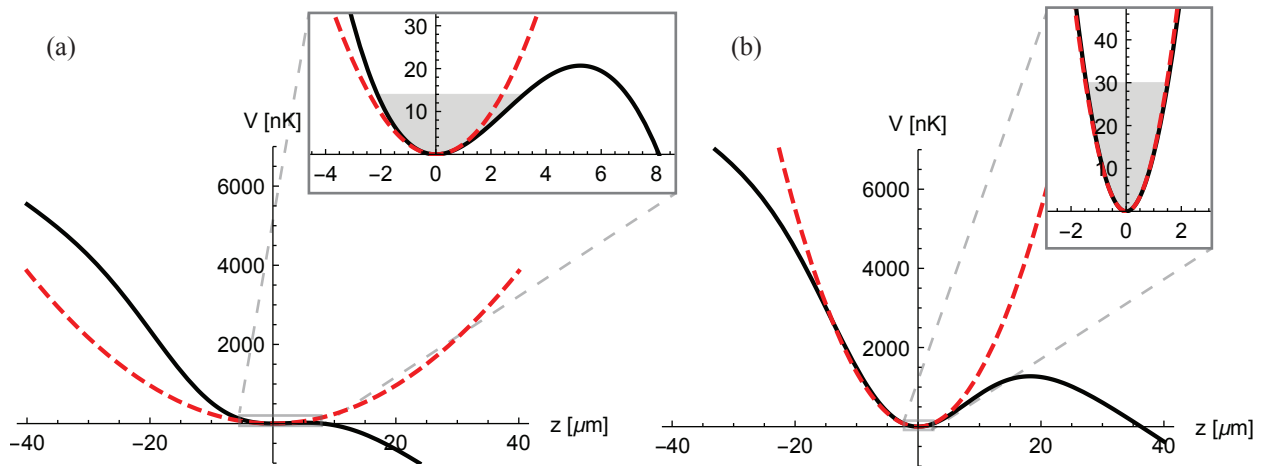


Figure 5.16: Cross section through the center of combined optical/gravitational potential at (a) low trap depth and (b) higher trap depth along the vertical axis. Dashed red line shows the harmonic potential corresponding to the curvature at the trap minimum. Insets show region near the trap minimum. Gray patch highlights a typical chemical potential for the given trap parameters, showing extent of a BEC in our experiments. When the chemical potential is a significant fraction of the trap depth, the condensate extends beyond the harmonic region at the very center of the trapping potential, both with gravity (shown) and without (not shown).

of trap depth to peak optical intensity increases as the z -focus becomes tighter, which means that at constant trap depth and temperature, the extent of the transverse optical potential explored by the atoms increases as the z -focus becomes tighter, possibly exposing the atoms to the more aberrant parts of the beam.

Characterizing such a trap is more complicated than just measuring trap frequencies, so we swapped out some lenses in our system to make the z -focus less aggressive and the x -focus more so. The resulting trap is long and narrow (and still quite thin vertically), and much more harmonic, as indicated by the condensate Thomas-Fermi profile and in-trap oscillations. A condensate in the modified trap is shown in Fig. 5.15. Importantly, in the updated trap, Eqs. (5.4) and (5.5) are satisfied. Having a well-calibrated trap potential is also very useful in calibrating the trap depth, as we shall see.

5.7 Measuring the trap depth

We infer the trap depth using a model of the trapping potential, including the effects of gravitational and optical forces, wherein beam parameters are set to values consistent with observed trapping frequencies and measurements of the trapping light.

The model for the trap includes the optical potential of a Gaussian focus and the linear

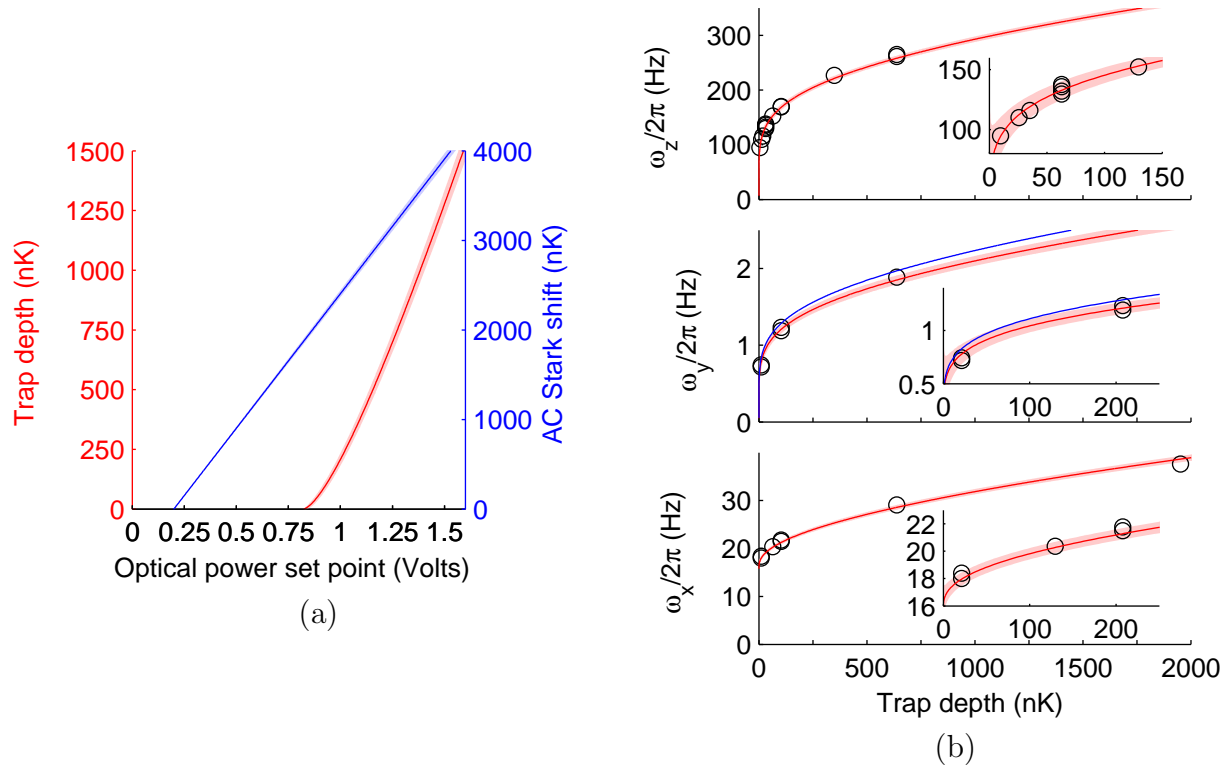


Figure 5.17: Fitting trap parameters. We use the measured trap frequencies to fit the peak AC stark shift, beam waists, and Rayleigh length. The resulting model predictions are shown in red for (a) the trap depth and (b) trap frequencies, where the light red patch is the error in the estimate implied by the fit. The blue line in the central panel of (b) is the predicted ω_y based on the Rayleigh range implied by the beam waist w_z .

potential of gravity

$$U_{\text{trap}} = \frac{-U_0}{\sqrt{1 + (y/y_{Rz})^2} \sqrt{1 + (y/y_{Rx})^2}} \exp\left(\frac{-2z^2}{w_z^2 (1 + (y/y_{Rz})^2)} + \frac{-2x^2}{w_x^2 (1 + (y/y_{Rx})^2)}\right) - mgz, \quad (5.15)$$

where the trapping laser beam propagates along the y -axis, U_0 is the maximum optical potential depth (peak AC Stark shift), and w_z (w_x) and y_{Rz} (y_{Rx}) are the beam waist and Rayleigh length, respectively, along the z -axis (x -axis). For an ideal Gaussian focus, $y_{Rz} = \pi w_z^2/\lambda$. For our trap, $w_z \ll w_x$ and thus a good approximation is to take $y_{Rx} \rightarrow \infty$.

Including the effects of gravity, the trapping potential no longer has a global minimum, as shown in Fig. 5.16, and at low enough power does not even have a local minimum. When a local minimum exists, the trap frequencies are related to the curvature of the trap potential

Parameter	Fit or inferred value	Measured value
Beam waist, w_x	210 μm	$\sim 200 \mu\text{m}$
Beam waist, w_z	22 μm	$\sim 25 \mu\text{m}$
Peak AC Stark shift, U_0	-3000 nK/V	
Optical power	145 mW/V (at atoms)	133 mW/V (after chamber)

Table 5.2: Trap parameters. The beam waists, w_x and w_z , along with peak AC Stark shift, U_0 , are fit to the Gaussian beam model with gravity, using the measured trap frequencies. From these fits, the optical power is derived from the expected AC stark shift, $152 \mu\text{K}/(\text{mW}/\mu\text{m}^2)$.

at the local minimum, which is shifted downward by gravity,

$$m\omega_i^2 = \left. \frac{d^2 U_{\text{trap}}}{dx_i^2} \right|_{x=0, y=0, z=z_{\text{min}}}, \quad (5.16)$$

and the trap depth is just the difference between the potential at the local maximum and the local minimum,

$$U = U_{\text{trap}}|_{x=0, y=0, z=z_{\text{max}}} - U_{\text{trap}}|_{x=0, y=0, z=z_{\text{min}}}. \quad (5.17)$$

The main parameters of our model are the beam waists and optical potential depth, $U_0(V)$, as a function of the optical power set point (in volts). The inputs to our model are the measured trap frequencies along all three axes, $\{\omega_i(V)\}$, where $i \in \{x, y, z\}$, as a function of the optical power set point.

The vertical trap frequencies depend on the vertical beam waist and optical potential depth only, so we first fit for w_z and $U_0(V)$ using the measured $\{\omega_z\}$. Next, we fit for w_x using the measured $\{\omega_x\}$, holding w_z and $U_0(V)$ constant. In principle, this determines y_{Rz} , however we can also choose to fit for y_{Rz} using the measured $\{\omega_y\}$; the difference is minor. Fig. 5.17 shows the trap frequency data and resulting fits. The middle panel of Fig. 5.17b shows the fit-free estimate of ω_y in blue, while the estimate with y_{Rz} as a free-parameter is shown in red. In estimating trap depth and trap frequencies, we use the model with fitted y_{Rz} , though the difference is negligible, especially at low trap depths.

As a check of consistency, we compare our inferred trap parameters to what we expect based on our knowledge of atomic physics and what we can measure about the trapping beam. Tab. 5.2 compares fitted/inferred and calculated/measured values. It's difficult to make precise measurements of the trapping laser beam before it enters the chamber, as it is larger than any of our cameras in the vertical direction in order to make a tight focus, and optical access is limited. Still, the beam waists we infer are consistent from what we expect from rough measurements of the beam size and focal length.

We measured the optical power as a function of set point voltage after the chamber, and we found the relationship to be stable and linear. The measured slope varied by about 2% over many months (within our measurement error), and the zero-power voltage was fixed

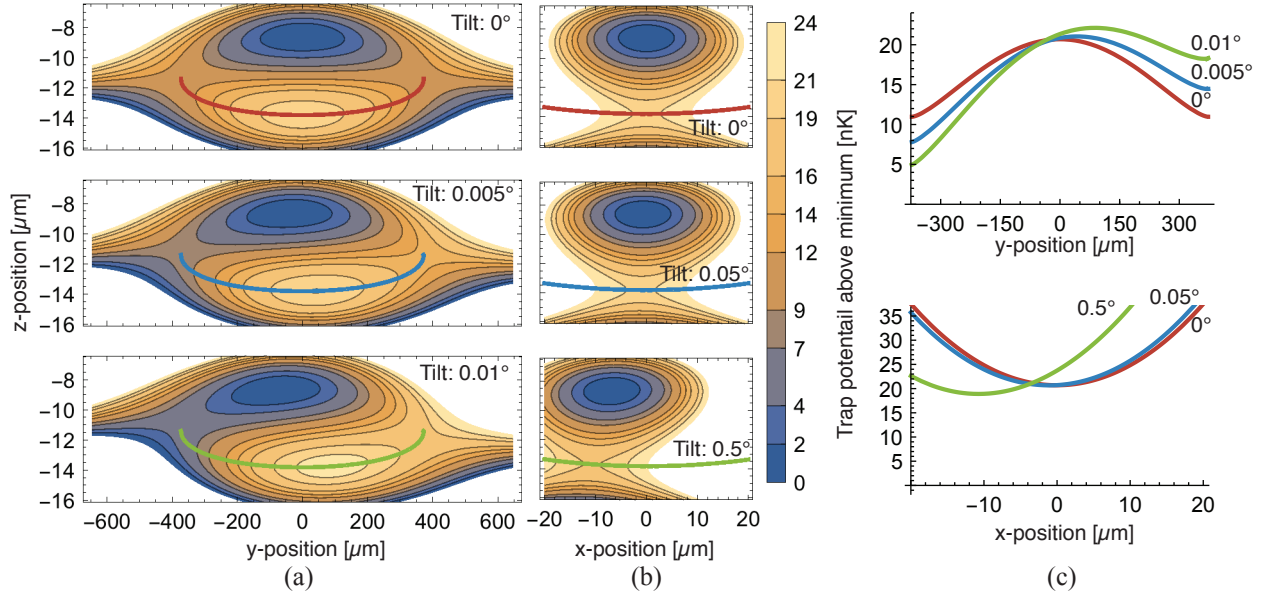


Figure 5.18: (a,b) Contour plots show the combined gravitational and optical potential that confines the atoms at very low trap depth. The optical focus is at the origin, but the trap minimum is offset, primarily along the z -axis. Minimum trap depth and center position are sensitive to trap tilt along both y - (a) and x -axes (b), but much more so along y , where the trap local minimum disappears entirely at a few hundredths of a degree tilt. The trap potential along the lip of the trap (red, blue, and green lines in [a] and [b]) is shown in (c).

by our transimpedance amplifier. The measured slope was about 10% less than the value inferred from the trap frequencies and the expected AC stark shift of the 1064 nm laser on ^{87}Rb in its ground state. The discrepancy is likely accounted for by additional losses as the beam exits the chamber, the difficulty of capturing the full beam power on a meter (owing to the large beam size), and the calibration of our high-power optical power meter.

5.8 Leveling the ODT

Our procedure for leveling the ODT *in situ*⁹ takes advantage of the extreme sensitivity of the trap depth and the position of the trap minimum to tilts, as shown in Fig. 5.18. The optical table on which our experiment rests floats on air cushions, and the precise height and tilt of the table as a whole can be adjusted. The tilt can also drift over time, so this procedure had to be repeated at regular intervals.

First, we choose a trap setting V_2 that is as low as possible, but that reliably still has a condensate present. We evaporate to V_2 and then ramp the trap back up to $V_1 \gg V_2$ over a

⁹This procedure is relevant for a trap that is already mostly aligned. Upon the initial installation, or if the trap is strongly tilted, a different procedure needs to be used.

few seconds. With these settings, we take an image of this condensate and mark its center with a persistent cross hairs in our imaging software, for future reference. With the trap at V_1 , the optical potential dominates gravity along the in-plane dimensions, so the cross hairs marks the optical trap center. Next, we evaporate to V_2 and image the condensate. If the condensate center has shifted with respect to trap center, we assume gravity is pulling it and can work out which way the trap is tilted based on how the directions in the image relate to the physical world. In our setup, conveniently, “up” is parallel, roughly, to the propagation of the trapping laser.

The trap tilt can be adjusted using optics, but if we are well aligned through the center of all the optics, there is no single knob that will level the trap without potentially making something else worse, e.g. the vertical position of the trap relative to the chamber. The floating table lets us adjust the trap tilt without changing the relative alignment of any other components of the apparatus. Ideally, the optical beam was installed very close to parallel to the table, with the table level to gravity (using a sensitive bubble level), and only small tweaks are needed. This was the case in our experiment.

The floating table has three adjustment points, one on each of three legs. The fourth leg is self-adjusting to maintain the constraint that all four corners are supported. Once we know about which axis we want to tilt the table, we need to figure out which knob to turn (i.e. which corner to adjust). There are three axes about which the table can tilt, each passing through two of the adjustment knobs. No axis runs through the self-adjusting leg. To rotate around one of these three axes, we tweak the knob at the leg that does not define the axis.

The self-adjusting leg cannot be entirely ignored. It is important that as the table is adjusted, the self-adjusting leg be able to respond by either raising or lowering. Therefore, it should not be at a limit of its range of motion. If it reaches a limit, which can be ascertained by inspecting the leg, the height of the table can be adjusted by turning all three adjustment knobs in equal measure.

Once a tweak has been made, we reinspect the condensate at V_2 . If it is centered, we move on, if not, further tweaks may be needed. If the condensate is well centered at V_2 , the trap may be further leveled at a lower trap depth V_3 using the same procedure. At very low trap depths and when the chemical potential is a large fraction of the trap depth, the number of atoms in the condensate may be a more sensitive probe of the trap alignment. Thus, to reach a final precise alignment we optimize the tilt of the trap (along both x and y) for condensate number.

Fig. 5.18 reveals something interesting about the trap at low depth/high chemical potential: the trap depth varies substantially as a function of the in-plane position. Along the y -axis, the lip of the trap is lowest at the edges, where the density is lowest. The opposite is true along x . In our work, including our published write-up of this work [62], we have typically reported the trap depth as would be measured vertically through the center of the gas (at $x = y = 0$). However, the minimum energy required for an atom to leave the trap may be substantially (as a fraction) less than the trap depth defined in that manner. The average energy of an atom lost is probably somewhere in between.




Figure 5.19: Calibrating the circularly polarized trap. (a) We use the observed trap frequencies of atoms in the $|m_F=1\rangle$ (red) and $|m_F=-1\rangle$ (blue) states to relate voltage control levels of the circularly polarized trap, and thus the AC Stark shifts, to the well-calibrated linear trap. The fractional difference in the slopes for ± 1 atoms is 0.025, close to the expected value of $2 \times 0.011 = 0.022$. (b) The resulting trap depths are shown, with the lowest control voltage setting used in this work indicated by the vertical dashed gray line. Dashed black line is the trap depth for atoms in the $|m_F=0\rangle$ state. Light colored patches show the systematic error carried over from the (linear) trap depth calibration. Dark colored patches are an additional (independent, essentially negligible) systematic error from the calibration in (a).

5.9 Setting up a state-dependent trap

Most of our experiments are done with linearly polarized trap light, which couples equally to the atoms regardless of their Zeeman state (ignoring the small tensor Stark shift). For some experiments, however, we will want to vary the trap depth for atoms in different states. One way to achieve this is with magnetic field gradients. However, our apparatus lacks the ability to control gradients and bias fields independently, possibly leading to complications. Another approach is to use circularly polarized trap light, which we calculate can apply a differential Stark shift of about 1.1% between atoms in the $|m_F=0\rangle$ and $|m_F=\pm 1\rangle$ states if the spin quantization axis of the atoms (defined by a bias magnetic field) aligns with the quantization axis of the light. In our case, the bias magnetic field and optical trap axes are nearly aligned.

The differential Stark effect, measured to be about 1.2% in our system, is small compared to what can be achieved with magnetic field gradients, however a small fractional change in Stark shift can translate into a large fractional change in trap depth when the trap is shallow, around 30% at our lowest trap depths, as shown in Fig. 5.19. With the same circular polarization, and all other experimental parameters fixed, we can change the trap

potential of the majority atoms, by using either $|m_F=1\rangle$ or $|m_F=-1\rangle$ atoms, while keeping the trapping potential for the $|m_F=0\rangle$ atoms constant.

Changing the trap polarization from linear to circular is straightforward, but with the change in polarization, the relations between ODT control voltage, optical power, and AC Stark shift are no longer known. However, a single scaling parameter captures the effects of all of these uncertainties, as the trap geometry is unaltered. Thus, to calibrate the circularly polarized trap, we measure the trap frequency for atoms in the $|m_F=\pm 1\rangle$ states at two different control voltages.

5.10 Noise performance and trap depth limits

The optical trap level is maintained by a home built PI controller. The setpoint is established by the 13-bit DAC of our National Instruments PXI-6723 control card. With outputs ranging from -10 V to $+10$ V, this card gives us 2.4 mV resolution in defining the set point. The trap level is measured by picking off a fraction of the trapping light after it passes through the chamber and focusing it onto a Hamamatsu S11499 photodiode connected to a SRS SR570 current preamplifier. The trap level is adjusted by varying the diffracted optical power through an AOM via a variable RF attenuator that adjusts the driving power.

The SRS preamplifier is an expensive solution to a simple problem, but it allows us to easily tune various parameters—bias voltage, transimpedance gain, and bandwidth—for optimal stability, noise, and dynamic range. We examined the performance of our feedback loop by measuring the noise on an out-of-loop photodiode with and without feedback. We found that very few of the knobs we could tune had much impact on the noise, as measured with a fast FFT spectrum analyzer. The largest improvement to our setup was replacing an older Thorlabs photodiode with the Hamamatsu, which is optimized for the infrared. Ultimately, integral gain and transimpedance gain were set to give the largest possible dynamic range (trap levels are varied from 7–8 V to less than 1 V in a typical experiment) and to make the trap level stable over its entire range of operation.

This latter requirement turns out to dictate a level of feedback gain that is just sufficient for the trap to reach its highest setpoint¹⁰. The open loop response of the optical trap level is non-linear owing to the variable RF attenuator and AOM. If the integral gain is set too high, the optical trap power can oscillate as the level is changed within a certain window, even if it is stable in another window. We monitor the optical power on an oscilloscope to make sure it is not manifestly unstable.

In spite of the unsophisticated manner in which we optimize our optical trap, we find it to be very stable. At the lowest trap depths, the stability of the temperature we measure suggests the trap level is stable to much better than the 2.4 mV resolution of our control card, which corresponds to a nearly 20% change in trap depth at our lowest values. More-

¹⁰Recall that a real integrator can only integrate up to a finite level before saturating, and so the DC gain of the feedback loop is not, in fact, infinite, as it would be for an ideal integrator.

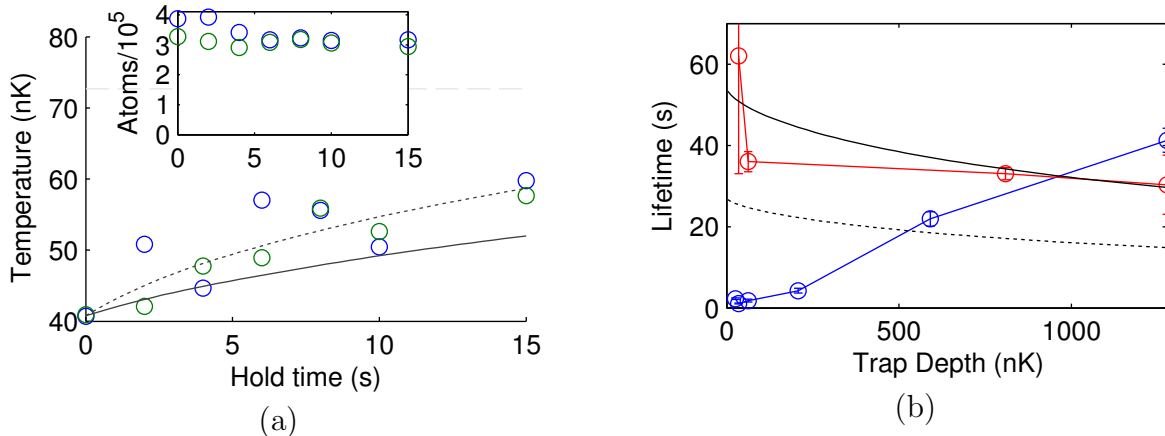


Figure 5.20: Heating rate and trap lifetime in the optical dipole trap. (a) The gas is cooled via cycled decoherence cooling from the temperature indicated by the dashed gray line and then allowed to heat. Two repetitions of the experiment are shown in green and blue, with temperatures extrapolated to zero magnon number using the slopes derived from the data of Sec. 5.19. During heating, no magnons are present. The solid line shows the expected heating based on peak scattering rate. The dotted line is based on heating at twice the expected scattering rate. (b) The lifetime of the majority (red) and magnon (blue) gas at different trap depths, along with statistical error in the measured lifetimes. Black dotted and solid line are inverse scattering rates from (a).

over, lifetime and heating rates seem to be nearly consistent (within a factor of 2) with a fundamental process: spontaneous scattering from the far-detuned trap.

Each photon scattered in the trap should deposit its recoil energy, $E_R = 97$ nK, in the gas at a rate that varies with the AC Stark shift U , ignoring counter-rotating terms,

$$R_{\text{sc}} \approx \Gamma \frac{U/\hbar}{\Delta},$$

with Γ the excited state linewidth and Δ the detuning [83, 67] of the trap light from the relevant atomic resonance. Fig. 5.20a shows the temperature of the gas as it sits in the optical trap at constant power with evaporative cooling turned off by precooling with cycled magnon decoherence cooling. The heating rate that we observe is consistent with about twice the rate expected owing to spontaneous scattering from the optical trap¹¹. This result suggests the heating rate could be improved somewhat, but it is approaching a fundamental limit given the current trap.

The source of additional heating beyond scattering from the ODT is not known. Possible sources include three-body collisions, optical intensity fluctuations, magnetic field noise, and mechanical noise. Three body decay is evident in the initial atom number decay after optical

¹¹Initially we thought the heating rate equaled the expected value, but more careful calibration caused us to correct our temperature estimates upwards, yielding the present estimate.

evaporation but is generally not seen at longer times. The number of atoms present during heating does seem to decline initially before leveling off, as shown in the inset of Fig. 5.20a.

We also measured the lifetime of the gas in the trap, which after initial three-body decay, was consistently 35-45 seconds and consistent with one-body loss. One might expect to lose one atom for each U_{trap} worth of energy deposited¹², however this is not what we observe. For example, as shown in Fig. 5.20b, at a trap depth of 1000 nK, the lifetime of roughly 40 s is consistent with losing one atom for every two scattered photons, since the solid black line corresponds to a scattering rate of half the observed value. Two scattered photons corresponds to about 200 nK, rather than the 1000 nK that the simple model predicts, once again indicating some additional heating.

Historical note: we initially performed the experiment of Fig. 5.20a in which magnons present during the hold time and found that the temperature seemed to equilibrate to a much lower value than expected. This prompted us to perform the experiments that revealed magnon-assisted evaporative cooling.

5.11 Magnetic field gradient control

In many ways, our experiments depend on reliable, repeatable control of magnetic fields. Unfortunately, as mentioned in Sec. 3.6, the chamber where we do science was not designed with exquisite control of magnetic fields in mind. Typically, in a sensitive experiment like ours, one would at minimum like to have control over all three bias fields and all five (independent) gradients. Our experiment, however, can only control one gradient via a single spherical-quadrupole-producing coil. However, as described well in Sec. 4.4.2 of Ref. [49], we are able to control the magnetic field gradients seen by our atoms by canceling one in-plane gradient and maintaining a magnetic bias field along the direction of the other in-plane gradient. Our scheme has the benefit of not requiring any additional coils beyond the ones used for our MOT, but the disadvantage that it constrains the bias fields that we can apply and couples external bias field fluctuations to gradient fluctuations.

Sensitivity to bias fields is somewhat onerous because in general, environmental bias fields fluctuate much more than gradients. This can be mitigated by operating with higher bias fields, which reduce the effect of background fluctuations on the orientation of the overall field. However, at high bias fields, the Zeeman energy increases and the RF fields required to drive transitions between Zeeman states are at a higher frequency. Unfortunately, as our (improvised) science chamber does not have an integrated internal RF antenna, we resort to coupling RF into our system from the outside, and in this configuration higher frequencies are more strongly attenuated, reducing the reliability with which they can drive magnetic transitions. High bias fields also increase the quadratic Zeeman shifts, which can be undesirable. By acquiring data when the magnetic environment is most quiet (late night on weekends when the elevators are not moving), we have been able to control gradients to better than 1 mG/cm.

¹²Atom lifetimes measured in the magnetic trap exceed several minutes.

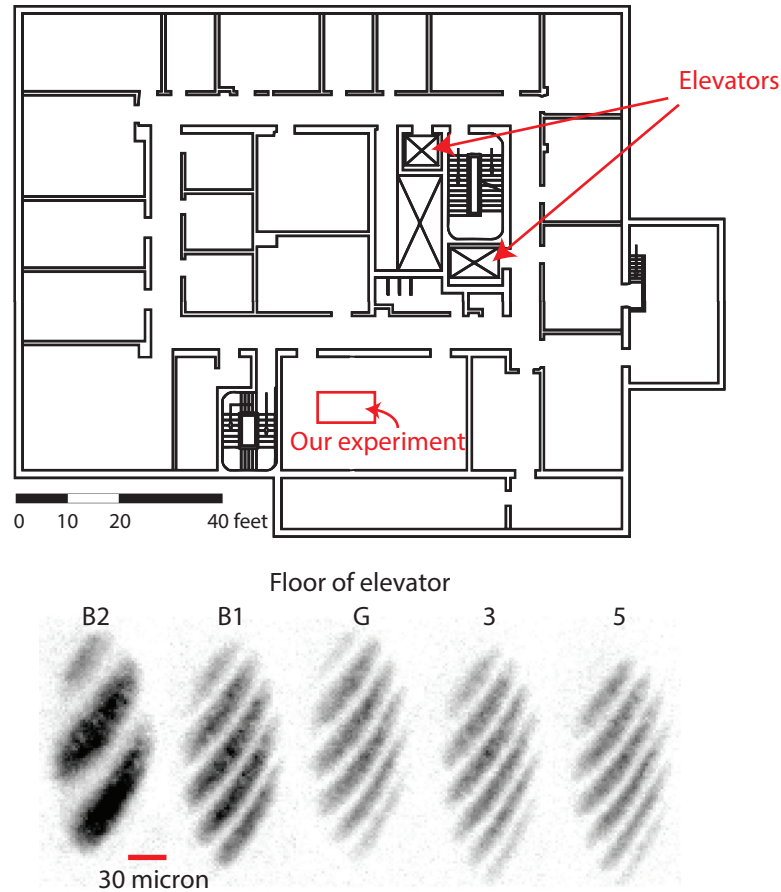


Figure 5.21: Top: The floor plan of Birge subbasement, indicating the location of our experiment table and the elevators. Bottom: Magnetic field gradients are measured by Ramsey interferometry with a 20 ms time between pulses. The magnetic field gradient present at the atoms varies from roughly 10 mG/cm to 25 mG/cm as the nearby freight elevator moves from the subbasement to the top floor.

We were able to improve on this situation by implementing active magnetic field stabilization of the slowly varying bias fields. To do this, a reference measurement is taken of the magnetic field value at a particular time in the experimental sequence where the fields under experimental control are stable in a “reference” configuration, for example during optical evaporation. The field is measured by a magnetometer placed near the atoms, but external to the chamber. Then, in subsequent runs of the experiment, the currents through large bias coils are adjusted to maintain the reference field value as long as the applied experimental fields are in the reference configuration. If the experiment requires fields to deviate from the reference configuration (e.g. to change the magnetic field gradient), the current through the bias coils is held constant at the last best value. The magnetic environment relevant to the experiment is thus held constant from run to run, even as the elevators move or faraway

superconducting magnets are slowly cycled.

To perform field stabilization, we use six coils in three pairs oriented along the axes of our optical table. The coils are large and rectangular, varying from 1–2 m across. Pairs are 1–3 m apart, with each pair controlled by a 0–10 amp, 0–25 V Kepco ATE power supply. Each coil is centered as well as possible on the location of the atoms in the chamber. Ideally, these coils would be in a Helmholtz configuration, however, as these coils were added after the chamber was constructed and were designed to minimally disturb a functioning experiment, the members of each coil pair are not equidistant from the plane of the atoms and their dimensions were allowed to deviate from the ideal. The number of windings in each coil was set such that the same common current (max 10 amp) and voltage (max 25 V) could cause each pair to generate a maximum field large enough to compensate bias field drifts up to about 0.03 Gauss with minimal magnetic field gradient at the position of the trapped atoms. Thus, opposing coils that were not equidistant from the atoms necessitated different numbers of windings. In the end, each coil was composed of 4–13 windings of 18 AWG wire. Winding the wire *in situ* was facilitated by using a single-turn of ribbon cable and joining the ends with the ribbon offset by one.

Three-axes of magnetic field control requires (at least) a three-axis measurement of the field. We achieve this by placing a Bartington Mag690 three-axis fluxgate magnetometer outside the main chamber, axes aligned with the coils. The magnetometer is close enough to the position of the atoms that the stabilization field is roughly constant in the space between them. This magnetometer can accurately measure fields up to 1 Gauss in magnitude. Thus, it's important that the reference field at the magnetometer not exceed this value.

The output of the magnetometer, ± 10 V full scale for each of three channels, is filtered by active double-pole ~ 20 Hz low-pass filters and digitized by a 16-bit Texas Instruments ADS7825 ADC. The power supplies are each controlled by a 0–1 V analog voltage, which we generate with three MAX541 16-bit DACs. In between the ADC and DACs sits an Arduino AVR-based microcontroller system, which can sample the ADC and write to the DACs at a rate of roughly 50 Hz for each channel while running a control and feedback logic loop. The Arduino is programmed to accept two digital control inputs that set it to either

- (a) hold its analog output at the most recent value;
- (b) measure a reference field and update the reference set point;
- (c) feedback to the control voltages to achieve and maintain the reference field set point;
- or
- (d) hold its analog output at a predetermined fixed value, typically zero amps.

The arduino also accepts commands over a serial interface that allow the set points to be manually adjusted and to manually tweak or override the current settings for both (a) and (d) control scenarios.

Much of our experimental sequence was optimized with no stabilization fields, so we begin our experimental sequence (MOT loading, magnetic evaporation, and initial loading




Figure 5.22: With feedback on, the bias magnetic field is more stable, and thus RF transitions are more stable. Here the populations of the three Zeeman states are shown after a $300 \mu\text{s}$ $\pi/2$ pulse, with the elevator in motion between experimental cycles.

of the optical trap) with the stabilization fields off, which corresponds to stabilization control configuration (d). After the ODT is loaded, during the initial optical evaporation and as we adjust the bias field and gradients to their reference values, we set the stabilization coils to configuration (a), which causes the stabilization field to ramp over 100 ms from zero Amperes to the current that last achieved the reference set point. Once the reference fields are stable, we enter configuration (c) to compensate for any changes to the magnetic field environment. We continue feedback until the experiment compels us to change the fields from the reference configuration, at which time we switch back to (a). Whenever the reference field configuration is changed or recalibrated, we set the stabilization currents to three Amperes. This assures that field variations can be tracked in future runs of the experiment by our stabilization coils, which are driven by unipolar power supplies. Once the reference levels are recalibrated, the reference field level is saved as the new set point by running the experiment with stabilization configuration (b) instead of (c).

With the stabilization coils integrated into the experimental sequence, bias fields are well stabilized at the one second timescale to better than 1 mG at the atoms. As a result, the applied, or rather, mostly canceled magnetic field gradients are quite stable—we can easily maintain an in-plane gradient of less than 0.5 mG/cm—from run to run and even day to day. The stability of the bias field means that RF transitions are much more stable, as shown anecdotally in Fig. 5.22.

Long RF pulses ($>500 \mu\text{s}$) are still limited by higher-frequency noise and sub-mG variations of the bias field at the atoms. The dominant source of higher-frequency noise is in phase with the line voltage and can largely be eliminated by synchronizing critical RF pulses with the 60 Hz line. Regardless, such pulses need to be recalibrated frequently.

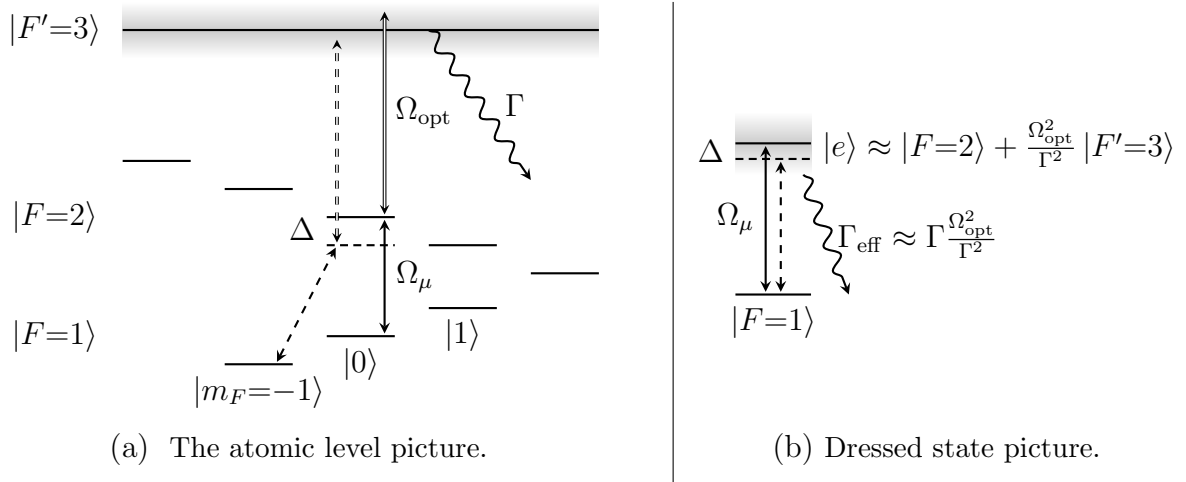


Figure 5.23: Our initial scheme for purging magnons as a level diagram (a) and viewed in a schematic dressed-state picture (b). There is a second dressed state that is unresolved in this picture, wherein the optical transition is not saturated.

5.12 Purging magnons from the trap (for cycled cooling)

Our strategy for purging magnons from the trap has evolved over time. Our usual method of purging atoms is illustrated in Fig. 5.23a. We simultaneously apply resonant microwaves to transfer atoms to the $F=2$ ground-state manifold and imaging light to cycle atoms on the $F=2 \rightarrow F'=3$ transition, causing them to accumulate sufficient momentum to leave the trap. However, in general, the dynamics of a three-level quantum system with two drives and dissipation is not trivial, and perhaps it's not surprising that this scheme for purging atoms can have some side effects and limitations.

In the limit that the optical drive is resonant and large compared to the microwave drive, but still small enough to avoid saturating the optical transition, $\Gamma \gg \Omega_{\text{opt}} \gg \Omega_{\mu}$, we can consider the $|F=2\rangle$ states to be lightly dressed by the optical field, admixing a small amount of the excited $|F'=3\rangle$ states. The resulting admixture has an effective line width and decay rate Γ_{eff} , as shown in Fig. 5.23b. The microwave drive that we apply during purge is resonant with atoms in the $|m_F=0\rangle$ state and is detuned for atoms in the $|m_F=-1\rangle$ state by $\Delta = \omega_L \approx 100$ kHz, the Larmor frequency. In this simplified scenario, where we also note $\Delta \gg \Omega_{\mu}$, the ratio of “bad” scattering to “good” scattering is

$$\frac{R_{-1}}{R_0} \approx \frac{\Gamma_{\text{eff}} \Omega_{\mu}^2 / (4\Delta^2)}{\Gamma_{\text{eff}} \Omega_{\mu}^2 / \Gamma_{\text{eff}}^2} = \frac{\Gamma_{\text{eff}}^2}{4\Delta^2}. \quad (5.18)$$

Thus, if we are willing to excite 1,000 magnons before we excite a majority atom, then we would limit $\Gamma_{\text{eff}} \lesssim 5$ kHz. This requires $\Omega_{\text{opt}} \lesssim 50$ kHz, about 1/100 of the saturation intensity (which is about 6 MHz). Conveniently, our maximum microwave Rabi frequency is roughly

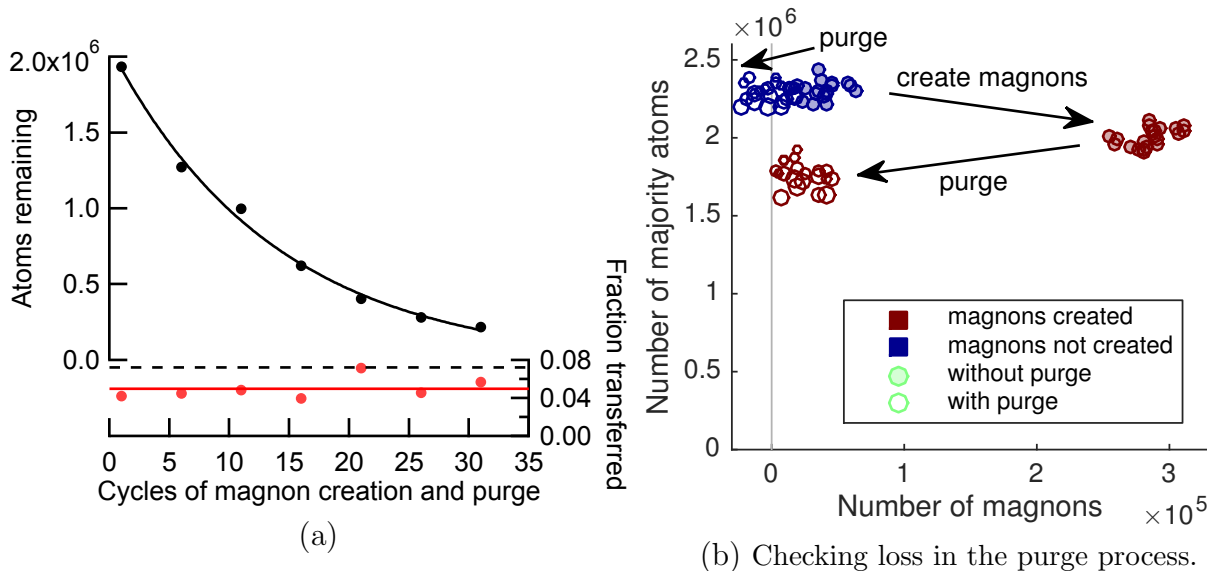


Figure 5.24: (a) Checking the purge process. We repeat cycles of magnon creation and purge. The final cycle does not include a purge. We simultaneously image the number of majority atoms and magnons after the last cycle to get both the atom number (top) and fraction of atoms converted into magnons (bottom, red points), which is roughly constant as the same RF pulse is used at each cycle. The atom decay rate (black dashed line) is compared to the magnon fraction and, being higher, indicates additional loss beyond the purged magnons. (b) A closer look at the loss of atoms in one purge cycle using the pump-probe procedure. Purge is applied (unfilled circles) to gases prepared with (red) and without (blue) magnons. Larger circles indicate more optical power in the purge pulse. When the purge sequence is applied to a gas without magnons (or with very few), almost no atoms are lost. When magnons are present, there is additional atom loss. Additional losses are slightly worse when the optical power is higher.

$\Omega_\mu \approx 5$ kHz, so that we can nearly saturate the microwave transition to the dressed excited state. In this scenario, the magnon lifetime with the purge is $\tau_0 \approx 1/(2\Gamma_{\text{eff}}) \approx 100 \mu\text{s}$, whereas the majority atom lifetime is $\tau_{-1} \approx 4\Delta^2/(\Omega_\mu^2\Gamma_{\text{eff}}) \approx 0.32$ s. Clearly this purge procedure is not non-destructive—0.32 s is a short lifetime by our standards—but ideally we only need to purge for several τ_0 to eliminate all of the magnons.

Our experience suggests, however, that the world is not ideal. For purging small numbers of atoms, this approach has worked well, approximately as expected. However, when a large number of magnons need to be purged, the effects of high optical density can become important: With a large number of atoms in the $F=2$ manifold, the purge light does not penetrate far into the gas, and the purge process takes many times longer than the simple analysis presented here would suggest. In this case, one might think of the atoms being purged in layers, one “OD” at a time. Purging a large fraction of the atoms can take several milliseconds.

Also, in this approach to purging, it is important that the light driving the optical transition be close to resonance, as even weak driving can result in a considerable AC Stark shift of the $|F=2\rangle$ levels. Such a shift could reduce the detuning from resonance of undesirable scattering events and decrease the rate of desirable scattering. Finally, our top-imaging beam typically operates near saturation intensity, and it is difficult to cut the intensity reliably by two orders of magnitude without feedback, so generally purging works best with our side imaging beam, which operates far below saturation intensity.

We were able to increase the reliability and robustness of the purging procedure by implementing a pump-probe procedure, rather than driving microwave and optical transitions simultaneously. In this scenario, we purge by first pumping the atoms with an (approximate) microwave π -pulse. Then we apply imaging light for a brief period, 50–150 μs depending on the optical depth of the atoms, to push the atoms out of the trap. To account for imperfect π -pulses, the procedure can be repeated several times. This purge procedure can be faster than our original method and is more robust to drifts in laser frequency and power. The final cooling data for our publication were acquired using the pump-probe purge method.

5.13 State purification and preparation

Atoms in the $|F=1, m_F = -1\rangle$ state, our majority gas, are magnetically trappable, and their presence in the imaging field of view can introduce unwanted aberration in our images of the magnon momentum distribution. Even if they remain 6.8 GHz out of resonance with the imaging light, their high density and moderate polarizability makes them an excellent lens that can cast a formidable shadow. Therefore, in order to image small numbers of magnons accurately, we begin an optimized purge process, similar to the one applied to trapped magnons in Sec. 5.12, but with microwaves resonant with the $|F=1, m_F = -1\rangle \rightarrow |F=2, m_F=0\rangle$ transition, removing the majority atoms from the field of view. The purge process begins about one millisecond after extinguishing the optical trapping light.

Purging majority atoms from the field of view presents some additional challenges beyond those encountered in purging the magnons. Unlike the trapped the magnon gas, the majority atoms are extremely dense, so considerations of high optical density are not just footnotes, they are front and center. The OD in the center of our BEC can easily exceed 50. In addition, the majority atoms far outnumber the magnons, and thus the constraints on the ratio of desirable to undesirable scattering events, such as that given for magnons in Eq. (5.18), are more severe.

We have used two different majority-atom purge procedures. The first procedure is a variation on the pump-probe purge that we adapted for use on the majority atoms. In this purge process, forces due to optical scattering remove the atoms from field of view. The second procedure relies on the magnetic field gradient used in MSF to push the majority atoms from the field of view. In this latter process, the majority atoms are transferred to a magnetically anti-trapped state. In the next few paragraphs I detail both of these purging procedures and discuss their relative merits.




Figure 5.25: Our ultimate scheme to remove $|m_F = -1\rangle$ atoms from the imaging field of view and to make $|m_F=0\rangle$ atoms compatible with momentum space focusing. Densely dotted levels are magnetically trappable.

The pump-probe optical majority purge Most of our data were acquired using a variation on the pump-probe purge, with alternating microwave (pump) and optical (probe) pulses. To reduce the optical depth, we used a pump power and pulse length that transferred only a fraction of the $|m_F = -1\rangle$ atoms into the $F=2$ manifold in each pump-probe cycle. The microwave pulses, probe power, probe duration, and number of pump-probe cycles could be varied to purge reliably while losing as few $|m_F=0\rangle$ atoms as possible.

In spite of this our attempts at optimization, the purge process removed between 10% and 40% of the $|m_F=0\rangle$ atoms, depending on the number of majority atoms being purged and on the particular mood of the experiment. These losses could possibly have been reduced by waiting longer before the purge sequence, at which point the density would be lower, or by using a magnetic field gradient to separate the clouds in momentum space. However, doing either of these things would have changed the timing of our momentum space focusing and imaging procedure, which we had already carefully calibrated. Crucially, the atom loss does not visibly impact the momentum of the minority atoms that remain. In counting magnons, we have to consider these losses.

The pump-sweep magnetic majority purge Eventually, after most of our cooling and thermometry data had been taken but before acquiring the state-selective trap data, we implemented a more efficient purge process that kept the same MSF timing. This process is depicted in Fig. 5.25. Enabled by better control and stability of the magnetic fields (Sec. 5.11), along with the ability to synchronize the levitation sequence with the line voltage, we found that a π pulse resonant with the $|F=1, m_F=-1\rangle \rightarrow |F=2, m_F=-1\rangle$ transition was stable enough to transfer reliably nearly all of the majority atoms. The $|F=2, m_F=-1\rangle$ state is magnetically anti-trapped, and so the action of the momentum space lens is to push atoms in that state out of the field of view.

Preparing the magnons for momentum space focusing In order to make the $|m_F=0\rangle$ atoms compatible with the levitating momentum space lens, we use a Landau-Zener (LZ) sweep to transfer them to the $|F=2, m_F=1\rangle$ state. This same sweep will also transfer any $|m_F=1\rangle$ atoms to the $|F=2, m_F=0\rangle$ state. In our apparatus, it is easier to ramp a magnetic field bias value than it is to ramp the microwave frequency, so we implement this sweep by fixing the microwave frequency and sweeping a magnetic bias field through the appropriate resonance. The ramp begins about 20 kHz detuned from resonance and moves to within a few kHz of resonance in about 0.5 ms. It then crosses to within a few kHz on the other side of resonance in 1 ms before darting to 20 kHz detuned in another 0.5 ms. We achieve more than 95% transfer efficiency with a microwave Rabi frequency of about 5 kHz.

With the magnons so prepared, we ramp on the momentum space focusing fields.

5.14 Data and data management

The primary data generated by our experiment are in the form of *images*. Associated with these images are innumerable experimental *settings*, most under computer control, along with even more *responses* of the environment to these settings. The environmental responses are, of course, highly correlated with the settings, otherwise our experiments would not work. But the responses are analog, even when the settings are digital, and subject to variations and fluctuations outside of our (formidable) control. We monitor many of the responses on our instruments: current clamps, power meters, magnetometers, error signals, oscilloscopes. Recorded or not, these measured responses are auxiliary data and potentially of use. Typically we discard these data, or account for them by eye or via a note in the lab book.

In a moderately complex study like the one described in this chapter, the volume of data we generate, and the variety of experimental settings, is fairly large. However, it is not generally so large that technological constraints force us to be clever in how we deal with it. We can store the raw data and the sequence files that generate it on disk indefinitely, allowing us to analyze it *post hoc* in whatever *ad hoc* way we choose. Auxiliary data and settings are hopefully noted in hard copy. If not, they may be safe to ignore. But maybe not. There have been times when I wished we could associate our data with, for example,

Figure 5.26: Some tables used to store data and analysis results. In the top row, two snapshots show tables that store experimental parameters and quantities pulled directly from primary data files derived from momentum space images and side images. There may be multiple images of from the side associated with each momentum space image, or *vice versa*. In the bottom left table, multiple different fits are applied to each piece of momentum space data, each yielding different derived quantities, such as temperature and condensate fraction. In the bottom right table, multiple different temperature results with different settings are combined to yield an extrapolated and/or average temperature.

the magnetic field measured in the lab at the moment it was acquired. However, the need has never been pressing enough to make it so, and the way we manage our data does not trivially lend itself to such coordination.

The ability to decide after the fact how to deal with primary data, and thus secondary and aggregated data that our analysis may generate, is a valuable luxury. Even though we may analyze the data as they come in, we frequently realize that the experiment is not what we think it is, and we are free to approach the data from an entirely different angle at a later time. Regardless, the way we organize data and analysis can have a big impact on our ability to tease out the patterns and trends they contain, on our ability to adapt them to testing evolving hypotheses, and on the way in which preexisting software tools can be applied to them.

In some ways, choosing how to manage data is akin to choosing a programming language for a project. A sensible default is to adapt what has been done before, or to use the tools with which one is familiar. In the programming space, in some senses— such as “Turing completeness”—all choices are equivalent. As a practical matter though, in terms of what sorts of things are easy, fast, or efficient, such choices can matter a great deal. In the sections that follow, the data, analysis, and results of our study will rightly take center

stage. However, the manner in which data and analysis were managed and stored will be a minor character, moving the story along. Let me introduce, briefly, that minor character now.

The primary data, the images, are acquired by *edcam* [2], a home built MATLAB-based image acquisition program. *Edcam* does basic processing of absorption images and produces files containing the inferred column density, as well as the raw image frames, should they be needed. These data files are stored on disk with a name determined by the operator, and are typically organized by date.

Information about the primary data, including metadata, such as variables relevant to the experimental sequence that generated a primary data file, and analysis results, were stored in an impromptu relational database [90] built on MATLAB's *table* data structure. A relational database makes a lot of sense for much of our data because there are strong relationships between different aspects of the data. For example, one or more images of the momentum space distribution of the atoms may be related to one or more images of the atoms from the side (to determine the atom number) and one or more images of the atoms *in situ*, all acquired with the same experimental parameters. These different pieces of primary data may be the source of many different derived quantities, such as the temperature or chemical potential. Additional quantities can be derived from these derived quantities, and so on, such as an average or extrapolated temperature. By taking care to structure the data in an intelligent way, writing code to analyze the data, perhaps in unforeseen ways, can become a lot easier.

Relational databases, along with other (possibly) relevant types of data management systems and concepts, are often covered in “programming for scientists”-type courses, such as Berkeley's “Python computing for science” [71]. Materials are often available online. Though such courses often focus on a particular programming environment, the principles of data management are quite general and might fill a much needed gap in the formal education of many experimental physicists. Keeping the many types of data in mind when building a new experiment might allow one to easily integrate environmental responses and other process data with the primary data, potentially allowing variability to be reduced via post-selection or more sophisticated data modeling.

5.15 Temperature fitting

Atom column momentum densities n_p are fit to a Bose-enhanced momentum distribution $n_p = b + Ag_j(e^x)$, with g_j the Bose function (polylogarithm) of order j and the argument

$$x = \alpha - (p_w - p_{w0})^2/2mk_B T_w - (p_t - p_{t0})^2/2mk_B T_t. \quad (5.19)$$

For non-degenerate gases, $\alpha = \mu/k_B T$ is a free parameter of the fit, along with the zero of momentum (p_{w0}, p_{t0}), the background level b , peak level A , and temperatures T_w and T_t along the weak and tight in-plane axes of the optical trap, respectively. For degenerate gases, $\alpha = 0$. If necessary, we also add a term that allows for a linear background gradient.

Generally, $T_w = T_t$ within the error of the fit. However, at low temperatures, several effects cause T_t to be an unreliable estimate of the gas temperature. Although the condensate expands primarily along the unfocused vertical direction, at very low temperatures the condensate expansion along the more tightly confined in-plane direction can be manifest (Fig. 5.9b). Also, the tilt of the trap in the tight direction is more difficult to calibrate, leading to a (fictitious) systematic upward shift in the apparent T_t by the mechanism explained in Sec. 5.4.4. Finally, Eq. (5.19) is expected to describe a Bose gas under the assumption that the semiclassical condition $k_B T \gg \hbar\omega$ holds along both the tight and weak axes of the trap. However, in our system $\hbar\{\omega_w, \omega_t\}/k_B \sim \{0.05, 1\}$ nK, and thus at the lowest temperatures we achieve this approximation is not valid along the tight axis. Thus, while we fit the full two-dimensional momentum densities to Eq. (5.19), throughout this work we estimate temperatures by $T = T_w$ alone.

5.15.1 Masking the condensate

Eq. (5.19) describes the normal component of the gas only, so we must exclude regions of the column density that include condensed atoms. Our final data set consists of nearly 5000 images with condensates and non-condensates of different shapes and sizes. Automating much of the fitting process, including finding and masking the condensate, if one is present, was quite important. Our procedure for masking a condensate is shown, in part, in Fig. 5.27. The algorithm for establishing the condensate mask is approximately as follows:

Compute the curvature of the image A Laplacian of a Gaussian (LoG) filter is applied to the image of the atom column density. The standard deviation of the Gaussian filter can be adjusted but is generally set to roughly twice the imaging resolution. The resulting image gives the curvature, C_i , of the smoothed image at each pixel i .

Mark points below a curvature threshold The distribution of curvature values differs markedly when a condensate is present and when one is not, as shown in Fig. 5.27b and 5.27c. Without a condensate, the distribution of curvatures is approximately symmetric. With a condensate, a long tail of negative curvatures is present. Candidate condensate pixels are established by marking pixels with a curvature less than $-\max(C_i)/\beta$, where β is a tunable constant of order unity that sets the sensitivity of the threshold.

Unmark regions less than a minimum size Noise in our images can cause isolated pixels or small groups of pixels to be tagged in regions where there is no condensate, or even no atoms. Thus, we unmark candidate pixels that are part of a group of candidate pixels smaller than a certain size, using the standard “closing” image morphology operation. We used a threshold radius of 3 pixels. We call the resulting marked pixels the mask.

Expand the mask To make sure the mask covers the whole condensate, we expand it by a variable amount using a “dilation” operation. We may want to mask a region

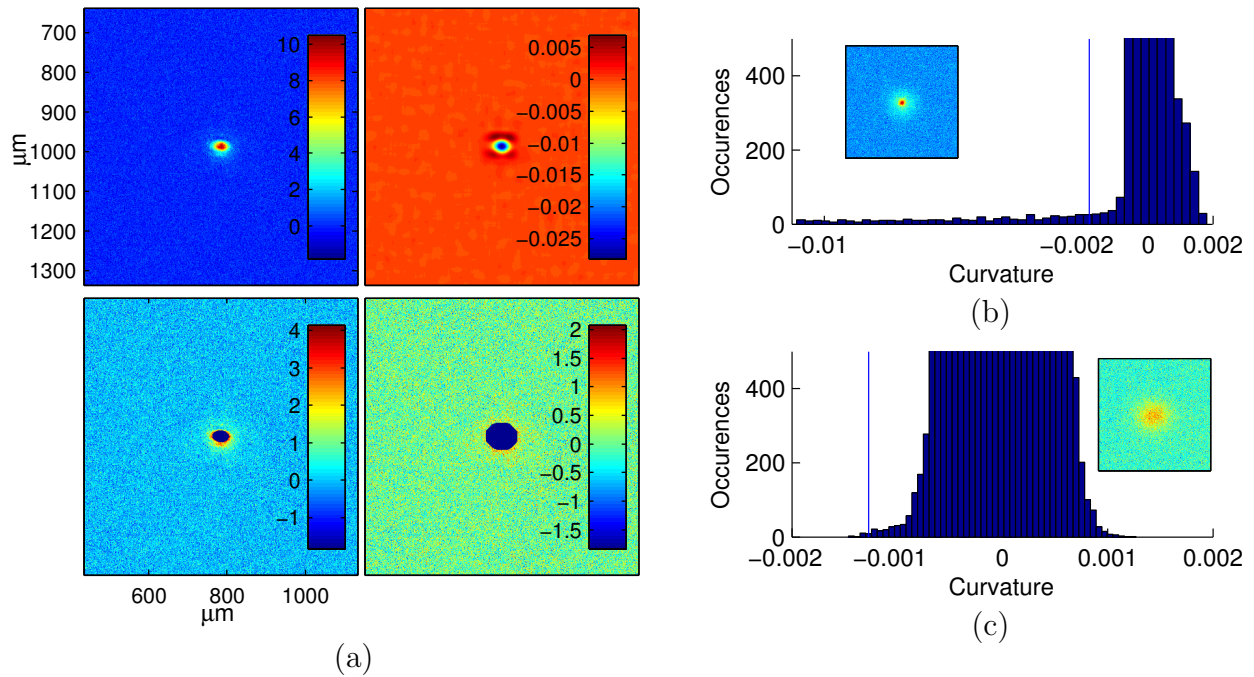


Figure 5.27: Masking the condensate automatically. (a) Top row: the image of the transverse momentum distribution, with color scale indicating density (left) and its curvature (right). Bottom row: The mask (left) and expanded mask (right) are shown in dark blue, superimposed on the transverse momentum density distribution. The distribution of curvatures is also shown both with (b) and without (c) a condensate. The masking threshold calculated from the distributions is indicated by the vertical blue line. Insets show the corresponding transverse momentum density distributions. Color scales assign dark red to maximum unmasked density or curvature in each image.

much larger than the condensate, for example to just fit the wings of the momentum distribution, so the amount by which we expand the mask may be quite large.

Merge with minimum mask Even if there is no condensate detected, we may want to mask the central part of the distribution. Thus, we combine the automatically determined mask with an adjustable minimum-size mask centered on the density distribution.

This masking procedure has been remarkably successful at finding the condensate in a wide range of our data, including momentum space, *in situ*, and time-of-flight images. The tuning constants have proven to be quite robust, and require only minor tweaks depending on the data.




Figure 5.28: Assortment of fits varying the mask size. Fits using the Bose function g_2 are in red while fits using $g_{1/2}$ are in blue. The mask size (x -axis) is normalized by the standard deviation of the selected fit, which is indicated by the square box. The first two frames are typical of samples with a condensate present, while the last frame is typical of samples without a condensate. Notice that both functions underestimate the temperature at smaller mask sizes (when a condensate is partially unmasked). At intermediate mask sizes, where the Bose-enhanced parts of the non-condensed distribution are present, the $g_{1/2}$ fits tend to overestimate the temperature. Fits with g_2 generally converge to the correct temperature (as indicated by the wings of the distribution, which are insensitive to the fit function employed) at smaller mask sizes.

5.15.2 Selecting a fit

In order to have reliable temperature estimates, we must verify that the temperature is insensitive to the size of the masked exclusion region. Thus we fit each image multiple times with masks of varying size. In addition, it is not clear *a priori* which Bose function we should be fitting to. At high temperatures, $k_B T \gg \mu$, the non-condensed atoms (including magnons) have little overlap with the condensate, and the statistics should be those of atoms in a harmonic trap. In such a case, the function g_2 should describe the momentum column density. At low temperatures, when $k_B T \ll \mu$, the magnons are confined to the region of the condensate and see an effective “square well” potential, implying a momentum column density approaching $g_{1/2}$. However, $g_{1/2}(e^x)$ diverges unphysically as $e^x \rightarrow 1$, and $e^x = 1$ at the center of a quantum-degenerate gas. This incongruity suggests that some other physics should keep the unphysical $g_{1/2}$ behavior from being realized. For example, weak interactions and finite size should prevent a true divergence. In addition, more recent work in our lab on the condensation of magnons suggests that the timescales for true equilibration of the magnons at the longest length scales is longer than the lifetime of our gas.

Fortunately, the high-momentum tails of the momentum distribution n_p are insensitive to the particular Bose function employed ($g_j(e^x) \rightarrow e^x$ as $e^x \rightarrow 0$). Thus, it is important to make sure that some of our fits include only the high-momentum data points. In addition, we will

also fit our data with two candidate Bose functions, $g_{1/2}$ and g_2 . We can then determine which Bose function best applies to our data by seeing which function most consistently estimates the temperature and magnon number as the mask size is varied. If we determine the proper Bose function, we can increase our signal to noise by including more low-momentum data in the fits while maintaining temperature consistency with the unbiased high-momentum tails of the distribution.

Both g_2 and $g_{1/2}$ tend to underestimate the temperature when the mask is very small and parts of the condensate are included in the fitting region. As the mask grows, the temperature predicted by g_2 tends to rise before it plateaus. In contrast, as the mask size grows, the temperature implied by $g_{1/2}$ will tend to rise above the asymptotic temperature value before converging on the true temperature from above. In both cases, once the mask is too large, the fits become unreliable as the signal to noise drops to zero. This behavior can be seen in the examples of Fig. 5.28 and in panel 10 of Fig. 5.29, where the temperature predicted by g_2 (red points) and $g_{1/2}$ (blue points) are shown versus mask size.

In general, when fitting condensed samples, g_2 seems to converge to the true temperature value faster than $g_{1/2}$. Often, especially at low temperatures, they converge equally well. The fact that g_2 underestimates and $g_{1/2}$ overestimates the temperature when the central part of the degenerate distribution is included in the fit suggests that $j \sim 1$ or $3/2$ might fit the data better. Indeed, in our ongoing (as of this writing) work exploring magnon condensation, this seems to be the case. While the choice of Bose function is not critical in estimating the temperature using the tails of the momentum distribution, it can impact the inferred condensate fraction a great deal. The more highly peaked Bose functions—those with smaller j —tend to ascribe more of the atoms in the central region to the non-condensed fraction.

Much of our thermometry data were obtained by imaging entirely non-condensed samples. Far from degeneracy, both g_2 and $g_{1/2}$ are equally good at describing the momentum distribution. This is the typical case for much of our data, including the third panel in Fig. 5.28. Even for samples approaching degeneracy, we noticed once again that in general g_2 seemed to describe the observed distribution better than $g_{1/2}$, even at low temperatures. This behavior is a bit puzzling and is being investigated in more detail in ongoing work. I hypothesize a partial explanation: First, as mentioned previously, the timescales required to achieve thermal equilibrium at long length scales seem to be quite slow in our system. This has also been observed in other experiments in our lab [25] and elsewhere [15]. The result of this may be that at long length scales, where the difference between $g_{1/2}$ and g_2 behavior is manifest, the gas is unaware, in some sense, that it is approaching degeneracy. In addition, when $k_B T < \mu$ and $g_{1/2}$ may be expected to apply, the effects of trap anharmonicity in the vertical direction become relevant, softening the box somewhat. Regardless, for this study, reliable temperatures could be obtained by yielding to the well-thermalized long length scales (e.g. high momentum), as necessary.

Our procedure for efficiently fitting and examining the fits leverages the relational data storage described briefly at the end of Sec. 5.14. When data were initially acquired, the location of the primary data files and the experimental parameters are entered into a table

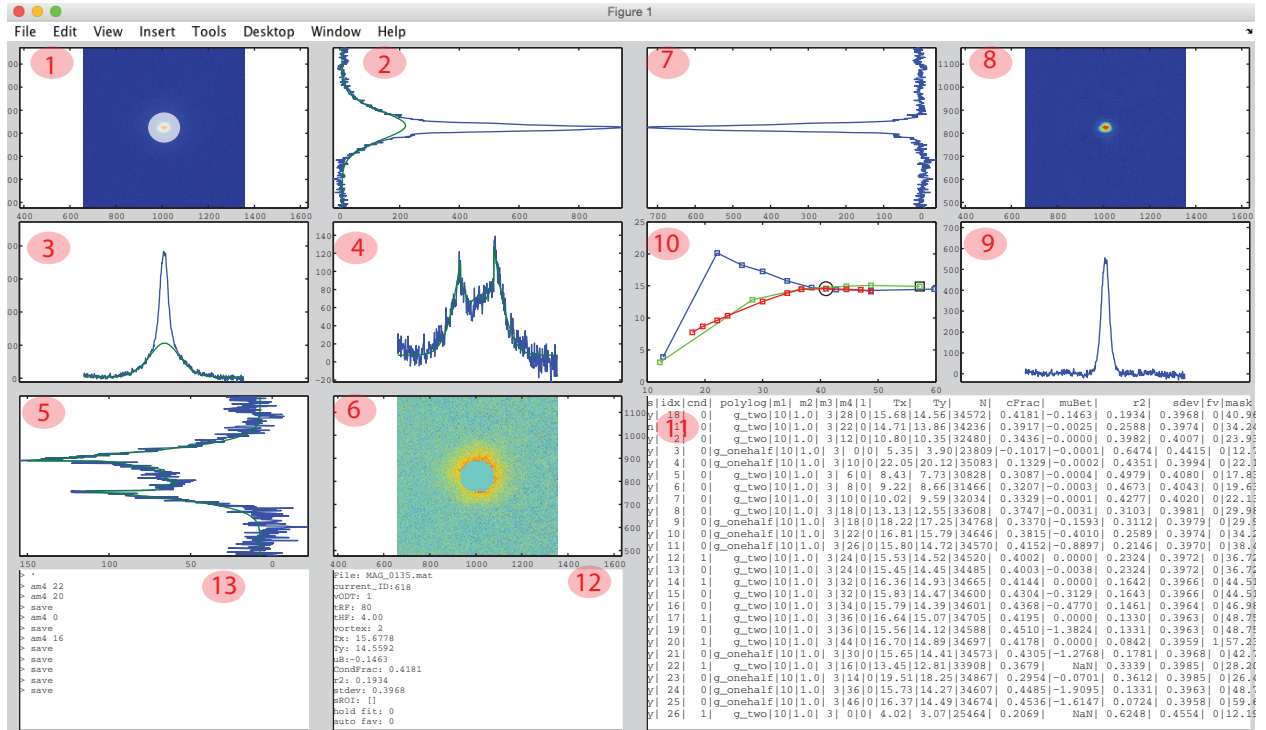


Figure 5.29: The “fit maker” was used to check fits, apply new fits, and select the favorite fit for each datum. This makeshift GUI is just an interactive MATLAB figure window. Panels 1–3: The original transverse momentum distribution and integrated profiles with integrated fits. The mask used for the fit is superimposed on the momentum distribution. Panels 4–6: Corresponding momentum distributions and integrated profiles with the masked data excluded. These panels show only the data used to produce the fit. Panels 7–9: The implied condensate column density and integrated profiles are obtained by subtracting the fitted non-condensed density. Panel 10: A graphical summary of the different fits performed on this datum shows the temperature vs. mask size. The three different colors differ in parameters other than mask size. Red and green points use g_2 Bose function with degeneracy parameter α free and fixed to zero respectively. Blue points use $g_{1/2}$ function and α free. The currently displayed fit is marked with a circle. The favorite fit is marked with a square. Panel 11: Numerical summary of different saved fit results and parameters. Panel 12: Information on the experimental parameters used for this datum, as well as a few fit highlights and configuration parameters of the fit maker. Panel 13: A command line allows new fits to be generated and saved, for fits to be selected as favorite, for navigation between data, and for settings of the fit maker to be changed. Time saving settings included automatically selecting the last-viewed fit as favorite and locking the fit settings while navigating between different data so that, for example, good fit settings do not need to be reentered manually for each item of related data.

and marked as new. A simple script is called to examine any new data and to add a standard battery of desired fits (which depend on the experimental parameters) to another table. Another script monitors this table of fits and would-be fits for unprocessed rows and calls the fitting routine to fill in any missing fit results. Each fit can take several seconds to perform, and at times, such as when a new battery of fits is to be applied to all of the data, many thousands of new fits might need to be performed. With the data tables stored on a central server, the fitting routine can easily be parallelized. To make viewing and recalling fits faster, the results of calls to the fit procedure are cached in a hash table whose key is uniquely determined from the data and fit settings.

While our fitting routing is quite robust and well automated, picking a particular set of fit parameters that produce the best unbiased fit is not. We need to check that the choice of Bose function and mask size are reasonable and that the fit results don't have a strong dependence on particular, somewhat arbitrary, choices. In some data, we need to decide whether to compensate for a linear background, or not; we may need to decide whether to fix the degeneracy parameter α (owing to the presence of a condensate), and to be sure this choice does not bias the data. And we need to do this for several thousand different images. To make this process relatively easy, we built a simple tool for browsing the data while adding, comparing, and selecting fits: the “fit maker,” shown in Fig. 5.29. The fit maker shows a summary of all of the fits that have been performed on a particular image along with details and views of a particular fit. It includes a capability to add new fits with varying parameters, and to select a particular “favorite” fit for use in subsequent analysis. All of the data and fits employed in this work were examined and selected using the fit maker.

5.16 Thermalization

In order to function as a thermometer and to cool the gas, the magnons have to be able to exchange energy with the majority gas at a rate sufficient to achieve a suitable approximation of thermal equilibrium. When $k_B T > \mu$, the thermal excitations of the majority gas disperse quadratically, similarly to the magnons, and the magnon gas would be expected to thermalize in step with the majority gas itself. However, when $k_B T < \mu$, the thermal excitations of a homogeneous Bose-Einstein condensate are linearly-dispersing phonons, and thus thermalization of the magnons with the majority gas requires collisions between quasi-particles with very different dispersion.

We found that an applied magnetic field gradient could expedite thermalization between initially coherent spin components. With no gradient (though a small background curvature remains, roughly 50 mG/cm² along the long axis of the trap), the out-of-equilibrium two-component condensate can persist for several seconds when the temperature is sufficiently low. A less dramatic demonstration of the effects of a magnetic field gradient is shown in Fig. 5.30, where the “melting” of the condensate is shown to be expedited by moderate gradients. The optimal gradient to achieve thermalization in our system is moderate, on the order of 0.2–2 mG/cm, depending on the temperature of the gas. We avoid larger gradients




Figure 5.30: Images of the magnon gas *in situ* at an intermediate temperature in an older trap geometry are shown with varying magnetic field gradient and hold time. At short times, a dense condensate is seen, reflecting the *in situ* distribution of the initially spin-polarized gas. In the presence of low gradients (top row), the dense condensate persists to later times, “melting” slowly. Higher gradients initially facilitate faster melting, as the condensate moves more quickly towards the edge of the trap. When the gradient is too large, however, the magnons re-condense in a reduced trap volume.

for several reasons. First, a large gradient has the potential to deposit extra energy in the condensate, leading to atom loss and low-order collective motion of the gas that takes time to dissipate. Additionally, recall that the magnons have an effective magnetic moment [47], so a gradient that is large compared to the temperature of the gas divided by the size of the trap reduces the trap volume for magnons by confining the magnons to the edge of the trap, lowering the magnon condensation temperature. This effect can be seen in Fig. 5.30, wherein the condensate is seen not to melt at the larger gradients.

A gradient of 0.2 mG/cm corresponds to an energy of approximately $h \times 10 \text{ Hz} = k_B \times 0.5 \text{ nK}$ across the longest dimension of our gas and was sufficient to produce a saturated non-condensed fraction of the $|m_F=0\rangle$ state in 100–2000 ms, depending on the temperature of the gas. Thermalization of magnons created out of a degenerate gas at different temperatures are shown in Fig. 5.31. In this thermalization experiment, the total time after the end of the evaporation ramp is constant, so that effects of thermalization of the entire gas following evaporative cooling are removed.

When magnons are created at an intermediate trap depth before the final evaporation

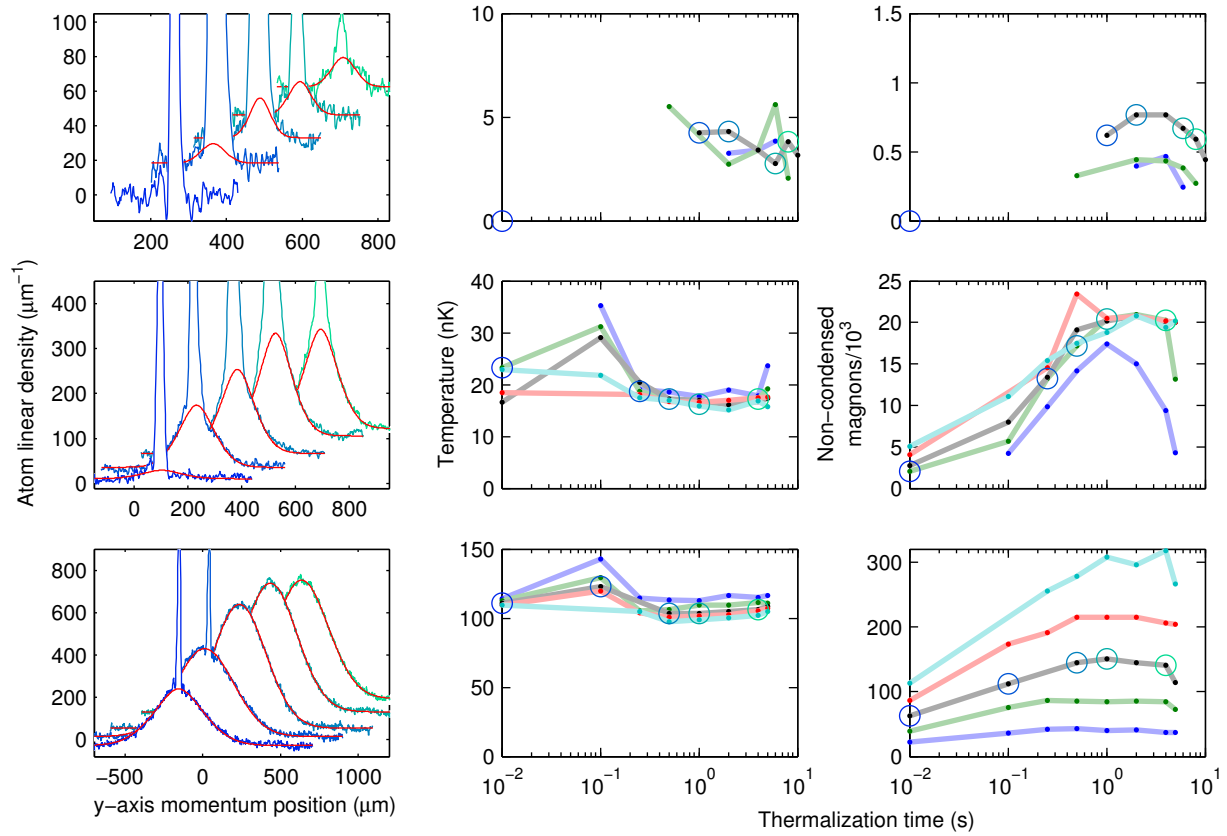


Figure 5.31: Thermalization of magnons of magnons created at trap depths of $k_B \times 35$ nK (top row), $k_B \times 210$ nK (second row), and $k_B \times 1275$ nK (bottom row), with total time following evaporative cooling ramp held constant. In the first column, integrated profiles (blue-green lines) and fits to the thermal component (red lines) are offset for clarity. The points corresponding to each fit in column one are marked by same-color circles in the second and third columns. Several features of magnon thermalization can be seen here. At the lowest trap depth, the temperature cannot be determined until after nearly one second of thermalization, at which point the temperature has reached steady-state and does not depend on the number of magnons. At this trap depth, the number of magnons does not seem to saturate (likely due to high loss rates) and the peak number of magnons seems to vary with the size of the initial magnon condensate. At the intermediate trap depth, the temperature can just barely be determined from the initial momentum distribution. Note the scatter in initial temperatures owing to the very low signal. The temperature reflected by the magnons initially rises by an amount inversely proportional to the number of magnons created, but the final temperature achieved is lower when more magnons are created owing to decoherence cooling. In contrast to the low trap depth case, the thermal magnon number saturates to a value independent of the initial condensate number, as long as the initial condensate number is sufficiently large. At high trap depth, similar patterns of thermalization occur, but are more pronounced. The initial temperature is well determined. At this trap depth, we typically do not create enough magnons to saturate the thermal gas.

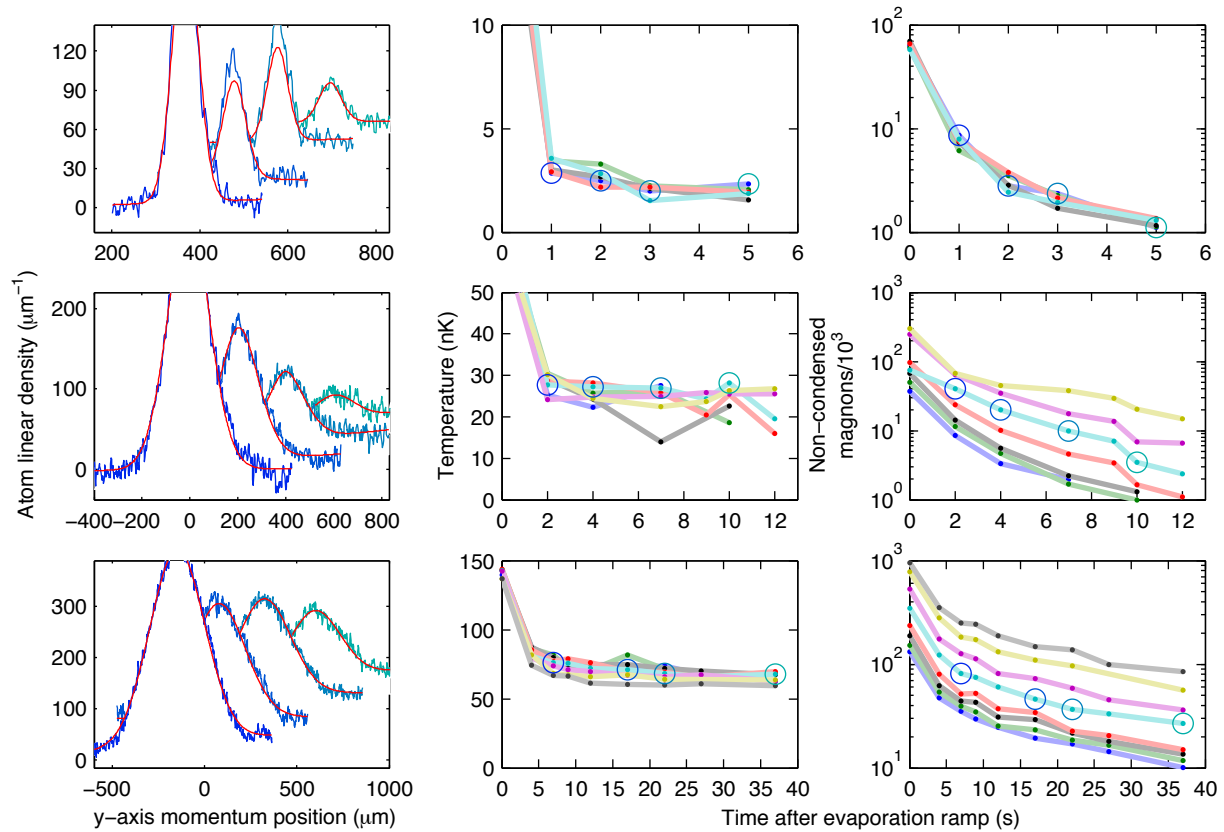


Figure 5.32: Thermalization of magnons created in a non-degenerate sample at a high trap depth before a several second ramp to a final trap depth of $k_B \times 35$ nK (top row), $k_B \times 210$ nK (second row), and $k_B \times 810$ nK (bottom row). In the first column, integrated profiles (blue-green lines) and fits to the thermal component (red lines) are offset for clarity. The points corresponding to each fit in column one are marked by same-color circles in the second and third columns. At low trap depths, the temperature signal is generally larger than when magnons are created at the final trap depth, as in Fig. 5.31. In the second and third fits in the top row, slow thermalization of the low energy magnons is evident in the messy distribution near the peak. However, the higher momentum wings still give a consistent temperature signal. At higher trap depths, the effects of magnon-assisted evaporation can be seen. Note also that steady-state temperatures are actually reached faster at lower trap depth. This is likely due to the smaller size of the gas and the increase in evaporative cooling power as the ratio of effective trap depth to temperature drops.

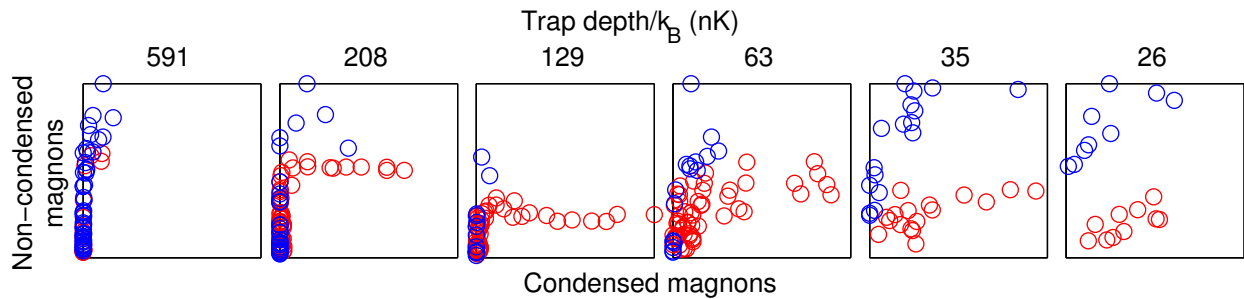


Figure 5.33: Saturation of the thermal gas. Number of non-condensed magnons and condensed magnons created at an intermediate trap depth (blue) and final indicated trap depth (red) are plotted with x - and y -axes set to the same scale in each frame. Scale differs from frame to frame, with fewer magnons at lower trap depths. All points are normalized by the trap volume, which depends on the number of majority atoms. At moderate trap depths, saturation behavior is clear as the number of non-condensed magnons increases to a threshold before the appearance of condensed atoms. At high trap depth, we do not create enough magnons to saturate the gas. At low trap depths, the higher signal of the magnons created before the final phase of evaporation is apparent. For magnons created at the final trap depth, the number of non-condensed magnons varies with the condensate number.

stage, thermalization is indicated by the achievement of a steady-state temperature that varies with the majority trap depth even when the magnon trap depth is fixed (see Sec. 5.17) and that is consistent with that of the majority gas alone (when it is measurable) and that of magnons created in the degenerate sample under similar conditions. Fig. 5.32 shows the temperature and number trajectory of magnons created at an intermediate trap depth. The saturation of the magnon gas apparent in Figs. 5.31 and 5.32 is explored further in Fig. 5.33.

We speculate that thermalization of the magnons occurs primarily at the edges of the condensate where low condensate density, and hence smaller local chemical potential, implies a better match between majority and minority spin dispersion. In addition to facilitating decoherence, the applied gradient may aid in thermalization by driving the transport of low-energy magnons to the edge of the condensate.

Magnon-phonon interactions are expected to occur even in the homogeneous BEC in the first order beyond the Bogoliubov mean-field theory (the Beliaev theory) [68], though I am not aware of any calculated estimates or direct measurements of the thermal phonon-magnon cross section. Our data are not sufficient to produce such an estimate, as our images do not provide reliable estimates of the number of condensed magnons, and the high magnon loss rate at the lowest trap depths makes accounting precisely for magnon production rates difficult. Moreover, as discussed in Sec. 6, our condensate is not well approximated as uniform, or even locally uniform, so frameworks that depend on the local density approximation are unlikely to apply.

The physics behind the decoherence and thermalization of spin impurities, and the variables that affect it, remains a compelling direction of future research. Previous experiments

on degenerate $F=1$ ^{87}Rb gases have found magnetic field inhomogeneity to facilitate decoherence and thermalization of spin populations, typically on timescales of several 10's–100's of milliseconds [12, 39, 37]. In our work on contrast interferometry with magnons, we observed coherence between magnons in different momentum states out to several hundred milliseconds [47] with a negligible gradient. However, none of these experiments claim to have operated in the highly-degenerate phonon-dominated regime $k_B T < \mu$. On the other hand, in strongly-interacting Fermi gases, decoherence of spin impurities has been well characterized and is limited by a fundamental lower-bound on the spin diffusivity, \hbar/m [5, 36]. Thus, in contrast to weakly interacting Bose systems such as ours, in the strongly interacting Fermi gas, the role of the gradient in driving diffusive spin transport is well understood.

5.16.1 Real-space magnon thermalization

When the magnons are initially created, their real-space distribution matches that of the harmonically trapped majority gas. However, at low temperatures, the equilibrium spatial distribution of non-interacting (dilute) magnons is that of free particles confined to a box. To observe the transition, we imaged magnons *in situ* as they thermalized. Indeed, in the data shown in Fig. 5.34, the faint shadows of the gas at long equilibration times do seem flatter than the initial profile. Unfortunately, the combination of low signal owing to atom loss and the small width of the cloud in this data set make a quantitative demonstration of the flat magnon potential difficult. In more recent work, we have observed *in situ* profiles of thermalized magnons that demonstrably match the predicted profile of magnons in a box better than magnons in a harmonic trap by employing a trap with rounder transverse dimensions [18].

Initially, our *in situ* observations of magnon thermalization showed a lot of unexpected clumping of the magnons, as shown in the left panel of Fig. 5.34. Further investigation revealed the presence of vortices in the majority gas (the vortices could be seen in the majority gas after a long time of flight as density dips in the condensate distribution). By evaporating more slowly, we could eliminate the vortices, and the degree of clumping of the magnons during thermalization was greatly reduced, as shown in the right panel of Fig. 5.34. The reason that some clumpiness remains, even without vortices, is a mystery, though it could be related to the emergence of a length scale that appears in many of our group's experiments on condensation in spin mixtures [25].

Much of our data in which magnons are created in a degenerate sample at the final trap depth have been acquired using similar evaporation parameters to those where we observed vortices. To determine whether vortices impact the temperature indicated by the magnons, we repeated some of our previous experiments with vortex-free evaporation ramps. Other than the *in situ* character of the thermalizing magnons, the presence of vortices does not seem to matter.




Figure 5.34: Magnons are created in a degenerate gas at low temperature and imaged *in situ*. When the majority gas contains vortices owing to a fast evaporation ramp, the magnons clump as they thermalize. Without vortices, the magnons clump much less. The distribution of magnons at long times should be more uniform as their 3-D density is expected to be constant in the trapping volume.

5.17 State dependent trap

Our data show clearly that the magnon gas thermalizes and reaches a steady-state temperature. However, we need to verify that the temperature reached by the magnons reflects energy exchange with the majority gas and not just independent evaporation and self-thermalization of the magnons. To do this, we vary the trapping potential depth of the majority gas, and thus its temperature, while holding the trapping potential depth of the magnons constant.

Varying the trapping potential depth of the majority gas changes the number of majority atoms and thereby changes the chemical potential and *effective* trap depth of the magnons, which is the potential energy depth minus the chemical potential, $U_{\text{eff}} = U - \mu$. If the magnons thermalize independent of the majority gas then we expect the magnon temperature to vary in proportion to their effective trap depth, even when the trap potential depth is constant. On the other hand, if the magnons thermalize with the majority atoms, the magnon temperature should vary with the majority atom trap depth. Fortunately, these two competing effects—self- and co-thermalization—pull the magnon temperature in opposite directions.

Fig. 5.35 shows clearly that co-thermalization is the dominant effect, as at constant magnon trap potential depth the magnon temperature varies counter to the effective magnon trap depth and with the majority atom effective trap depth (and trap potential depth). However, the data also suggest that the temperature reflected by the magnons is not strictly proportional to the estimated majority atom effective trap depth. There are several possible explanations for this that might warrant further study.




Figure 5.35: Temperatures indicated by the magnon momentum distribution in the state-dependent trap have a positive correlation with the majority effective trap depth and a negative correlation with the magnon effective trap depth at constant magnon trapping potential depth. Here, same-color data connected by a dashed line were acquired at the same magnon trapping potential depth, but with varying majority atom trapping potential depth that was either lower (squares), higher (circles), or the same (diamonds) as the magnon potential depth. Different majority atom numbers result in different effective trap depths for the magnons even when the trapping potential depth is constant. The magnitude of the temperature changes is discussed in the main text. Vertical error bars show standard error of the mean over several measurements and horizontal error bars show systematic uncertainty in the trap depth predicted by the model of Sec. 5.7.

First, variations in the temperature at the same majority atom trap depth (but varying magnon effective trap depth) could be due to magnon assisted evaporation, which we have documented at higher trap depths. We have not observed that temperature varies with the number of magnons at these low trap depths, however, the dynamics of evaporation at such low trap depths is not well understood, and it is possible that magnon-assisted evaporation does occur but saturates at low magnon numbers. Second, the trap depth itself is not precisely calibrated, or even well defined, at such low trap depths. The systematic error in our estimates of trap depth are proportionally large. Moreover, as explored in Sec. 5.8, the trap depth is not constant across the trap and can be very sensitive to tilts. The data displayed in Fig. 5.35 was acquired with a tilt in the less-sensitive short in-plane axis that was noticeable in the momentum-space expansion of the cloud orthogonal to the direction used to infer temperatures.

5.18 Error analysis

5.18.1 Statistical error

One of the powerful features of our magnon thermometer is that its high signal-to-noise ratio allows us to estimate temperatures from a single image. In essence, each magnon independently samples the collective momentum distribution, and each image of the gas constitutes the results of hundreds to hundreds-of-thousands of identical, simultaneous experiments, giving us a formidable atomic signal, as well as atom shot noise. Each individual atom is probed by a large, random number of photons, yielding a large optical signal, along with the inevitable¹³ photon shot noise. Our camera is sensitive, but relies on stochastic quantum and thermal processes that can introduce additional noise. Finally, technical noise—including dust and vibrations—has the potential to further spoil our precious temperature signal. The noise in our imaging system is dominated by photon shot noise [49] and when imaging a sufficient number of atoms, the atom distribution can be determined at the atom shot noise limit, at least at moderate length scales [46].

By default, our temperature fits do not take noise into account. In effect, this means that the fitting algorithm, which is based on a non-linear least-squares optimization (MATLAB's *nlinfit*), assumes that the error on each pixel is equal and drawn from a normal distribution. However, this is not the case. While the imaging noise is roughly equal across the fitting region insofar as the imaging beam intensity is roughly constant across the region, the atom shot noise varies considerably. We need to verify that ignoring the variation in error is not biasing our fits. In addition, when errors are not explicitly included in the fitting routine, the fitting routine must determine from the data and the fit itself what size of a deviation from the least-squares result is statistically meaningful, or not, once again employing assumptions that may not hold up.

To better understand the uncertainty of our temperature estimates, and to verify that the fits are not systematically shifted by having done the bulk of our fits without specifying the pixel-by-pixel errors, we perform two additional fits. First, we refit the lowest temperature data using our selected fit settings with noise-weighted errors. We estimate the noise at pixel i as

$$\text{Var}(n_i) = \sigma_i^2 = \tilde{n}_i + \text{Var}(\{n_j\}),$$

where \tilde{n}_i is the estimated number of atoms in pixel i obtained from the previous non-noise-weighted fit and the $\{n_j\}$ are a selection of column densities obtained from pixels without atoms present, but with comparable imaging optical power. The first term accounts for atom shot noise and the second for imaging noise, which is dominated by photon shot noise. As shown in Fig. 5.36, the temperature estimates with and without noise weighting are not substantially different.

Next we use the results of this fit to construct a simulated image that is perfectly described by our model. To do this, we use the parameters of the fit to produce a (noiseless) column

¹³Unless we can somehow use highly squeezed states for imaging.




Figure 5.36: The difference between temperature estimates from fits with and without uncertainty weighting is minor for the 271 lowest temperatures, generally much less than the uncertainty in the fit itself.

density profile. We then add to each pixel independent Gaussian noise with variance σ_i^2 and refit the resulting (noisy) data. The errors on the parameters of the fit of the simulated data can be compared to the errors on the parameters of the fit to our actual data. If our experimental data are well described by our model, and their noise is as expected, the simulated and actual fits should be essentially identical. Note that our noise model neglects correlations in the atom shot noise between different pixels, which might be the case owing to finite imaging resolution. The experimental and simulated errors are compared in the left panel of Fig. 5.37.

We can also analytically estimate what the error in our temperature estimates should be given our noise in the limit that each pixel represents an independent measurement of the momentum distribution and that all parameters of the fit except the temperature are fixed. Obviously, this latter requirement is not entirely valid, and in reality the variances of and covariances between different fit parameters probably cannot be ignored. Still, in this simple limit, the expected error in the temperature estimated from pixel i is derived from the fit function n_p as

$$\text{Var}(T_i) = \text{Var}(n_i) / \left(\frac{\partial n_p}{\partial T} \right)^2.$$

The overall variance in the parameter T estimated from a weighted average over the individual measurements (i.e. each pixel) would then be

$$\text{Var}(T) = 1 / \sum_i \text{Var}(T_i)^{-1}. \quad (5.20)$$

The calculated and experimental uncertainties are compared in the left panel of Fig. 5.37. The calculated error is suggestively close to a factor of two less than the experimental and simulated errors. It is not surprising that the calculated error is less given the assumptions made in the calculation, but one might wonder if a factor of two is missing somewhere. In a careful examination, none was found.




Figure 5.37: Comparison of experimental, calculated, and simulated statistical uncertainty in the single-shot temperature estimates. In the left panel, the uncertainty derived from our actual data assuming atom shot noise and photon-dominated imaging noise matches well with that derived from perfect simulated data perturbed by the assumed noise. In the right panel, the simple model of uncertainty described in the test underestimates the actual noise observed by a factor of two. The dashed gray line has a slope of two, the solid gray line has a slope of one.

The temperatures were measured many times at each trap depth, and if the error estimates given by our fitting routine are correct, they should accurately reflect the variation in the independent repeated experiments. In statistics, a simple (and like much in statistics, much abused) test of such consonance is the reduced Chi-square statistic,

$$\chi_{\text{red}}^2 = \frac{1}{N-1} \sum_i \frac{(T_i - T)^2}{\text{Var}(T_i)}.$$

$\chi_{\text{red}}^2 \sim 1$ when the individual variances $\text{Var}(T_i)$ are consistent with the scatter of the (assumed to be independent) data. Here, $N-1$ is the number of degrees of freedom, N measurements less one degree of freedom to determine the uncertainty-weighted average temperature T . For example, if the T_i are independent and drawn from identical normal distributions with known standard deviation $\sigma = \sqrt{\text{Var}(T_i)}$, the numerator is just $N-1$ times the unbiased estimate of σ^2 and the expected sum would be one.

As shown in Fig. 5.38, at the lowest temperatures that we access, the scatter of the data are consistent with the uncertainty assigned by the fitting routine, with the χ_{red}^2 hovering around unity. One outlier shows gives a $\chi_{\text{red}}^2 \sim 4$. Closer inspection reveals this deviation to be due to a few fits with uncommonly low error, but differing estimates of the temperature. Regardless, in all cases, using weighed and unweighted means give very similar temperatures. At higher temperatures, a weighted mean no longer describes the data as the temperature




Figure 5.38: The χ_{red}^2 statistic is plotted versus the weighted average over many runs with the same settings. In general, at the lowest temperatures shown $\chi_{\text{red}}^2 \sim 1$ as expected when the error assigned by the fitting routine matches the scatter of the data.

varies with the number of magnons, as described in the Sec. 5.19. In cases where the uncertainties are well matched to the scatter in the data, such as our lowest temperature points, we can confidently reduce our uncertainty in the weighted mean of the independent measurements just as in Eq. (5.20).

5.18.2 Systematic error

There are several sources of systematic error in our system, most of which have already been discussed. The dominant source of systematic error arises from the uncertainty in the MSF magnification (Sec. 5.4.3), which results in a constant 6% uncertainty in the temperature. Another source of systematic error that we consider is in the imaging resolution. Uncertainty in the magnification of our imaging (Sec. 5.5.2) system contributes a relatively minor 2% uncertainty to the temperature estimates. The combined uncertainty of these two effects is about 6.5%, which is what the value used to compute the uncertainties quoted here and in our publication [62].

Assuming a Gaussian thermal distribution, the effect of finite imaging resolution (Sec. 5.5.1, approximately $8 \mu\text{m}$ in our system) should be to introduce a constant positive offset of 0.2 nK in the measured temperature. We do not correct our reported temperatures for this effect (though we note this effect in the methods section of our paper), which means that our reported temperatures likely overestimate the true temperatures by this amount. An alternative approach would have been to correct the temperatures for the imaging resolution and incorporate the uncertainty in our imaging resolution (and the uncertainty introduced by the various ways in which one might choose to correct the data) into the systematic error. We chose to avoid additional data processing and to error on the side of overestimates.

An additional source of error pertains to the variable effective mass of the magnon, m^* , which we previously measured to be 3% heavier than the bare ^{87}Rb mass in our system [47], though in a different trap geometry. The origin of such a large shift is not entirely

known. Some amount of the shift, 0.3% can be accounted for by considering interactions one order beyond the usual mean-field Bogoliubov theory (the Bilaev theory) [68]. We suspect the remaining portion may owe something to dipolar interactions. Regardless, it is not obvious how this shift in the magnon mass affects the interpretation of the MSF data. Do the MSF data reveal the *in situ* momentum distribution? In such a case the shifted effective mass m^* should be used to infer the temperature from the momentum width: A momentum p implies $E = p^2/2m^*$. Do the MSF data reveal the *in situ* velocity distribution? In this case, $E = m^*v^2/2$ and the momentum revealed by MSF would need to be corrected to account for the mass difference between the *in situ* magnons and the free atoms undergoing MSF. Perhaps neither momentum nor velocity is conserved, and the change in magnon momentum requires a dynamic description akin to the hydrodynamic theory discussed above. My intuition is that the magnon momentum is conserved as the condensate density drops after the trap is extinguished, and that the effective mass should be used to infer the temperature. However, we use the bare mass of ^{87}Rb in all of our estimates. We did not consider this affect in our initial analysis, and so it is not included in our estimates of the systematic error, though it probably should be. While not necessarily negligible, the systematic error due to the uncertainty in the magnon mass (and how to incorporate it), around 3%, would not be a dominant source of systematic error.

5.19 Temperature extrapolation

At higher temperatures, the presence of magnons clearly impacts the temperature. In such cases, we can still obtain a reliable and low-uncertainty estimate of the temperature of the majority gas without magnons by extrapolating the temperature signal to the zero-magnon limit. The way we do this in practice involves two steps, shown in Fig. 5.39.

In most of our data sets, we obtain temperatures with varying numbers of magnons and with varying hold times after evaporation (in the case that magnons are created at an intermediate trap depth) or after magnon creation (in the case that magnons are created in a degenerate sample at the final trap depth). The first step in our extrapolation procedure is to select the hold times at which the temperature is constant and at which sufficient magnons are present for reliable temperature fits. These data can then be aggregated and the variation of temperature with magnon number fit to a line. The zero-magnon intercept and its uncertainty are taken to be the magnon-free estimate of the majority gas temperature and the uncertainty of this estimate.

We perform a similar process to estimate T/T_c of the majority gas, in general with different hold times, as T/T_c and T are not in general constant at the same hold times. In fact, T/T_c generally rises as atoms are lost, while T remains constant. We select hold times where T/T_c is roughly constant and minimal.

Along with our measurements of T and T/T_c , we also quote the number of atoms and quantities derived from it, such as the chemical potential μ . The atom number, however, decreases slowly with hold time, so there are a range of atom numbers we could reasonably

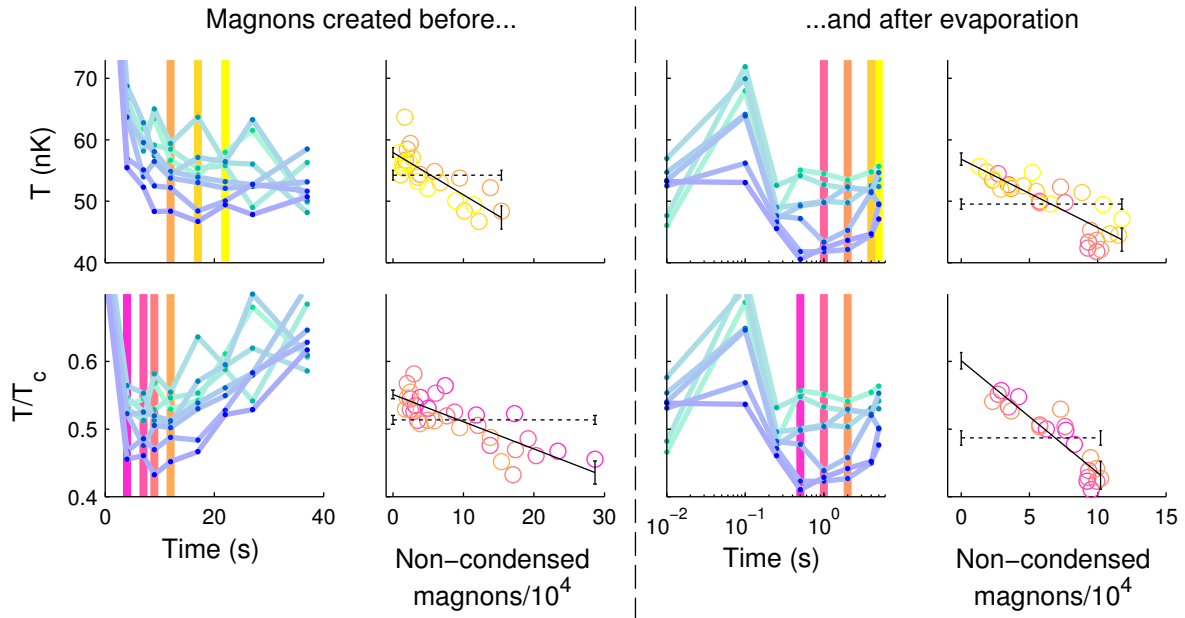


Figure 5.39: Extrapolating the temperature T and ratio of temperature to majority condensation temperature T_c to the zero magnon limit at a trap depth of $k_B \times 590$ nK. To obtain reliable, low-noise limits, we perform the extrapolation on data acquired at several hold times where T (T/T_c) is constant. The first (third) column shows T (T/T_c) as a function of hold time after the final evaporation (after magnon creation, with constant hold after evaporation). Different colored points and connecting lines indicate different initial magnon numbers. The vertical colored bars highlight the data used in the second (fourth) column, where T (T/T_c) is plotted versus the number of non-condensed magnons. The circle points are colored to indicate the hold time at which they were acquired, with colors matched to the bars in the first (third) columns. The linear fit to T (T/T_c), and uncertainties derived from the extrapolation, are indicated by the solid line and error bars at end points. The extrapolated T (T/T_c) and uncertainty is the left end point and error bar. The mean and standard deviation of the mean are indicated by the dashed black line and error bars on the end points. The linear extrapolation describes the data better than a simple mean.

state with each measurement. In general, when we quote a single number, we chose the largest measured atom number relevant to both the extrapolation of T and T_c . The range of choices for atom number is typically less than the systematic uncertainty we assign to our atom number estimates.

5.20 Optimizing cycled decoherence cooling

Several experimental parameters can affect the efficacy of cycled decoherence cooling, including the number of magnons created at each cycle, the thermalization time following magnon creation, and the magnetic field gradient. These parameters might not be constant from cycle to cycle. For example, as temperatures are lowered, the number of magnons created should drop with the critical number for magnon condensation. We optimized cycled decoherence cooling several cycles at a time by varying the parameters and plotting the resulting T/T_c and majority atom number achieved, as in Fig. 5.40. The best parameters are those where the slope of T/T_c versus majority atom number is steepest with the additional cycles. For simplicity, we did not optimize the gradient degree of freedom, but rather chose a gradient sufficiently small that the magnon box was effectively flat even at the lowest temperatures reached by decoherence cooling.

In the end, we found that fairly simple heuristics could be applied to roughly optimize cycled decoherence cooling without searching the entire parameter space. At each cycle of decoherence cooling, the number of magnons created should be small such that no magnon condensate persists upon thermalization. The thermalization time should be as short as possible such that only a negligible magnon condensate persists. In practice, we optimize cooling four to ten cycles at a time, breaking cooling into several stages with different parameters.

5.21 Two types of cooling

When we initially undertook our exploration of magnon decoherence cooling, we did not foresee the phenomenon of magnon-assisted evaporation. Our original protocols for both cooling and thermometry only included producing magnons at the final trap depth in the degenerate sample. It was only upon investigating heating rates in the trap that I suspected an additional cooling effect due to the evaporation of magnons.

The heating experiment involved performing cycled decoherence cooling to a sufficiently low temperature that evaporative cooling would cease. Magnons produced in the final step of cooling were retained for a variable number of seconds in order to probe the temperature of the gas as it heated in its trap. Surprisingly, I found the temperature to level-off at a lower value than the initial value reached by (non-magnon-assisted) evaporation. I hypothesized that the magnons improved the power of evaporative cooling, and our experiments in which we produce magnons at an intermediate trap depth proved that to be the case.

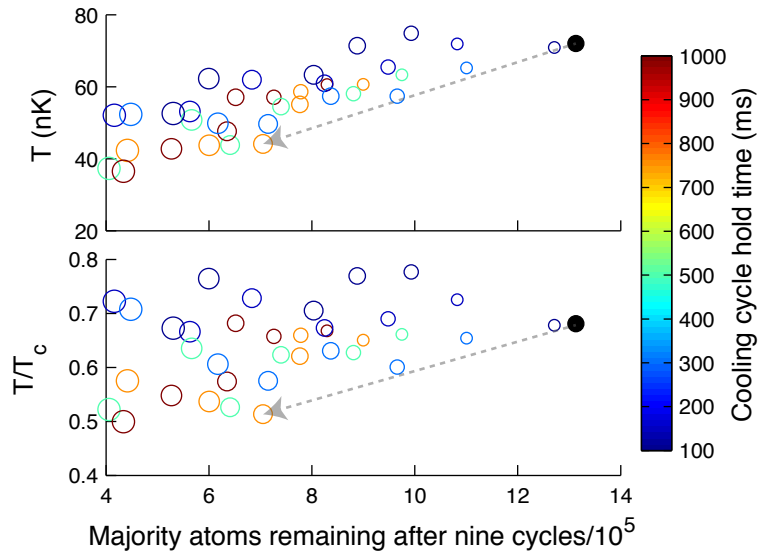


Figure 5.40: Optimizing cycled decoherence cooling at a trap depth of $k_B \times 810$ nK after nine cycles of cooling. The solid black point is the T (top) or T/T_c (bottom) achieved after one cycle of magnon creation and thermalization. The sizes of the open circles indicate the length of the RF pulse used to create magnons in each cycle, with longer RF pulses indicated as larger circles. Longer RF pulses convert more atoms to magnons in each cycle, leading to fewer majority atoms, but also potentially more cooling. Data points are color coded by the thermalization time employed in each cycle. Short thermalization times can be effective when the RF pulses are short. Longer thermalization times are needed when the RF pulses are longer. An optimal parameter choice would be to use the steepest slope, indicated by the dotted gray arrows. To continue the optimization, parameters are selected for the first nine cycles and new parameters are tried in subsequent cycles, generating a similar plot.

Initially, the discovery of magnon-assisted evaporation challenged our straightforward interpretation of our previously acquired magnon-decoherence data. Could the cooling we observed just be magnon-assisted evaporation turning on with the introduction of magnons? To tease apart the two types of cooling, we combined them, as shown in Fig. 5.41a. In this experiment, we evaporate with a large number of non-degenerate magnons present. Upon waiting some time at the final trap depth, we purge the non-degenerate magnons and begin cycled magnon decoherence cooling wherein the number of magnons created each cycle is less than the number of magnons that survived evaporation. As a result the power of any magnon-assisted evaporation is reduced. Regardless, decoherence cooling proceeds to lower the temperature as expected.

When magnons are first created at the final trap depth, it is likely that, initially, some magnon-assisted evaporation does occur along with magnon decoherence cooling. This is evident in Fig. 5.41b, which plots the energy removed per magnon as indicated by the slopes of the temperature extrapolation described in Sec. 5.19, wherein single pulses with varying

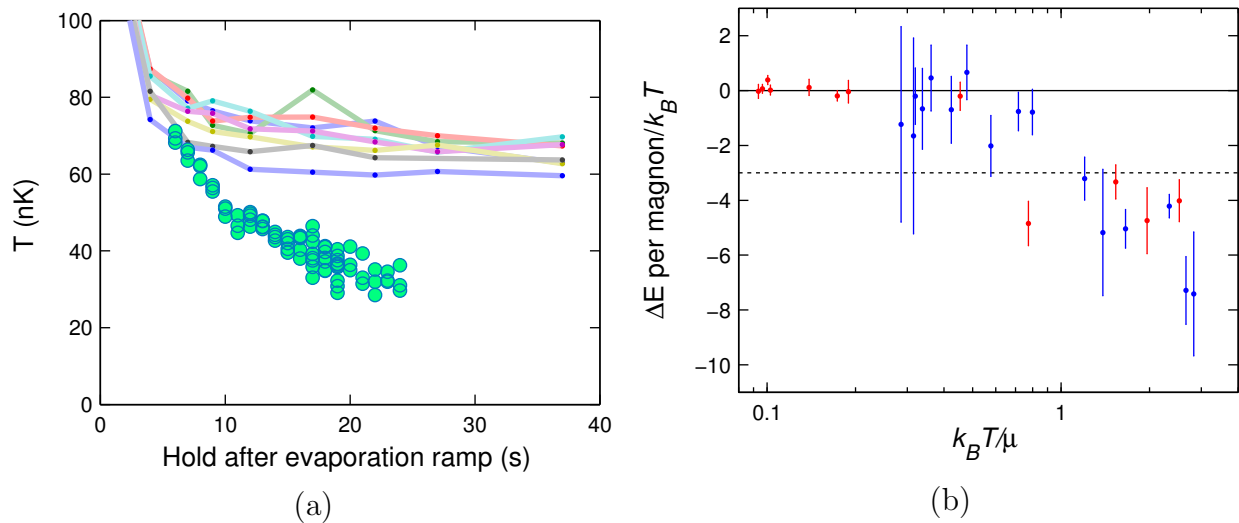


Figure 5.41: (a) Combining magnon-assisted evaporation and cycled decoherence cooling. Evaporation is performed with magnons present. Trajectories in which a variable number magnons are held at the final trap depth (points with lines between them) are compared to trajectories in which these magnons are removed from the trap and cycled decoherence cooling is applied using a smaller number of magnons than had been removed (filled circular points). Cycled decoherence cooling continues to cool the sample despite fewer magnons being present to assist in evaporation. (b) The energy removed per magnon for a single cycle of magnons created before (red) and after (blue) the final evaporation ramp in units of $k_B T$ are derived from the slopes and uncertainty of the extrapolated temperature versus magnon number. Decoherence cooling applies only to the blue data, while magnon-assisted evaporation applies primarily to the red data. Additional magnon-assisted evaporation likely also applies to the blue data, especially when $k_B T > \mu$.

numbers of magnons perform varying amounts of decoherence cooling or magnon-assisted evaporation. At higher temperatures, the energy removed per magnon in magnon-assisted evaporation is on the order of $4\text{--}6 k_B T$. At similar temperatures, the energy removed per magnon when magnons are created in the degenerate sample rises to $6\text{--}8 k_B T$, consistent with a $3 k_B T$ contribution from decoherence cooling in addition to some contribution from magnon-assisted evaporation.

5.22 Benefits of a square well potential

The fact that magnons disperse like free massive particles, as opposed to harmonically trapped or massless particles, can substantially benefit the signal to noise achievable by the magnon thermometer. This is shown dramatically in Fig. 5.42, where it is shown that the number of non-condensed magnons from which we derive our temperature signal at very




Figure 5.42: The number of non-condensed magnons estimated from fits over a wide range of data from our data set is compared the critical number for magnon condensation as computed for magnons confined to a harmonic trap with trapping frequencies derived from our optical trap and for magnons confined to a square well with volume equal to the calculated volume of the majority condensate under the Thomas-Fermi approximation. Runs with magnons created before the final evaporation ramp are plotted as dots while runs with magnons created at the final trap depth are plotted as open circles. Color varies with the ratio of the measured temperature to the calculated chemical potential of the majority gas.

low temperatures is more than an order of magnitude more than the maximum number of non-condensed particles that we would expect to exist in the harmonically trapped majority gas. Rather, the maximum number of non-condensed magnons scales very much like the prediction for particles confined in a flat box.

In the present work, we do not have as high confidence in our estimates of the non-condensed magnon number as we do the temperature. Unlike our estimates of temperature, the estimates of non-condensed magnon number can be shifted systematically by the incorrect Bose function in the fits of the transverse momentum distribution, as discussed in Sec. 5.15.2. We are exploring in better detail the characteristics of magnon condensation and thermal position/momentum distributions in ongoing work.

5.23 Thermometry and cooling results

Details of the results of our experiments and analysis are covered in Ref. [62], so here we include just a brief summary. Fig. 5.43 shows the temperature and level of degeneracy reached by evaporative cooling only, magnon decoherence cooling, and magnon assisted

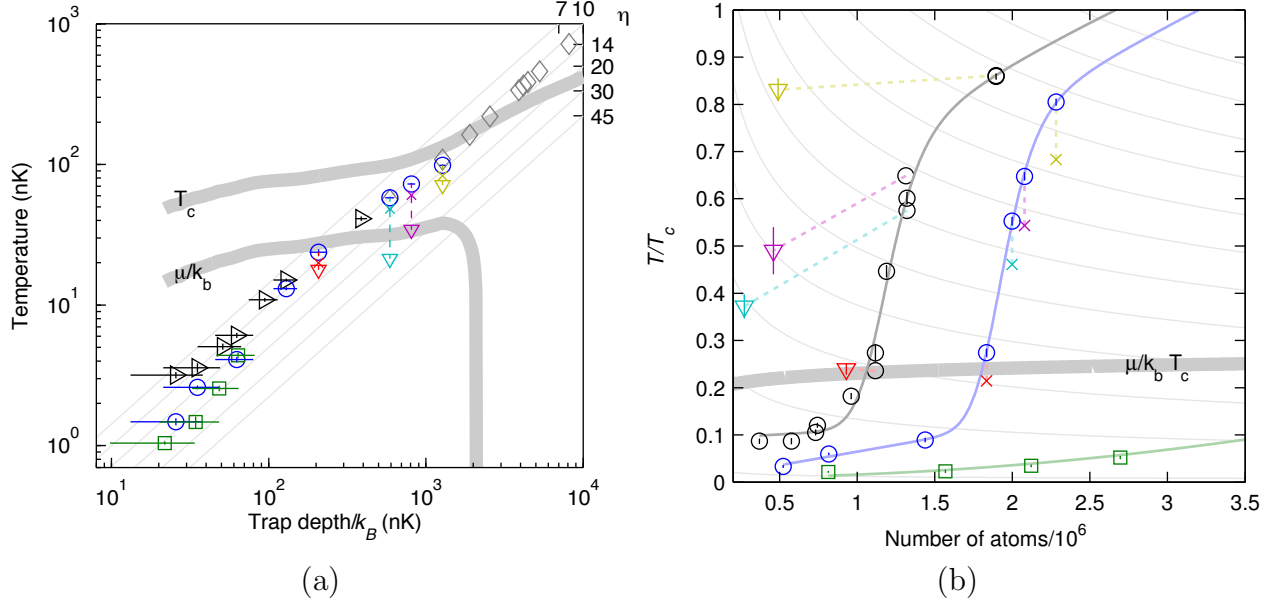


Figure 5.43: Thermometry and cooling results. Three runs (blue circles, green squares, black sideways triangles) differ in initial atom number (about 2.5, 3, and 5 million atoms at $T = T_c$, respectively). Magnons in black run were created at the final trap depth, blue and green at an intermediate trap depth. (a) T and (b) T/T_c are measured at various optical trap depths. Lower T and T/T_c are achieved in runs with larger initial atom number. T and T/T_c achieved by cycled decoherence cooling and magnon assisted evaporation are shown by colored downward triangle and cross (“x”) points, respectively, with the color coding the final trap depth in each of four runs. Dotted lines connect corresponding points at the same trap depth. In (a), thermometry using the majority spin population (light gray diamonds) agrees with thermalized magnon thermometry extrapolated to the zero-magnon values (circles, squares, triangles). Error bars show statistical uncertainty of extrapolated temperature (vertical) and systematic uncertainty in trap depth (horizontal). Thin diagonal gray lines show contours of η , the ratio of trapping potential depth to temperature. Thick gray lines show calculated T_c and μ/k_B in (a) and $\mu/k_B T_c$, calculated using the frequencies of the $k_B \times 60$ nK deep trap, in (b). (b) Efficiency of evaporation only (circles, squares, sideways triangles), with error bars showing statistical uncertainty of the extrapolated T/T_c , is compared to magnon-assisted evaporation (crosses) and cycled decoherence cooling (downward triangles). Solid lines connecting points are guides to the eye, with a steeper slope indicating more efficient evaporation. A variation of parts of this figure appeared in our publication, Ref. [62].




Figure 5.44: Close up of lowest temperatures, plotted against the effective trap depth, the trapping potential depth minus the chemical potential. Plotted this way, the evaporative cooling efficiency is seen to decrease slightly, though the decrease is not statistically significant, and different runs with different numbers of atoms (the different colors here) essentially collapse to a single trajectory.

evaporation. The key takeaways illustrated in this figure are as follows.

- Temperatures measured with magnons can be extrapolated to the zero magnon limit and agree with temperatures measured by conventional thermometers in the regime where conventional thermometers are applicable.
- Evaporative cooling is extremely effective at reducing temperature and entropy, even in the regime $T < \mu/k_B$. We measure temperatures as low as $T = 1.04(3)_{\text{stat}}(7)_{\text{sys}}$ nK, corresponding to $T/T_c = 0.022(1)(2)$ where T_c is calculated using the measured atom number, $N = 8.1 \times 10^5$, and optical trap frequencies. For this lowest-temperature gas, $k_B T/\mu \approx 0.07$. Its entropy per particle, as estimated in Ch. 6, is $S/N = 1 \times 10^{-3} k_B$, the lowest value ever reported for an atomic gas by two orders of magnitude.
- The evaporative cooling efficiency is typically characterized by the parameter η , which is the ratio of trap depth to temperature. In the regime $T < \mu/k_B$, one should consider the effective η , η_{eff} , the ratio of effective trap depth to the temperature. The effective trap depth is the difference between the trapping potential depth and the chemical potential. As shown in Fig. 5.44, η_{eff} is essentially constant over the entire range of our data.
- In the regime $T \gtrsim \mu/k_B$, magnon assisted evaporation yields colder, more highly degenerate samples without compromising the number of majority atoms at the final trap depth. When $T \lesssim \mu/k_B$ we cannot discern an additional cooling effect of magnon assisted evaporation.

- Cycled decoherence cooling allows low temperatures and entropy to be reached in a deep trap. We reach an temperatures as low as $T \approx 22$ nK, corresponding to $\eta \sim 30$, in a $k_B \times 700$ nK deep trap.
- The efficiency of cycled decoherence cooling, in terms of the number of atoms sacrificed to reach a desired temperature or entropy, is much lower than that of evaporative cooling. This makes sense, as each atom lost by evaporation carries away an energy of roughly $\eta k_B T$ while each atom lost in decoherence cooling carries away roughly $3 k_B T$ plus the chemical potential.
- In our experiments, decoherence cooling saturates at temperatures $T \sim \mu/k_B$ in spite of our initial expectations that decoherence cooling should be most effective in this regime because the gas heat capacity drops quickly as it becomes increasingly degenerate. We suspect the discrepancy lies in the slow speed of thermalization in this regime, which results in the magnons being unable to remove energy faster than it is deposited by heating processes.
- Magnon decoherence cooling can only barely reduce the entropy at $T/T_c \sim 0.9$. This is consistent with our expectation that the entropy should only be reduced by the decoherence cooling process for $T/T_c < 0.96$.

To derive this limit, consider that in an ideal Bose gas in a harmonic trap with geometric-mean frequency ω , the following relations for the saturated thermal atom number N_{th} , energy E , and heat capacity C , at the temperature T , apply:

$$N_{\text{th}} = \zeta(3) \left(\frac{k_B T}{\hbar \omega} \right)^3 \quad (5.21a)$$

$$E = 3 \zeta(4) \frac{(k_B T)^4}{(\hbar \omega)^3} = \frac{3 \zeta(4)}{\zeta(3)} k_B T N_{\text{th}} \quad (5.21b)$$

$$C = \frac{dE}{dT} = 12 k_B \frac{\zeta(4)}{\zeta(3)} N_{\text{th}}. \quad (5.21c)$$

Suppose we transfer $dN > 0$ atoms to the minority spin state (where dN is far below the critical number for magnon condensation), let them thermalize, and eject them. To lowest order, the majority gas loses energy $dE = -3k_B T dN$, and thus the temperature changes by $dT = dE/C = -3(k_B T)/C dN$. The non-condensed fraction changes as,

using Eqs. 5.21,

$$d\left(\frac{N_{\text{th}}}{N}\right) = \frac{1}{N} \left[dN_{\text{th}} + \frac{N_{\text{th}}}{N} dN \right] \quad (5.22a)$$

$$= \frac{1}{N} \left[\frac{dN_{\text{th}}}{dT} dT + \frac{N_{\text{th}}}{N} dN \right] \quad (5.22b)$$

$$= \frac{1}{N} \left[\frac{3N_{\text{th}}}{T} \frac{-3k_B T}{C} dN + \frac{N_{\text{th}}}{N} dN \right] \quad (5.22c)$$

$$= \frac{dN}{N} \left[\frac{3N_{\text{th}}}{T} (-3k_B T) \frac{\zeta(3)}{12k_B \zeta(4) N_{\text{th}}} + \frac{N_{\text{th}}}{N} \right] \quad (5.22d)$$

$$= \frac{dN}{N} \left[-\frac{3\zeta(3)}{4\zeta(4)} + \frac{N_{\text{th}}}{N} \right]. \quad (5.22e)$$

Thus, we must have $N_{\text{th}}/N < (3\zeta(3))/(4\zeta(4)) = 0.83$ to reduce the non-condensed fraction through such cooling. This corresponds to, for the case under consideration, $T/T_c < (0.83)^{1/3} = 0.94$.

In reality, the experimental upper limit of T/T_c may be somewhat less than 0.94 in cycled decoherence cooling as additional condensate atoms are lost in the magnon purge process.

Fig. 5.45 shows the results of our analysis of magnon decoherence cooling. The upshot of this analysis is as follows.

- At higher temperatures, when $T > \mu/k_B$, decoherence cooling works essentially as expected, with each magnon carrying away an energy of roughly $3k_B T$.
- At lower trap depths, the energy carried away by each magnon decreases, as indicated by the temperature change and using the heat capacity of a non-interacting gas. There are several effects that explain this, and our data do not allow us to easily untangle them. First, when $T \ll \mu/k_B$ the magnons are in an effective “box” potential and their thermal energy is expected to average only $3/2k_B T$ per particle. Second, the condensate heat capacity in this regime becomes that of a phonon gas, which drops more slowly with temperature than the non-interacting gas. We do not know of a good theory for the heat capacity in this crossover regime¹⁴. Finally, neither of these effects account for the quenching of decoherence cooling, which, as already mentioned, is likely related to slow thermalization of the gas.

¹⁴If we know the entropy in this regime, we can derive the heat capacity. The following chapter details calculations of the entropy at very lowest temperatures we encountered, 1–3 nK. Doing similar calculations at higher temperatures, in the 10–20 nK crossover regime, may be possible, though they would require accurately enumerating a much larger number of states of the trap, which may prove quite difficult.

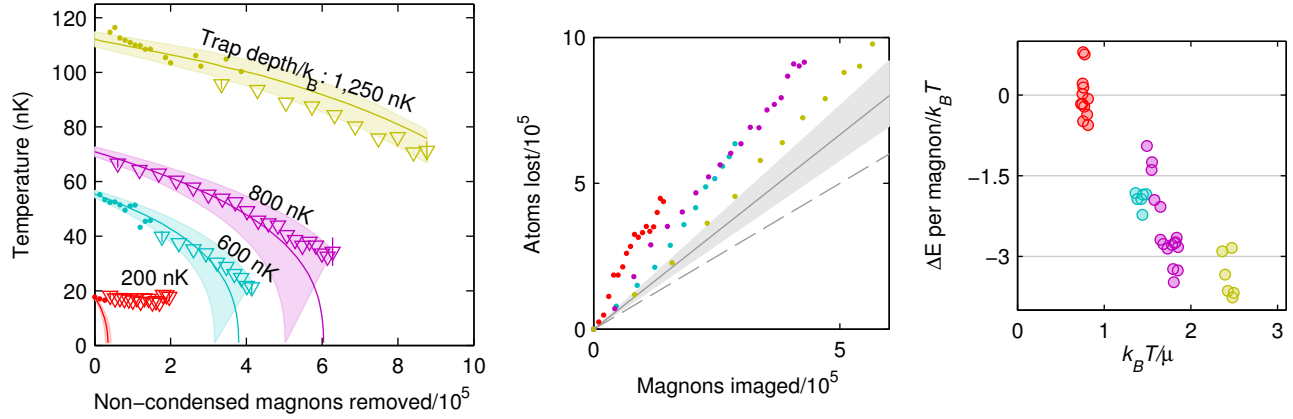


Figure 5.45: A closer look at decoherence cooling. (Left) Magnons are created at the final trap depth and cool the gas as they thermalize. Thermalized magnons can be removed from the trap and the cooling process repeated before measuring the temperature by imaging the momentum distribution of the thermalized magnons. Each non-condensed magnon removed from the trap takes energy from the gas. Here, decoherence cooling trajectories are plotted versus cumulative number of magnons removed at several trap depths, coded by colors that correspond to those of Fig. 5.43. Closed circles show temperatures after a single cycle of decoherence, but with varying numbers of magnons. Some additional magnon-assisted evaporative cooling may also be present. Open triangles show repeated cycles of magnon creation, thermalization, and purge, with representative error bars on first and last triangles giving statistical error over several repetitions. Solid lines show zero-free-parameter theory predictions assuming each non-condensed magnon removes $3 k_B T$ energy. Solid patches indicate the range of predictions included within the uncertainty in the cumulative number of magnons removed. (Center) We image $75\% \pm 10\%$ of the magnons present after each cycle of decoherence cooling. From these magnons, we extract the temperature. In a separate run, we also extract the total number of atoms remaining in the trap. Here we compare the number of atoms lost along each cycled cooling trajectory to the cumulative number of magnons imaged. If atom loss were due only to removing magnons, the number of atoms lost would lie in the gray region, which accounts for the $\sim 75\%$ efficiency of our magnon imaging sequence. Instead, the cumulative number of atoms lost in each decoherence cooling trajectory is larger than this owing to additional atom losses unrelated to decoherence cooling, such as finite trap lifetime and 3-body loss. The dotted line represents 100% magnon imaging efficiency and no additional atom loss. (Right) The amount of energy removed by each magnon, in units of $k_B T$, is plotted versus the ratio of $\mu/k_B T$, where each point is based on an estimate of the slope over 4 consecutive triangle points in the left frame. We observe that the net energy carried away by the magnons vanishes when $T \lesssim \mu/k_B$. A variation of parts of this figure appeared in our publication, Ref. [62].

Chapter 6

Calculating the entropy per particle

The entropy per particle of a Bose-Einstein condensate deep in the phonon regime, where $T \ll \mu/k_B$, can be calculated using the standard relations for Planck-distributed phonons, Eq. (5.3), provided one can make the local density approximation. For the gas of Ch. 5 at the lowest entropy settings, the chemical potential is $\mu \approx k_B \times 15$ nK, the speed of sound is $c = \sqrt{\mu/m}$, and density is $n = \mu \times m / (4\pi\hbar^2 a)$, with a the scattering length, and this calculation yields $S/N \approx 1.5 \times 10^{-5}$ at the point of highest density at the center of our trap at $T = 1.04$ nK. This naive calculation underestimates the entropy in our gas.

We cannot rely on the local density approximation to calculate the entropy per particle because the thermal phonon wavelength $\lambda = hc/k_B T \approx 50 \mu\text{m}$ is far larger than the transverse Thomas-Fermi radii of the condensed gas ($15 \mu\text{m}, 2.5 \mu\text{m}$). To understand the entropy, we will have to estimate numerically the spectrum of excitations along these dimensions. However, since the condensate radius along the longest axis is large (around $300 \mu\text{m}$), we can apply a one-dimensional local density approximation to determine the entropy. In this chapter I detail and build confidence in this important calculation.

6.1 Numerical objective

In general, the entropy S and number of atoms N in a Bose gas at equilibrium temperature T and with chemical potential μ can be calculated as [69]

$$S/k_B = \sum_i \frac{(\epsilon_i - \mu)/k_B T}{e^{(\epsilon_i - \mu)/k_B T} - 1} - \log(1 - e^{(\mu - \epsilon_i)/k_B T}) \quad (6.1a)$$

$$N = \sum_i \frac{1}{e^{(\epsilon_i - \mu)/k_B T} - 1}, \quad (6.1b)$$

where the many-body excitation spectrum of the gas is $\{\epsilon_i\}$. For example, plugging in the spectrum of single-particle excitations in a harmonic trap one can derive the well-known formula for the entropy of a Bose gas in a harmonic trap. These formula are accurate even

in the presence of interactions, provided one accounts for interactions properly in deriving the $\{\epsilon_i\}$.

The excitation spectrum of a weakly interacting (with s-wave interactions) trapped Bose gas is typically well described by the zero-temperature Gross-Pitaevskii theory [69], wherein the ground state order parameter Ψ_0 satisfies the time-independent Gross-Pitaevskii equation,

$$(T + V(\mathbf{x}) - \mu + g |\Psi_0(\mathbf{x})|^2) \Psi_0(\mathbf{x}) = 0, \quad (6.2)$$

with the kinetic energy operator $T = -\hbar^2 \nabla^2 / 2m$, confining potential V , and the interaction parameter $g = 4\pi \hbar^2 a / m$. The condensate density is $n = |\Psi_0|^2$. Note that in the limit $g \rightarrow 0$ Eq. (6.2) reduces to the time-independent Schrödinger equation for a single particle, with μ the energy eigenvalue. Also, the solution Ψ_0 is not necessarily unique; $\Psi_0 = 0$ is always a trivial solution.

For a particular static solution $\Psi_0(\mathbf{x})$, the small amplitude excitations can be written

$$\Psi(\mathbf{x}, t) = \left(\Psi_0(\mathbf{r}) + \sum_i [u_i(\mathbf{r}) e^{-i\omega_i t} + v_i^*(\mathbf{r}) e^{i\omega_i t}] \right) e^{-i\mu t / \hbar},$$

where “*” denotes complex conjugation, and the Bogoliubov amplitudes u and v satisfy the coupled eigenvalue equations

$$\begin{aligned} \hbar\omega_i u_i(\mathbf{r}) &= (T + V - \mu + 2g |\Psi_0(\mathbf{r})|^2) u_i(\mathbf{r}) + g \Psi_0(\mathbf{r})^2 v_i(\mathbf{r}) \\ -\hbar\omega_i v_i(\mathbf{r}) &= (T + V - \mu + 2g |\Psi_0(\mathbf{r})|^2) v_i(\mathbf{r}) + g \Psi_0^*(\mathbf{r})^2 u_i(\mathbf{r}). \end{aligned} \quad (6.3)$$

Once again, in the limit $g \rightarrow 0$ these equations are equivalent to the Schrödinger equation with eigenvalue $\hbar\omega_i + \mu$. The excitation energies relevant to Eqs. (6.1) are thus $\epsilon_i = \hbar\omega_i + \mu$.

This is our challenge: derive the spectrum of excitations of our gas in our trap, and employ Eqs. (6.1) to calculate S/N . Our approach is to numerically solve Eq. (6.2) for the ground state order parameter Ψ_0 and then to plug this into the eigenvalue equations Eqs. (6.3) to pull out the excitation spectrum. To begin, we need a way to discretize the relevant equations such that the solutions we obtain are sufficiently accurate and the numerics tractable.

6.2 Discretization strategy

The naïve way of discretizing involves choosing a Cartesian grid $\{\mathbf{x}_i\}$ with spacing Δx between points, on which one will define the spatially varying functions $\Psi_i = \Psi(\mathbf{x}_i)$ and $V_i = V(\mathbf{x}_i)$. The kinetic energy operator can then be approximated using finite-differences. To do this, one defines the one-dimensional difference operator ∂_x ,

$$\partial_x \Psi_j = \Psi(x_j + \Delta x / 2) - \Psi(x_j - \Delta x / 2),$$

so that

$$\partial_x^2 \Psi_j = \Psi_{j+1} - 2\Psi_j + \Psi_{j-1},$$

and relates it to the one-dimensional derivative operator ∇_x via the translation operator,

$$\Psi(x_j + \Delta x/2) = e^{\Delta x \nabla_x / 2} \Psi(x_j).$$

The difference operator can then be written

$$\partial_x = e^{\Delta x \nabla_x / 2} - e^{-\Delta x \nabla_x / 2} = 2 \sinh(\Delta x \nabla_x / 2),$$

allowing one to solve and expand to find

$$\nabla_x^2 = (\partial_x^2 - \partial_x^4/12 + \partial_x^6/90 + \mathcal{O}(\partial_x^8)) / (\Delta x)^2.$$

This expression allows the kinetic energy operator T to be written as a matrix to arbitrary accuracy (if one employs an arbitrary number of mesh points). The Gross-Pitaevskii equation can then be written in a discrete form as

$$\sum_i (T_{ij} + V_i \delta_{ij} + g |\Psi_i|^2 \delta_{ij}) \Psi_i = \mu \Psi_j. \quad (6.4)$$

Formulated in this way, it is not obvious what assumptions and approximations are made in the discretization, and what their impact might have on the numerics. It should be clear that the approximation to the derivative ∇_x in terms of a truncated series of finite differences ∂_x depends on the finite differences being small. More generally, recall that the Gross-Pitaevskii equation can be derived by minimizing the energy functional

$$E = \int \left(\frac{\hbar^2}{2m} |\nabla \Psi_0(\mathbf{x})|^2 + V(\mathbf{x}) |\Psi_0(\mathbf{x})|^2 + \frac{g}{2} |\Psi_0(\mathbf{x})|^4 \right) d\mathbf{x} \quad (6.5)$$

with respect to the conjugate order parameter Ψ_0^* , with the constraint that particle number is conserved. The accuracy of the finite difference method relates to the precision with which the integral in the energy functional (and also the particle number) can be approximated by a rectangle-rule sum:

$$\int f(x) dx \rightarrow \sum_i f(x_i) \Delta x.$$

The energy of a particular configuration $\{\Psi_i\}$ computed via the rectangle rule may overestimate or underestimate the true minimum energy. While with a small enough grid spacing this approximation can be very good, as we shall see, there are more nuanced discretizations that lead to a sum rule that achieves a higher degree of accuracy with fewer mesh points.

First, let's consider another common way to discretize an equation: in function space via a basis approximation. To discretize using N orthonormal basis functions $\{\phi_i\}$, we write

$$\Psi(\mathbf{x}) = \sum_i \Psi_i \phi_i(\mathbf{x}). \quad (6.6)$$

For convenience, we will once again consider a one-dimensional problem. We plug Eq. (6.6) into Eq. (6.2), multiply on the left by $\phi_j^*(x)$, and integrate to obtain the expression

$$\int dx \phi_j^*(x) \times \left[\sum_i \left(T + V(x) + g \left| \sum_k \Psi_k \phi_k(x) \right|^2 \right) \Psi_i \phi_i(x) = \mu \sum_i \Psi_i \phi_i(x) \right], \quad (6.7)$$

which we write more concisely as

$$\sum_i (T_{ij} + V_{ij} + gI_{ij}) \Psi_i = \mu \Psi_j. \quad (6.8)$$

The matrices T and V only need to be computed once (for time-independent problems), however the non-linear interaction matrix I needs to be computed for any change to the weights Ψ_i . Thus, solving Eq. (6.8) using an iterative procedure like Newton's method requires N multiplications of a different $N \times N$ matrix at every step, as with the finite difference approach. However, in the present basis function approach, calculating each $N \times N$ matrix requires an explicit numerical integration:

$$I_{ij} = \int dx \phi_j^*(x) \phi_i(x) \left| \sum_k \Psi_k \phi_k(x) \right|^2. \quad (6.9)$$

This integration can be performed to arbitrary precision, regardless of the number of basis functions employed.

An advantage of the basis function approach over the finite difference approach is that the solution $\Psi(x)$ is defined at every point x , and thus the energy of the configuration computed via Eq. (6.5) constitutes of a lower bound of the true ground state energy, as in a variational calculation. In addition, with the proper choice of basis, a relatively small number of basis functions may be required to achieve high accuracy. Regardless, it appears that the computational requirements—many many numerical integrations of Eq. (6.9) at every iteration—are quite high. The basis approach allows us to understand explicitly the assumed interaction energy and order parameter in all space, but at the cost of making local information cumbersome to calculate.

It turns out that the integral Eq. (6.9) can be calculated quite painlessly and accurately with the appropriate choice of basis functions, allowing us to combine the best of the finite difference and basis function approach. Suppose we choose orthonormal basis functions $\{f_i\}$ and mesh points $\{x_i\}$ such that

$$f_i(x_j) = \lambda_i^{-1/2} \delta_{ij} \quad (6.10)$$

and write

$$\Psi(hx) = \sum_i \Psi_i \lambda^{1/2} f_i(x),$$

where $\Psi_i \equiv \Psi(hx_i)$ and h is a scale factor that allows us to fit the range of $\{x_i\}$ to the region in which we want to find Ψ (i.e. hx is dimensionless). Then, remarkably, integrals involving Ψ can be computed very accurately via the quadrature rule

$$\int dx F(x) \approx \sum_i \lambda_i F(x_i). \quad (6.11)$$

In fact, this quadrature rule gives exact results for inner products of functions proportional to the f_i , such as the order parameter. With this basis, many of the virtues of finite difference return, for example the potential energy and the interaction energy are well approximated as diagonal matrices,

$$\begin{aligned} V_{ij} &= \int dx f_j^*(x) V(hx) f_i(x) \approx \sum_k \lambda_k f_j^*(x_k) V(hx_k) f_i(x_k) = V(hx_i) \delta_{ij} \equiv V_i \delta_{ij} \\ I_{ij} &\approx \sum_k \lambda_k f_j^*(x_k) \left| \sum_l \Psi_l f_l(x_k) \right|^2 f_i(x_k) = |\Psi_i|^2 \delta_{ij}, \end{aligned} \quad (6.12)$$

and while the kinetic energy matrix T is clearly not diagonal, we will chose functions $\{f_i\}$ such that it can be easily computed exactly. Note that with the scale factor, the condensate density is integrated as

$$n_0 = \int dx' |\Psi(x')|^2 = \int dx h |\Psi hx|^2 \approx \sum_i \lambda_i h |\Psi_i|^2. \quad (6.13)$$

With this straight-forward construction, called the Lagrange mesh method [6, 7]—the class of functions $\{f_i\}$, which vanish at all mesh points except one, are called generally called Lagrange functions— the discrete Gross-Pitaevskii equation reads

$$\sum_i (T_{ij}/h^2 + V_i \delta_{ij} + g |\Psi_i|^2 \delta_{ij}) \lambda_i^{1/2} \Psi_i = \mu \lambda_j^{1/2} \Psi_j, \quad (6.14)$$

where the scale factor h is needed to scale the kinetic energy to match the scaled order parameter. Once we have chosen our Lagrange functions, the equations we need to solve are no more complicated than those of Eq. (6.4), but we have gained some important features.

First, like in the finite difference case, we only need to keep track of the value of the order parameter $\Phi(hx)$ at the mesh points $x \in \{x_i\}$. However, in contrast to the finite-difference approach, the order parameter implied by the weights $\{\Psi_i\}$ is defined everywhere. As a result, the kinetic energy can be calculated exactly. In addition, the Lagrange mesh comes with a quadrature rule that permits a wide range of relevant integrals, including the interaction energy, to be approximated with a high degree of accuracy using fewer mesh points. Thus, the solutions to Eq. (6.14) are closely related to a variational solution to Eq. (6.5) on the chosen basis. The Lagrange-mesh method is closely related to a broader class of discretization schemes known as “discrete variable representations,” about which much has been written [42].




Figure 6.1: The Hermite basis Lagrange functions f_i (colored curves) with 10 mesh points and quadrature weights $\{\lambda_i\}$ (black points) shown above their mesh points. As desired, only one basis function is non-zero at each mesh point. The mesh points are more closely spaced near the origin. Such spacing is well suited to problems where the kinetic energy is largest near the origin.

6.3 The Hermite mesh

A key part of the success of the numerical methods we employ in this section is the selection of the right basis functions, and in the Lagrange mesh approach, their accompanying grid points. Different problems, with different boundary conditions, coordinate systems, and potentials will demand different bases. For the anisotropic harmonic oscillator, a mesh based on the Hermite polynomials with Gaussian weighting is particularly well suited. Recall that orthogonal polynomials are typically orthogonal with respect to a particular weighted inner product. For the Hermite polynomials $\{H_n\}$,

$$\int_{-\infty}^{\infty} dx H_m(x) H_n(x) e^{-x^2} = \sqrt{\pi} 2^n n! \delta_{nm}.$$

For $g = 0$ and with a harmonic potential, the solutions to Eq. (6.2) are in fact Hermite polynomials with Gaussian weighting, and these basis functions should constitute a perfect variational foundation.

The general procedure for constructing the Lagrange functions and mesh points from a basis of orthogonal polynomials can be found in the Refs. [6] and [7], including the particular case of the Hermite mesh. The procedure is straightforward, with the N mesh points $\{x_i\}$ corresponding to the zeros of the N th Hermite polynomial and the basis functions constructed from particular prescribed linear combinations of the remaining N basis functions:

$$f_j(x) = \lambda_j^{1/2} \sum_{i=0}^{N-1} \chi_i^*(x_j) \chi_i(x) = \frac{\chi_N(x)}{(x - x_j) \chi'_N(x_j)}, \quad (6.15)$$

with $\chi_i(x)$ the i th normalized, exponentially weighted Hermite polynomial. Care must be taken with the second, compact form for f_j in Eq. (6.15) because the singular denominator at $x = x_j$ can cause issues in finite-precision calculations. The quadrature weights are given simply by

$$\lambda_i = 2 / (\chi'_N(x_i))^2.$$

The 10 Lagrange interpolating functions $\{f_i\}$ corresponding to a mesh with 10 points are shown along with the weights $\{\lambda_i\}$ in Fig. 6.1. Finally, the kinetic energy matrix for the Hermite mesh can be written explicitly and exactly as

$$\frac{2m}{\hbar^2} T_{ij} = \begin{cases} (4N - 1 - 2x_i^2)/6 & i = j \\ (-1)^{i-j} (2(x_i - x_j)^{-2} - 1/2) & i \neq j \end{cases}. \quad (6.16)$$

If this seems like a lot of work to create a grid for a calculation, don't be fooled. The only results we actually require for our calculation are the mesh points, weights, and kinetic energy matrix. The interpolating functions are only needed if we want to understand the character of our solution in between, or outside, the mesh points.

6.4 Solving a simple 1D problem

To illustrate the method and to check our numerics, we solve Eq. (6.14) in one dimension with very low atomic density and with harmonic confinement, in which case interactions should be negligible and the resulting energies and excitation vectors should be those of a quantum harmonic oscillator. Solved with the constraint of constant (low) density, the ground state energy μ should be simply $\hbar\omega/2$. As written, Eq. (6.14) is non-linear matrix equation for Ψ_0 where μ is specified. In order to allow μ to be a free parameter and to constrain the atom number n (in the 1D case, the atom number is really a 1D linear density), we add an additional equation:

$$\sum_i h\lambda_i |\Psi_0|^2 - n = 0. \quad (6.17)$$

We then solve the system of equations via Newton's method. In practice, this is achieved by using *Mathematica's* `FindRoot` or *MATLAB's* `fzero`.

Once the ground state Ψ_0 and energy μ are determined, we plug them into Eq. (6.3) along with the matrices T and V . The resulting equation is linear in the u and v , and so its eigenvalues and eigenvectors can be computed via standard exact or approximate methods. Note that u and v must be solved for simultaneously, and thus the matrix equation must combine both lines of Eq. (6.3) and is thus twice as large in each dimension as the grid on which u , v , and Ψ_0 are discretized. In spite of the great deal of structure present in this matrix—its eigenvalues come in positive/negative pairs and it is anti-symmetric, for example—there do not seem to be extraordinarily efficient exact or approximate methods for its diagonalization, as there are for Hermitian and real-symmetric matrices.

The choices we have to make in solving for the Ψ_0 and μ are two: we must choose the number of grid points N and the scaling h of the grid. For convenience, we calculate the excitations and excitation energies on the same grid as we use in solving for the ground state. However, it is possible to choose different N and h for these two problems, as the value of Ψ_0 is determined everywhere by the interpolating Lagrange functions. Choosing a new grid for the eigenvalue problem would simply involve resampling Ψ_0 at the new grid points with high accuracy using the functions $\{f_i\}$.

To develop a sense of the role of the grid scaling h , we solve the problem using several values and compare the estimated excitation energies to their exact values in each case. In calculating the entropy, our primary concern is to compute system's energy levels accurately up to several times higher than the thermal energy. The spatial extent of the N th harmonic oscillator wave function scales as $N^{1/2}$, so to be able to accurately describe high (compared to N) energy excitations the mesh should grow accordingly. Thus, we scale the mesh such that $hx_N = sx_{\text{ho}}\sqrt{N}$, where x_N is the N th (largest) mesh point and $x_{\text{ho}} = \sqrt{\hbar/m\omega}$ is the harmonic oscillator length. The energy and excitation vectors are shown for $N = 50$ and three values of s in Fig. 6.2.

With the proper scaling, it is not surprising that most of the energies and excitation vectors are reproduced exactly¹. With $s = 1.5$ the spectrum is very accurate out past the 40th (of 50) eigenvalue. From Fig. 6.2b, we can derive a useful heuristic for finding a grid scaling: On one hand, the grid should be large enough in extent to fully contain the eigenvectors. The discretization forces the excitation functions to zero just beyond the edge of the grid, and if the exact solutions are not zero there, the energy implied by our interpolation will differ from the exact solution. If the mesh is too small, increasing the number of sample points will be of limited use in gaining additional accuracy. On the other hand, making the grid too large in extent will mean that the mesh does not have sufficient resolution to describe the excitation vectors. In such a case, adding more sample points (with the mesh extent constant) should benefit the calculation.

6.5 Higher dimensions

Solving problems in higher dimensions is relatively straightforward. In understanding the various options for simulating in higher dimensions, I found the thesis of McPeake [55] to be helpful. Here, we will continue to focus on Cartesian coordinates and the Hermite mesh, but bear in mind that for problems with cylindrical or spherical symmetry, other coordinate systems and meshes based on other Lagrange functions may be more suitable. To extend our current approach into two dimensions, we construct two one-dimensional meshes, perhaps with different numbers of mesh points, N_x and N_y , and different scalings, h_x and h_y , and combine them. The coordinate grid is a straight-forward outer product of the two 1D meshes, in the usual Cartesian style.

¹Indeed, for some scaling, the basis is constructed of the exact solutions, and the exact energies should be recovered.

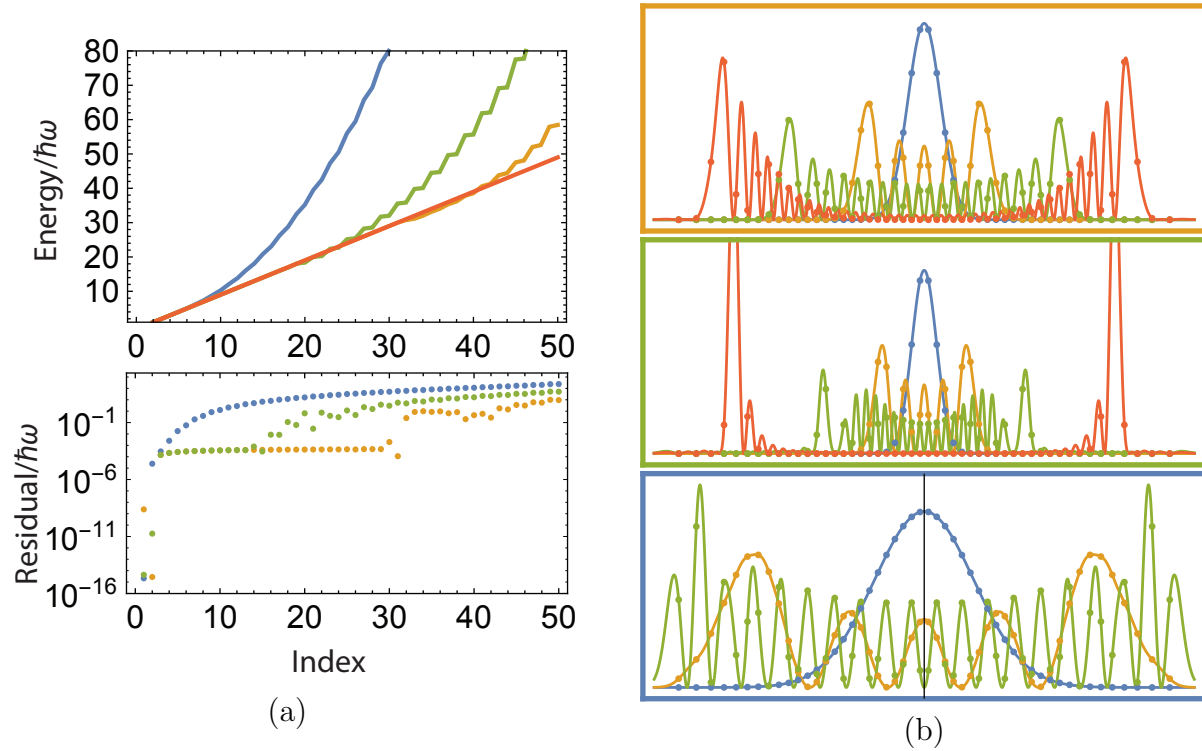


Figure 6.2: (a) The energy of excitations in a harmonic trap, in units of the harmonic oscillator energy $\hbar\omega$, are taken from the zero and positive eigenvalues of Eq. (6.3) discretized with $N = 50$ mesh points and scale h such that $hx_N = sx_{\text{ho}}\sqrt{N}$, where x_N is the N th (largest) mesh point, $x_{\text{ho}} = \sqrt{\hbar/m\omega}$ is the harmonic oscillator length, and $s = \{0.5, 1.5, 2\}$ for the blue, orange, and green curves and points, respectively. The red line is the exact harmonic oscillator solution. The residuals show the difference between the calculated and exact results. (b) The probability densities corresponding the ground state (blue), 5th (orange), 20th (green), and 42nd (red) excitations, calculated as $|u(x) + v^*(x)|^2$, for each of the scalings. The frame border colors are coded to the mesh scaling, in accord with (a). Lagrange interpolating functions are used to infer the value of the excitation functions in between grid points. The x -axes have different scalings; the blue ground state curve is the same width in all cases. The top panel has a well scaled mesh and captures the details of the exact excitation vectors well. The benefits of a well chosen basis are evident in the way features between grid points are well described. In the second panel, the mesh is too large and detail is lost in the center of the mesh. Wiggles are missing or are not of full height in the green and red curves. In the third panel, the mesh is too small.




Figure 6.3: A sample kinetic energy matrix for 1D 8 point mesh (left) and 2D 8×8 mesh (right). Matrix elements that are identically zero are left white. The 2D matrix is quite sparse.

The most natural way to formulate our equation is probably to consider our state vector as a rank-2 tensor and our Hamiltonian as rank-4 tensors, or something like that. In practice, though, we need to flatten everything into the same sort of matrix equation that we had in the 1D case, so that we can compute eigenvalues and solve the appropriate system of equations using the usual well-worn tools. There is not a unique way to do this, but there are lots of ways that don't work. I made a few errors in my first attempt at setting up the right equations in 2D, so for the sake of posterity I present my explicit construction.

First, we assign a linear index to each Cartesian point. To do this, we construct a mapping from linear index i to Cartesian indices (α, β) . I will refer to this map a few times, so let's give it a name, $g(i)$, $g : i \rightarrow (\alpha, \beta)$. This part is very straightforward; just number the grid points going down the columns first, for example. With this map, we can associate each index of our state vector Ψ_i with a Cartesian point (x_α, y_β) . Easy so far.

Constructing the diagonal interaction and potential energy matrices I and V is similarly straightforward, for example $V_{ii} = V(x_\alpha, y_\beta)$. Constructing the appropriate linearly indexed form of a non-diagonal matrix, such as the kinetic energy matrix, is a little more tricky. First we construct the appropriate 1D kinetic energy matrices, $T^{(x)}$ and $T^{(y)}$. From these, the linearly indexed 2D kinetic energy matrix can be generated, with $g(i) = (\alpha, \beta)$ and $g(j) = (\kappa, \lambda)$, as

$$T_{ij} = \frac{h^2}{h_x^2} T_{\alpha,\kappa}^{(x)} \delta_{\beta,\lambda} + \frac{h^2}{h_y^2} T_{\beta,\lambda}^{(y)} \delta_{\alpha,\kappa},$$

where $h = h_x h_y$ is the 2D scaling parameter. The state vector has $N_x N_y$ entries, and the kinetic energy matrix is $N_x N_y \times N_x N_y$, which means that for larger meshes the matrices get large. Fortunately, T is quite sparse, which can substantially speed up calculations and reduce memory usage. With the relevant matrices and state vector constructed, solving the 2D problem is exactly like solving the 1D problem, with one caveat. If the size of the mesh is

large, finding the exact eigenvalues of the resulting matrix Eq. (6.3) might be prohibitively slow. In this case, the desired number of approximate eigenvalues can be found via iterative methods.

6.6 Comparing to the Thomas-Fermi limit in 2D

The Gross-Pitaevskii ground state can be computed trivially in the Thomas-Fermi limit, and in the case of 2D harmonic confinement with cylindrical symmetry, the excitations can be computed analytically [28]. With ω the transverse trap frequency, the collective excitations of the trapped sample are given by

$$\epsilon_{nm} = \hbar\omega\sqrt{2n^2 + 2n|m| + 2n + |m|}, \quad n \geq 0 \in \mathbb{Z}, \quad m \in \mathbb{Z}.$$

To test the accuracy of our numerics, and to further understand the key parameters involved, we use our numerical procedure to compute the ground state and spectrum for a harmonically trapped system with transverse confinement, deep in the Thomas-Fermi regime, with $\mu/\omega \approx 250$. We use $N_x = N_y = 50$ points for each dimension of the mesh (2,500 total mesh points) and set the scaling parameters h_x and h_y such that the largest mesh point sits at 1.4 times the Thomas-Fermi radius. Cuts through the center of the resulting ground state density and excitation densities are shown in Fig. 6.4 along with a comparison of the calculated and exact energy spectrum.

The basis functions are clearly not as ideally suited to describing the Thomas-Fermi distribution as they are to the harmonic oscillator distribution, as small additional wiggles can be seen in the interpolation. The mesh points themselves fall perfectly on a Thomas-Fermi profile, which suggests that the ripples seen in the interpolation are related to the quadrature rule imperfectly approximating the integrated interaction and potential energies. In essence, the quadrature approximation seems to result in the ripples being ignored by the potential and interaction energy terms. The energy of the ripples is included in the kinetic energy of the ground state, but in the Thomas-Fermi limit the kinetic energy is negligible compared to the interaction energy.

Once again, we see good agreement with the expected spectrum result as long as the excitation vector amplitudes vanish before they reach the edge of the mesh. The excitations show a little bit of additional ripple between mesh points as well. The ripples evidently have negligible effect on the energy.

We can transform this 2D problem into a 3D problem easily, as long as we are happy with the third dimension remaining unconfined. To do this, we promote our 2D ground state solution to 3D by the trivial transformation $\Psi_0(x, y, z) = \Psi_0(x, y)$, that is, without any dependence on the third dimension. Next, we adjust slightly the ansatz from which Eq. (6.3) is derived, writing the excitations u and v as

$$u(x, y) \rightarrow u(x, y)e^{-ikz}, \quad v(x, y) \rightarrow v(x, y)e^{-ikz}.$$

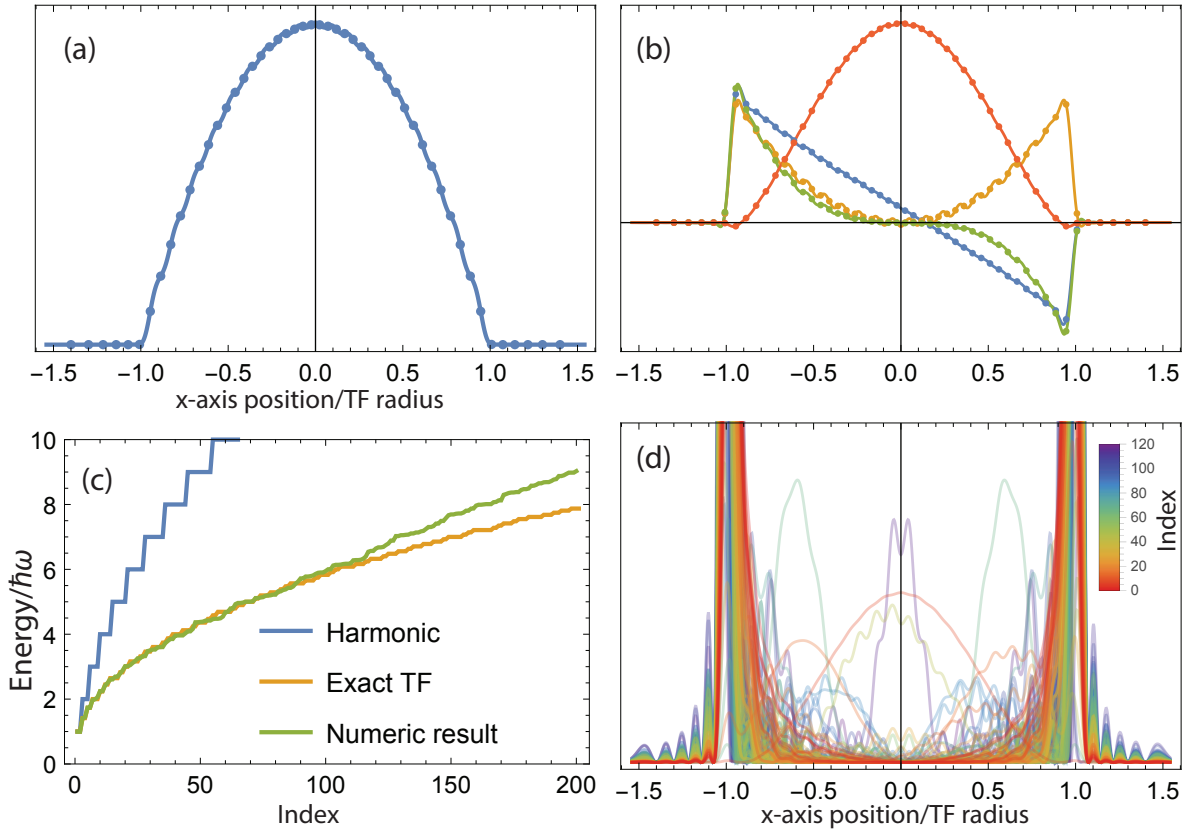


Figure 6.4: The Thomas-Fermi limit with transverse harmonic confinement. (a) Section of the ground state density $|\Psi_0|^2$ through the trap center. (b) Slices through the trap center of a few low lying excitations, shown as $|\Psi|^2 - |\Psi_0|^2$, where $\Psi = \Psi_0 + u + v^*$, including the center of mass mode (blue), scissors mode (green), and two quadrupole modes (red and gold). The excitations plotted in this manner illustrate the motion of the whole gas under each excitation. (c) The computed excitation spectrum matches the exact result for the first 100 lowest lying excitations. The chemical potential has been subtracted from the energy of the Thomas-Fermi and computed excitations. (d) Slices of the first 120 excitation vectors are shown as $|u + v^*|^2$, which highlights the probability density of the excitation. The excitations are generally concentrated in regions of low or zero condensate density. As in the harmonic case, the energy of the excitations deviates when the excitation vectors retain some amplitude near the edge of the mesh (blue curves are plotted behind red curves).

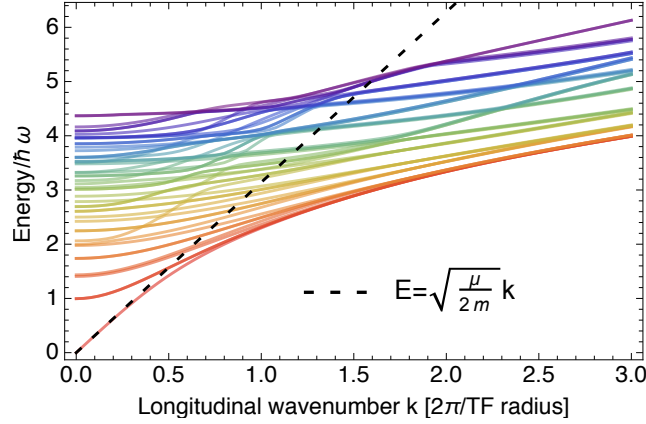


Figure 6.5: Spectrum of the cylindrical Thomas-Fermi condensate in three dimensions showing the lowest 50 transverse excitations at each longitudinal wavenumber, color coded in order of energy. The initial slope of the pure longitudinal excitation is consistent with the expected phonon dispersion, with effective harmonic channel speed of sound $c = \sqrt{\mu/2m}$, until the phonon wavelength approaches the transverse diameter, at $k = 1/2$. Even at higher wavenumbers, crossings with the transverse excitations seem to be evident.

The resulting eigenvalue equation, which keeps terms to linear order in u and v , is the same as before, but with the substitution

$$\mu \rightarrow \mu - \hbar^2 k^2 / 2m. \quad (6.18)$$

We obtain the 3D excitation spectrum in the unconfined limit by computing the eigenvalues of the adjusted eigenvalue equation with a range of longitudinal wavenumbers k . The 3D spectrum of excitations for the cylindrical harmonically trapped gas in the Thomas-Fermi limit is shown in Fig. 6.5, where the expected speed of sound is seen in the low energy longitudinal excitations.

As we move into numerical territory without precise results to compare to, we can now understand how to be confident in our results. First, we should rely only on ranges of the energy spectrum where the excitation vectors are fully contained by the mesh and do not have such high frequency wiggles that the mesh cannot accommodate them. This later condition can be tested for by increasing the mesh density and looking for improved accuracy. Based on our experiences in two extreme limits—Thomas-Fermi and non-interacting—we will choose a scale such that the largest mesh point along either axis satisfies

$$x_{\max} = \max(1.4 x_{\text{TF}}, 1.5 \sqrt{N_x} x_{\text{ho}}), \quad (6.19)$$

with x_{TF} and x_{ho} the Thomas-Fermi radius and harmonic oscillator lengths, respectively, and N_x the number of mesh points. This heuristic for defining the mesh scale will allow us to find the ground state and excitation spectrum at a range of densities.




Figure 6.6: Slices through the center of the condensate density along x (left) and z (right). Our solution is in blue, while the Thomas-Fermi and harmonic ground state densities, with the same total number of atoms, are shown in yellow and green respectively. Horizontal axes are scaled by the Thomas-Fermi radius in the appropriate direction, x or z .

6.7 Simulating our (harmonic) trap

We are primarily interested in the spectrum of our gas at the lowest trap depth settings, where we achieved the lowest T/T_c . This trap has trap frequencies $(\omega_x, \omega_y, \omega_z) = 2\pi \times (18, 0.78, 109)$ rad/s and on average we achieved a temperature of 1.04 nK with 8.1×10^5 atoms. Clearly, no axis of our trap is unconfined. Can we simulate the full 3D trap?

Simulating the full 3D trap is computationally prohibitive for us. The weak axis trap frequency corresponds to a lowest energy of 0.03 nK. Using the 1D Thomas-Fermi energy spectrum,

$$\epsilon_n = \hbar\omega\sqrt{n(n+1)}/2, \quad n \geq 1 \in \mathbb{Z},$$

we can estimate that in order to compute energies up to ten times the temperature of interest, we would need to have nearly 500 accurate eigenvalues along the z axis alone. In the Thomas-Fermi limit, which is most valid along the z axis, this would require a mesh with a few thousand points along one axis. By contrast, we expect that 50 mesh points should be sufficient along x and 20 points more than sufficient along z . The full 3D calculation would then require on the order of a million mesh points. While not impossible to deal with, as the matrices involved would be quite sparse, the resulting calculations would take a very long time on a personal computer.

Instead, we will make use of Eq. (6.18) to turn a 2D solution into a quasi-3D solution, employing what is known as the local density approximation (LDA). Under the LDA, we assume that properties like the entropy and energy can be defined locally, and that global properties can be estimated by integrating slowly varying local ones. For this sort of description to make sense, the density of the gas should change slowly compared to the wavelength of the thermal excitations. At 1 nK, the thermal phonon wavelength at the center of our trap is about $50 \mu\text{m}$. This is larger than either of the x or z dimensions, but much shorter

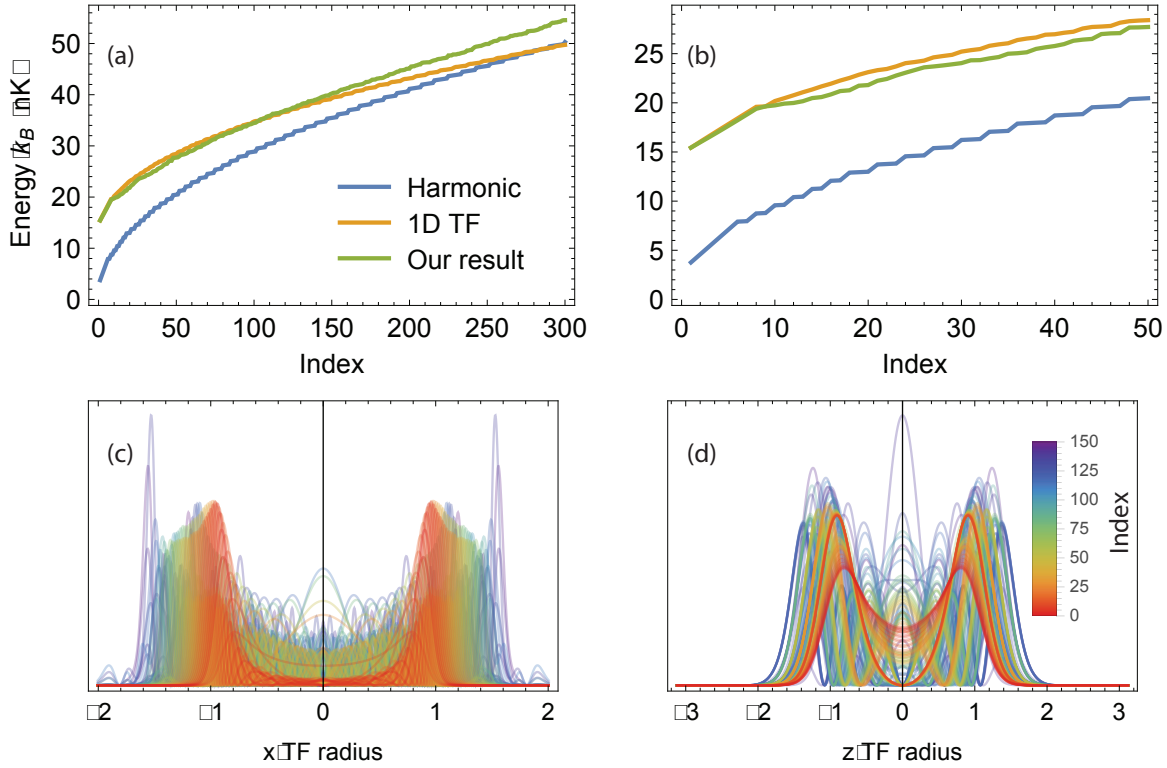


Figure 6.7: Spectrum of our trap, modeled as harmonic, with the longitudinal wavenumber $k = 0$. (a) The 300 lowest energy excitations are enumerated and compared to those of a 2D anisotropic harmonic oscillator and of two independent 1D Thomas-Fermi gases of the appropriate sizes. The chemical potential μ is included in the energy of the calculated and TF spectra in order to emphasize the interaction energy compared to the harmonic oscillator estimate. (b) Only the first 50 energies are shown. The lowest energies match the TF model exactly, as the lowest energies are excitations along x only. At higher energies, no agreement should be expected. (c,d) Slices of the eigenvectors, plotted as $|u + v^*|^2$, along the x and z axes, respectively.

than the weakly confined y dimension, which has a Thomas-Fermi radius of roughly $340 \mu\text{m}$. Near the edge of the trap, along y , the speed of sound drops and the thermal wavelength gets shorter. The thermal de Broglie wavelength for free particles is about $10 \mu\text{m}$. This situation lends itself well to the LDA.

We begin by looking at the spectrum of excitations at the (longitudinal) trap center. To do this, we use $N_x = 50$ and $N_z = 20$ mesh points along x and z , respectively. We set the chemical potential to the value predicted in the Thomas-Fermi limit, $\mu/k_B \approx 14.5 \text{ nK}$. The scaling factors employ the heuristic Eq. (6.19), which leads to scaling limited by the harmonic oscillator term along both axes. The 2D spectrum (longitudinal wavenumber $k = 0$) and eigenvectors are shown in Fig. 6.7. The eigenspectrum initially has a Thomas-Fermi-like character, with the excitation vectors concentrated at the edge of the gas and the energies

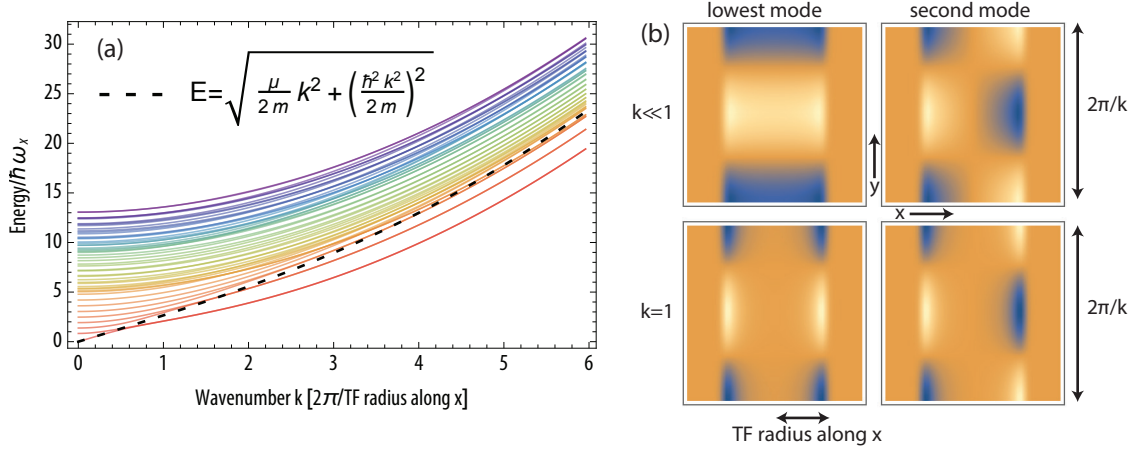


Figure 6.8: (a) Spectrum of our trap, modeled as harmonic, with the longitudinal wavenumber plotted in units of $2\pi/x_{\text{tf}}$, with x_{tf} the Thomas-Fermi radius along x . The dotted line shows the 1D Bogoliubov dispersion. As energy levels cross the Bogoliubov dispersion from above, they tend to become degenerate with other levels. (b) One such degeneracy is visualized by plotting a 2D cross section of the density excitation, plotted as $|\Psi|^2 - |\Psi_0|^2$, in the xy plane. Axes are not to scale. At low k , the lowest mode is a pure longitudinal excitation, with slightly higher amplitude at the edge of the gas than in the middle. The second mode energy is dominated by a dipole excitation along x that is concentrated near the edge of the gas. When k is higher, such that the longitudinal wavelength is on the order of the Thomas-Fermi radius along x , the longitudinal excitations are highly localized at the edges of the gas. The excitations confined at each edge are uncoupled, and can have any relative phase, leading to the degeneracy.

growing somewhat like a simple model that adds the energy of two independent 1D Thomas-Fermi gases. At higher energies, however, the energy grows more like that of a harmonic oscillator, and the excitation vectors look more like those of the harmonic oscillator. Only the first 50 or so eigenvalues will be needed to enumerate energies up to ten times the temperature (above the chemical potential).

The quasi-3D spectrum at the trap center is shown in Fig. 6.8. At low longitudinal wavenumber k , the lowest energy excitations scale with k as expected for phonons in a transverse harmonic channel. At higher values of k , the energy spectrum bends upwards in accord with the expected Bogoliubov dispersion. The spectrum has many interesting features. For example, as the longitudinal excitations become shorter in wavelength, lower energy transverse excitations form degenerate groups.

To calculate the entropy per particle, we first calculate the entropy per length. To do this, we select some scale length L and calculate the energy spectrum for each $k = j\pi/L$, $j \in \{0, \pm 1, \pm 2, \dots\}$, up to some cutoff k_{max} . The entropy S can then be calculated using Eq. (6.1a), where one must be careful to recall that the computed energy eigenvalues are already relative to the chemical potential, and the entropy per length is simply S/L . Note




Figure 6.9: Entropy at the longitudinal center of the trap at the lowest depth setting. Entropy per particle is plotted against (a) the maximum energy eigenvalue considered and (b) the temperature, with two different values of ϵ_{\min} (see text). In (a), $k_B \epsilon_{\min} \approx 0.1$ nK. Including eigenvalues up to ten times the temperature is more than sufficient to saturate the entropy estimate.

that the ground state is excluded from the sum in Eq. (6.1a). The entropy per particle can then be found by dividing S/L by the total 1D density $n = n_{\text{th}} + n_0$: $S/N = S/(Ln)$, where the non-condensed 1D density is $n_{\text{th}} = N/L$, with N given by Eq. (6.1b), once again with the ground state excluded, and the condensate density is given by Eq. (6.13).

The entropy per particle at the longitudinal trap center is shown in Fig. 6.9 as both a function of temperature and a function of the maximum energy included in the sum of Eq. (6.1a). The entropy largely saturates after energies roughly five times the temperature are included. In general we calculate energies up to ten times the temperature. This threshold energy dictates the number of eigenvalues we need to calculate at each value of k , as well as the maximum longitudinal wave vector k_{max} . In Fig. 6.9b, it is notable that in the regime of temperatures between 0.2 nK and 2.5 nK the entropy per particle grows like the square of the temperature. This may be due to the fact that the thermal excitations of the gas are effectively 2D. The lowest energy excitation along the tight z axis corresponds to roughly 5 nK, which is largely frozen out in this temperature range.

If one is not careful, the length L employed in the discretization of k can affect the estimate of the entropy. The lowest energy above the ground state that our calculation considers corresponds to a longitudinal excitation with wavenumber $k_{\min} = \pi/L$. In order for the entropy estimate to be independent of L , the variation in the relative occupations of the lowest lying states should be small. Equivalently, the energy of the lowest lying state, $\epsilon_{\min} = ck_{\min}$ in regions with a condensate, the c the longitudinal speed of sound, should be much less than the temperature, $k_B \epsilon_{\min} \ll T$. For our calculations, we use L such that $k_B \epsilon_{\min} \approx 0.1$ nK. As seen clearly in Fig. 6.9b, lowering ϵ_{\min} does not change the entropy estimate at 1 nK.




Figure 6.10: The calculated peak density in our trap as a function of μ is compared to the Thomas-Fermi and non-interacting ideal Bose gas predictions.

6.8 Entropy in our harmonic trap

We can now calculate the total entropy of our gas in a harmonic trap with the trap frequencies corresponding to our lowest trap depth settings, with $T = 1.04$ nK, by using the local density approximation and by assuming that the chemical potential varies in accordance with the Thomas-Fermi limit along the long y axis. To do this, we calculate the entropy and number of atoms as in Sec. 6.7 at a range of chemical potentials, yielding the entropy and number density as functions of μ : $\frac{S}{L}(\mu)$ and $\frac{N}{L}(\mu)$. We then write the total entropy of the gas as

$$S = \int_{-\infty}^{\infty} \frac{S}{L} \left(\mu_{\max} - \frac{1}{2} \omega_y^2 y^2 \right) dy, \quad (6.20)$$

and the total number of atoms as

$$N = \int_{-\infty}^{\infty} \frac{N}{L} \left(\mu_{\max} - \frac{1}{2} \omega_y^2 y^2 \right) dy, \quad (6.21)$$

where μ_{\max} is chosen to yield the experimentally measured value of N . The entropy per particle of the gas is then S/N . We need the Thomas-Fermi approximation along y in order to define the way the chemical potential varies with trapping potential. We do not need to rely on the Thomas-Fermi approximation along x and z , however, and the peak chemical potential μ_{\max} may not take on its Thomas-Fermi value. In fact, as shown in Fig. 6.10, for a given value of μ , the atom density is generally somewhat less than the Thomas-Fermi prediction. As μ increases, the Thomas-Fermi approximation becomes more valid.

The range of μ at which we calculate the number and entropy densities clearly needs to include μ_{\max} . In principle, there is no lower bound chemical potential implied by the limits of integration on Eqs. (6.20) and (6.21). In practice, as shown in the left hand panels of Fig. 6.11, both number density and entropy density fall off exponentially for $\mu < \hbar\omega_x/2 + \hbar\omega_z/2$, and thus the integrals can be truncated (or extrapolated) without consequence.

While the overall entropy per particle is significantly higher than the expectation for a non-interacting Bose gas—we calculate $S/N \approx 0.0012$ whereas with the same atom number

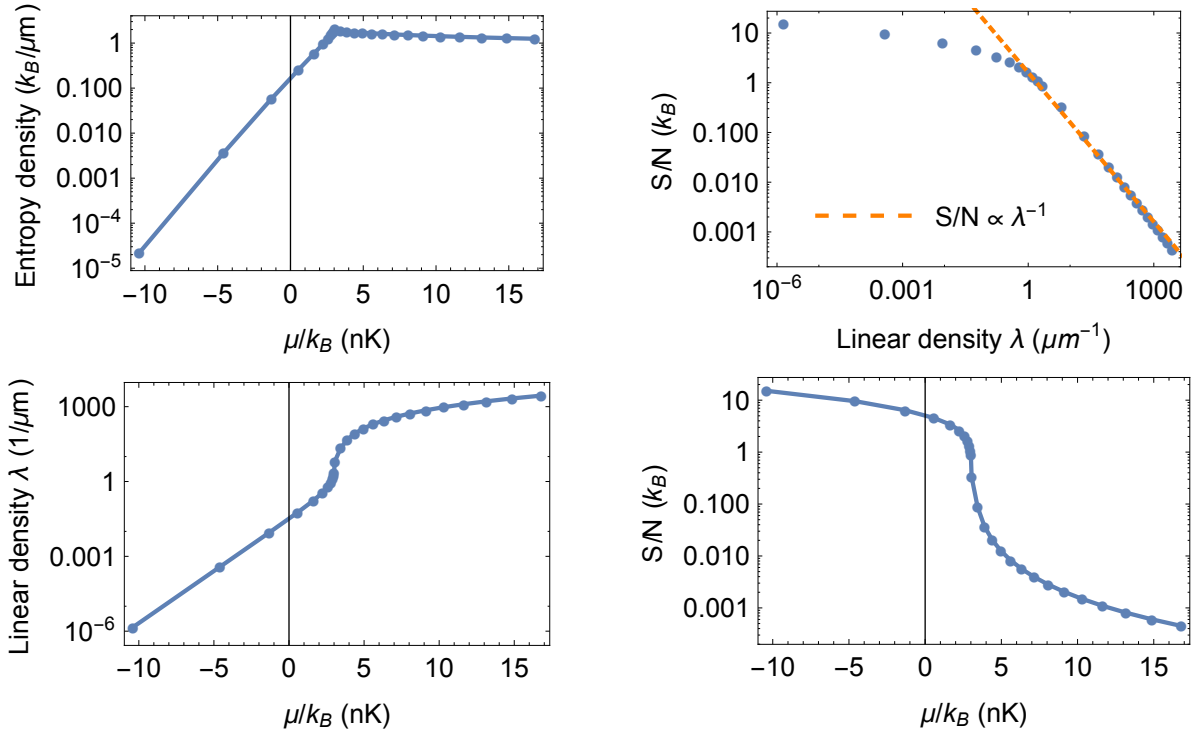


Figure 6.11: Entropy density, entropy per particle, and linear atom density in a harmonic trap with our trap frequencies. The gas becomes quantum degenerate for $\mu > \hbar\omega_x/2 + \hbar\omega_z/2 \approx k_B \times 3\text{nK}$. The entropy density is roughly constant in the condensed region, consistent with the expectation for an ideal Bose gas. The entropy and atom density both fall off exponentially at negative μ .

and temperature in the non-interacting limit we would expect $S/N \sim 0.00003$ —in many ways, the entropy behaves similarly in the non-interacting and interacting gases. For example, in the degenerate non-interacting gas, the entropy is independent of the total atom number, and thus the entropy per particle varies like $1/N$ (in Eq. (5.1) note that $T_c^3 \propto N$). In our interacting gas, the entropy per particle drops only slightly faster than $1/N$ in the region of degeneracy. Equivalently, entropy density is roughly constant as μ varies in the regime of degeneracy.

The overall entropy per particle and the entropy per particle at the one-dimensional (along y) trap center are within a factor of two of each other. As noted in Sec. 6.7, the entropy at the trap center scales with the square of the temperature around 1 nK. Indeed, we can verify that varying the temperature by 10% changes the overall entropy by 20%. Thus, our estimate of the entropy, to one significant digit, is not affected by the systematic and statistical error in our thermometer. It is likewise not affected by the atom number uncertainty at this level of precision, as the peak μ scales weakly with the atom number.




Figure 6.12: Line profiles of the excitation densities through the center of the trap, shown as $|u + v^*|^2$, with the profile of the trapping potential superimposed. The zero-level of the densities is shown roughly at the chemical potential level of the trapping potential. Along the x axis the potential is nearly harmonic. Along the z axis the trapping potential is not harmonic and many of the excitations are unconfined, with non-negligible amplitude beyond the lip of the trap.

6.9 The trap is not harmonic

We have calculated an estimate of the entropy per particle at equilibrium in a harmonic trap with trap frequencies corresponding to those in our system at the lowest trap depth. However, as shown clearly in Secs. 5.8 and 5.7, the optical trap deviates significantly from harmonicity in the region where the condensate is confined when the chemical potential is a considerable fraction of the trap depth, as it is at the settings we are considering. How might taking the anharmonicity of the trap into account affect our estimate of the entropy?

Estimating the entropy in the anharmonic trap is complicated by a few factors. First and foremost, the gravity-tilted optical trap is not confining and lacks a global minimum. Irrespective of this, by setting the initial guess for the ground state such that it has no amplitude beyond the Thomas-Fermi radius, Newton’s method has no problem finding the necessary meta-stable ground state Ψ_0 . The excitation spectrum, however, includes many states that are clearly not confined, as shown in Fig. 6.12. These states pose a problem because their energies are not reliable—the excitations are clearly not contained to the mesh—and regardless, atoms occupying these states would quickly leave the trap.

We will take two main approaches to work around this issue. One approach is to post-select excitations that do not have significant amplitude outside the confining region of the trap. Physically, this approach may be justified by noting that because these states couple strongly to the continuum, the lifetime of atoms in these states would be short, and thus at steady state they are likely to be less occupied than their energy (which we don’t trust anyway) might suggest. In taking this approach, we decide how much excitation density we permit to be unconfined, and reject energy levels that exceed this threshold. The metric

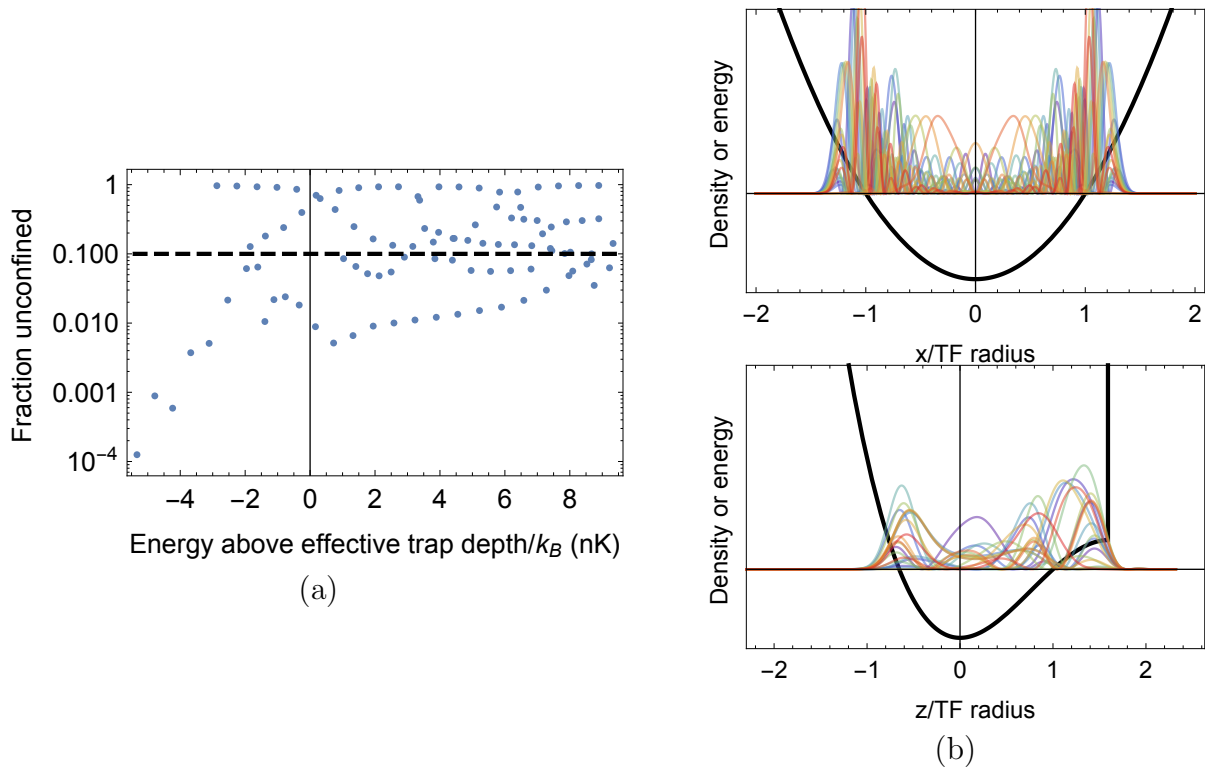


Figure 6.13: Two approaches to dealing with the lack of confinement in the anharmonic trap. (a) By computing a measure of the confinement (see text) associated with the calculated excitations, we can exclude from our entropy estimate states that fall below a certain threshold of confinement. Here we show the unconfined fraction f of the first 100 states at $\mu/k_B \approx 15$ nK and with longitudinal $k = 0$, plotted against the energy of the state referenced to the effective trap depth. Many states with energy larger than the effective trap depth are confined; such states are excited predominantly in the mostly-harmonic x direction. One of the thresholds that we use is indicated by the dashed line. (b) Alternatively, we can make the anharmonic trap confining by adding a wall at the trap lip. The line profiles shown are analogous to those of Fig. 6.12.

Approach	S/N
Harmonic trap	0.0012
Anharmonic wall	0.0012
Anharmonic $f = 0.1$	0.0012
Anharmonic $f = 0.02$	0.0011
Anharmonic $f = 0.01$	0.0011

Table 6.1: Summary of different entropy per particle estimates. There is not much variation.

we use to gauge the level of unconfinement is the fraction of the excitation density that is beyond the lip of the trap, in the continuum \mathcal{C} :

$$f = \frac{\int_{\vec{x} \in \mathcal{C}} d\vec{x} |u(\vec{x}) + v^*(\vec{x})|^2}{\int_{\vec{x}} d\vec{x} |u(\vec{x}) + v^*(\vec{x})|^2}.$$

The unconfined fraction f is shown for the first 100 excitations in Fig. 6.13a. The choice of f is somewhat arbitrary, but fortunately the outcome is not particularly sensitive to this choice.

Another approach is to restore confinement while preserving most of the anharmonicity of the trap. There are many ways to do this. One naïve and simple approach is illustrated in Fig. 6.13b: add a hard wall at the trap lip. The ground state in this trap is unchanged, but the excitations are confined. With this approach, we can calculate the entropy as we do with any other trap, with no postselection of the computed excited states. In either approach to calculating with anharmonicity, we continue to approximate the trap as harmonic along the long axis y and we assume that the trap geometry along x and z does not depend on the y coordinate. With this assumption in place, we need only vary μ to estimate the entropy density at various points along the y axis.

The entropy estimated by taking account anharmonicity in any of these ways is essentially unchanged over the harmonic case, as seen in Tab. 6.1. The entropy is dominated by the lowest energy excitations—longitudinal phonons and a few of the lowest lying transverse excitations. As shown in Fig. 6.14, the lowest energies are nearly identical in the harmonic and anharmonic traps. The entropy density is generally somewhat higher in the anharmonic trap than in the harmonic one, but not enough to make a large difference overall.

It is interesting and important to note the difference in entropy density between the two anharmonic approaches at high chemical potential. When using the postselection procedure, the entropy density drops rather quickly above a particular value of μ . This should make sense, as the zero point energy of the excitations is the chemical potential, and thus with larger μ the excitations need less energy to leave the trap. At larger μ , the effective trap depth is lower. At smaller μ , the effective trap depth increases and more excitations can be considered confined. Here we see the significance of our assumption that the transverse trap geometry, and hence the trap potential depth, is constant, regardless of the value of the




Figure 6.14: (Left) The energy spectra calculated in the anharmonic trap, with two different approaches to the confinement problem, are compared to the results in the harmonic trap. At low energy, the difference between traps is small. (Right) The entropy density is plotted against the chemical potential, for the same traps. The anharmonic trap has slightly higher entropy density in general. The postselection procedure causes the entropy density to drop when the the chemical potential is large.

longitudinal coordinate y . In fact, as seen clearly in Fig. 5.18, the trap depth is not constant along the y axis. In fact, the trap potential depth decreases away from the trap center. As a result, we are likely overestimating the entropy away from the trap center.

Why not vary the transverse trap potential with μ ? This is certainly possible, though because we do not know ahead of time the peak chemical potential μ_{max} that results in the correct total number of atoms, we may have to repeat the calculation a few times with different μ_{max} . But our first guess at μ_{max} is likely to be good, as the peak values of μ that we find generally don't deviate much from the Thomas-Fermi prediction. One could do this calculation, and it is very likely to find a slightly smaller estimate for the entropy.

We have not done this version of the calculation, though, because we would rather estimate the entropy as conservatively as is reasonable. It is not obvious how states that are coupled to the continuum should contribute to the entropy at steady state. Based on the high per-atom loss rate in our trap at the lowest trap depths, we know that at steady-state many atoms do occupy states that are coupled to the continuum, and these states are likely to have lower occupancy than the temperature might suggest. Fully accounting for the effect of continuous evaporation in the estimated entropy likely requires machinery that is beyond the scope of this work.

6.10 Adjustments to the chemical potential

In our work we generally report T_c and μ as calculated by standard formulas for Bose gases in the non-interacting and Thomas-Fermi limits, respectively. In particular,

$$k_B T_c = \hbar\omega \left(\frac{N}{\zeta(3)} \right)^{1/3} \quad (6.22)$$

$$\mu = \frac{\hbar\omega}{2} \left(\frac{15Na}{a_\omega} \right)^{2/5}, \quad (6.23)$$

where ω is the geometric mean trap frequency and $a_\omega = \sqrt{\hbar/(m\omega)}$ is the associated harmonic oscillator length. Several effects can shift the true values of T_c and μ . These simulations allow us to estimate some of them.

The value of μ is shifted by two main effects. In the Thomas-Fermi approximation, kinetic energy is neglected compared to the interaction energy. As a result, Eq. (6.23) underestimates the true value of μ for finite values of Na/a_ω [13]. In addition, the anharmonicity of the vertical trapping potential results in a larger trap volume than predicted by ω , skewing the calculated value of μ upward. Our calculation of the trap entropy via μ_{\max} takes into account both of these opposing effects and estimates the peak value of μ to be shifted upwards by barely more than 1% of the harmonic Thomas-Fermi value, with each effect contributing no more than a 5% shift when considered individually.

The value of T_c is likewise shifted by two main effects: trap anharmonicity and interactions. Both effects are expected to be small. Trap anharmonicity should reduce T_c , similarly to μ , by a few percent at the very lowest trap depths. At higher trap depths, the anharmonicity is reduced. Interactions are estimated to produce a fractional shift [21] in T_c of

$$\frac{\delta T_c}{T_c} = -1.32 \frac{a}{a_\omega} N^{1/6} \approx -0.02. \quad (6.24)$$

The resulting combined shift is on the order of the systematic uncertainty in our temperature measurements.

Chapter 7

The future

The future of our experimental apparatus is bright as a new crew of smart and motivated personnel have taken over as its stewards. It takes years to learn the ins and outs of an apparatus as complex (and haphazard?) as ours, but already my colleague Fang Fang has shown an exceptional ability to produce impressive and compelling data. The work she has lead, already mentioned as Ref. [18], to be published soon, continues our work with magnons by exploring the details of their condensation, or as it turns out quasi-condensation. In that work, we leverage our ability to fully image the gas magnetization to map the emergence of a magnon condensate phase as a function of temperature and other experimental parameters. As we expected from our other work with magnons, magnons do indeed condense in a box-like potential. In our most recent work, we can see a clear crossover from harmonic to box-like condensation as a function of the ratio of temperature to chemical potential.

The new members of our experimental team are already looking beyond magnons to an improved apparatus capable of producing molecules of rubidium and lithium, eventually in an optical lattice. This work involves reviving some of our early dual-species efforts and, likely, a major overhaul of our main vacuum chamber and the surrounding optics. I will not be an integral part of that effort, but I look forward to seeing many impressive results in the future.

Appendix A

Winding the slower

The procedure for winding contiguous (single-wire) layers of the slower was roughly as follows. The general idea is that positions along the slower are referenced to particular point, the home position, which we place at the end of the main section of the slower, in terms of turns of the chuck with the fourteen-per-inch threading gears engaged. All of the windings on a particular piece of stock are referenced to the same home position. Here a one inch outer diameter tube hosts the base layer of coils (1/8 inch wire) on which are wound the main section (1/16 inch wire, to the “left” of the home position) and the boost coils (1/8 inch wire, to the “right” of the home position). The stretch coils and the anti-gradient coils are on separate pieces of stock with a larger OD.

To begin winding a layer with a new conductor, the starting position for the first turn of wire is identified in terms of the number of fourteen-per-inch turns from the home position. The chuck and slower are then rotated the appropriate number of turns, with the threading gears engaged at fourteen-per-inch, to bring the winding tool into position. The starting position of the layer is then marked with an indelible felt-tip marker. To protect against shorts, the prior layer of wire is coated with a single layer of kapton tape, and to hold the new layer in position, the kapton tape is coated with a thin layer of thermally conductive epoxy. We used JB Weld.

Before affixing the wire, we consider how to achieve the appropriate winding pitch. The design of the slower necessitated turns of the thin wire with one-quarter density (shown in the winding diagram, Fig. 2.10, as one coil every four slots), which could not be made with the lathe directly. Such turns require some manual intervention; however appropriate positioning of the lathe tool (always moved by turning the chuck with gears engaged) is still important. In this case, rotate the chuck (no wire attached) to bring the winding tool to the point at which half- or single-density windings commence. Then, set the gears to the closest approximation of quarter-density that is available (we used half-density) and turn the chuck backwards a number of full rotations equal to the number of quarter-density turns. Next, affix the new wire lead, which passes through the winding tool, at the previously marked starting position with a padded hose clamp. At this point, the winding tool and the wire position on the slower do not line up precisely. Next, turn the chuck to lay down the




Figure A.1: Winding the slower. (a) An overview of the first half of the wound slower. One-quarter, one-half, three-quarter, and full density turns are clear. (b) Close-up of the a retaining ring used to step the number of layers. (c) The slower being wound on the lathe. The wire spool was supported (support not shown) and allowed to freely rotate to reduce tension on the wire. Tension could be controlled by the winding tool. (d) The winding tool, made of delrin, had a slot for both the 1/16 inch and 1/8 wire. The slot precisely positioned the wire on the slower and kept it properly oriented.

appropriate number of quarter-density turns. Using the underlying layers as a guide, adjust the spacing of the fresh-laid wire to approximate one-quarter density. The winding tool is presently in position to begin laying half- or single-density turns, and the wire and winding tool should line up precisely. Change gears as necessary and continue laying down turns according to the winding diagram, using the threading functionality of the lathe to guide the wire placement.

When quarter-density turns are not required, no shenanigans are needed and the wire can be affixed directly at the position of the winding tool when it is set to the starting position. If multiple layers are to be put on top of one another such that the number of layers jumps at the starting position by more than one, a retaining ring should be affixed gently at the starting position. Such a ring has a slot for the lead wire to pass through and abuts the first turn on the layer. When using such a ring, be careful with the wire insulation. Delrin or teflon rings are ideal. Excessive tension on the wire lead should be avoided. Set the gears

for the appropriate winding density and begin counting turns of the chuck.

Each run of wire in the main section of the slower includes at least one out-and-back. The taper in the number of turns makes getting leads out easy, as shown in Fig. A.1. A large retaining ring is fixed at the home position at the end of the main section to allow the wire at this position to stack. If all is done right, the left-to-right (low field to the left) layer should reach the stop at the home position after the correct number of turns. The direction of travel of the lathe tool is then reversed, and the next layer is placed in similar, but reverse, fashion. Of course, when laying wire, the chuck is always turned in the same direction; only the direction of travel of the lathe tool is changed.

After our first (few) attempts at winding the slower, we discovered shorts between some of the layers. The polyimide enamel coating of the small wire was not as robust as the kapton tape coating of the larger wire (even though kapton is made of polyimide). Shorts are avoided by wrapping a layer of kapton tape between layers and by maintaining the tension with which the wire is wound to just above the level where the wire is held in place and lays flat. Shorts are found by two techniques. Shorts between layers produced by separate conductors can be identified by finite conductivity between conductors where there should be none. Shorts between layers composed of the same conductor can be found by anomalous (low) resistance between ends of the conductor, or by abrupt changes in the measured magnetic field.

Appendix B

Chip fabrication proposal

An overview of the chip fabrication project goals, and its status when I assumed the lead role on the project, can be found in the following three pages. This document was developed as part of a request for proposal in an attempt to outsource the chip fabrication project to Teledyne Scientific. The document overstates slightly our capability (at the time) with respect to electroplating and planarization—in reality, we could only reliably achieve void-free electroplating of the larger, lower-aspect-ratio bias chips, and we had only attempted to remove excess copper material from chips with copper on one side—but it is a helpful overview of the project nonetheless.


Teledyne’s proposal came back indicating, essentially, that their capabilities with regard to fabrication roughly matched our own, and that they had little to no experience with the parts of the project that we had yet to figure out, namely the metalization, lapping and polishing of copper, and bonding of the final chip stack. For a few hundred thousand dollars, they would be willing to make a “best effort” attempt at developing a successful process and producing the desired product, but they could make no guarantees. The time frame they proposed, if I recall correctly, was roughly a year.

Introduction: For our new experiment where we want to create a magnetic ring trap for ultracold atoms, we need very precise magnetic fields, i.e. magnet coils. These can be achieved by microfabricated chip coils, controlling its dimensions to the required degree of precision on the order of one micrometer. For the specific magnetic field configuration, we require four distinct double-sided chip coils with very high aspect-ratio microfabricated wires (120 μm deep) to carry $\sim 4\text{A}$. Furthermore we would need some bonding/stacking of these 4 chip coils on a micron level, where two of these chip coils need to be bonded together, respectively and subsequently bonded with a 2mm thick thermally conductive spacer (sapphire? silicon?). However, we should differentiate between these two fabrication steps. In the first phase we are looking for getting a number (maybe 2-10 wafers) of these microfabricated chip coils cut into dies. We will try to perform the bonding/stacking ourselves in the lab with thermally conductive epoxy. However, both quality and quantity is probably improved if we can also outsource this step. Therefore, in parallel to our own effort, Teledyne should also work on this second phase, which involves die bonding (potentially with Au thermal compression bonding) and wire/lead attachments. But let us first define the requirements on the chips themselves.


Chip Coil Specifications: In general, each of these four different coils is a microstructured conductor on a substrate formed by two flat spirals connected through the chip substrate by a via. A schematic drawing of a single chip coil is shown here without the substrate. Leads coming in on the top surface from a soldering pad (not shown), spiraling inward in about 20 turns (counter-clockwise in this example), connected to the other side through the substrate by the via, spiraling outwards again (also counter-clockwise when viewed from the other side), and leading out in parallel to the incoming lead to a second solder pad on the bottom side of the substrate surface.

To introduce our terminology we like to call the smaller coils "Curvature Coils" and the larger ones "Antibias Coils". For one assembly we need two Curvature Coils of opposite helicity and two Antibias Coils of opposite helicity. Their dimensions are the following:

	Curvature Coils	Antibias Coils
Number of Turns	21	27
Wire Width [μm]	20	50
Wall Thickness [μm]	10	10
Trench Depth [μm]	120	120
ID [mm]	2	4.5
OD [mm]	3.3	7.9



The dies carrying these coils should have a thickness of $\sim 525 \mu\text{m}$. Their size is about $28.6 \times 8 \text{ mm}$. The exact dimensions and location of the spiral coils are depicted below. In the center of each coil we need a hole through the substrate for optical access. Furthermore the location of the soldering pads is off-axis to allow one corner to be cut out, which facilitates solder connections in the bonded final assembly. The location of alignment marks (chip-to-chip and stack-to-stack) is indicated as small squares. This is where we can have backside illumination with IR light for chip glueing in our lab.



This page should define the basic requirements we have on these chip coils. Hopefully without biasing you too much, on the next page, however, we would like to show you some examples of the chips that were made in the Berkeley Microlab by Tom Purdy and Dan Brooks from silicon with electroplated copper. But in principle our coils could be made with other materials, different processes, order of processes, etc. The finished chip coils have to withstand temperatures of $\sim 200\text{C}$ for several days (vacuum bakeout $\sim 30\text{C/h}$) and -80C (daily cycle from/to room temperature in $\sim 1\text{h}$) without mechanical failure, i.e. cracking and plastic deformation of the conductor spirals.

Masks: We do have masks for 4" wafers made by photo-sciences with mirrored, clear drawn features (material 5090CRSL, product code F10CS5090). There is one mask per Curvature/Antibias coil for holes that are etched all the way through (via, central hole, L-shaped cut to break off one corner plus alignment marks) and one mask each for patterning the features (spiral, leads, solder pads plus alignment marks). The "features masks" are for front + backside and opposite helicities are already integrated. A fifth mask was made with 50 x 50 um squares (10 um in between) for Au thermal compression bonding. The data format is DXF and GDSII. I will send along those files for you. Maybe you could use them in the fabrication, maybe we/you will have to draw new masks.

Deep Etching: First the holes were etched almost all the way through (-120um) before the actual features were etched down about 120um on each side. For the central hole only a 70 um rim was etched to avoid to much contamination. Etching has to be very accurate, because it leaves behind very high-aspect ratio features (walls between turns are 10 x 120um).

Central Hole: If you have the capabilities to drill (etch, machine) the central hole as one of the last steps in the fabrication process this could be advantageous. It would eliminate the copper layer on its sidewalls seen in our chips. The dimension and position of this hole are less critical. It is only for optical access.

Tolerances: What are reasonable tolerances you can achieve in the alignment of the front and backside masks? What will be your tolerances on the smoothness of the bottom of the etched trenches?

The via is offset from the last turn to accomodate its size (60 x 120 um). However, if you could make the via as slim as 20um it could be integrated in the last turn! This would allow us for instance to increase the size of the central hole and have a larger NA for optical access.











Insulation: Before backfilling the trenches with copper an insulating thermal oxide layer (1um) was put on. How can we specify and measure the leakage current between turns and from the top to the bottom spiral? Would it be possible to bias the substrate?

Plating: The growing (electroplating) process has to be done in a way that leaves no voids in the high-aspect ratio trenches and the vias. What could be the tolerances on defects (not visible on the surface) from the growing process (most likely reflecting the roughness of the etched trenches)? Can you analyze it by taking cross-sections or do you have other means?


Adhesion: To improve adhesion the leads were spread out in parallel and additional "anchors" around the soldering pad were implemented. Do you have other ideas on how to improve adhesion? (How) can we specify it?

Lapping/polishing: Can you do double-sided wafer polishing? Or would you have to cut the dies first and then do the polishing? Please note the drastic differences in local densities (areas of very high and very low copper). So far Tom has achieved surface roughnesses across the coil of 100's of nanometers. Which tolerances could you meet?

Add-ons: Is it possible (and would it make sense) to put an insulation layer on top of the chips? Out of curiosity, how difficult would it be to integrate additional gadgets, like a platinum temperature sensor, a Hall probe or a microwave patch antenna at 6.8 GHz?

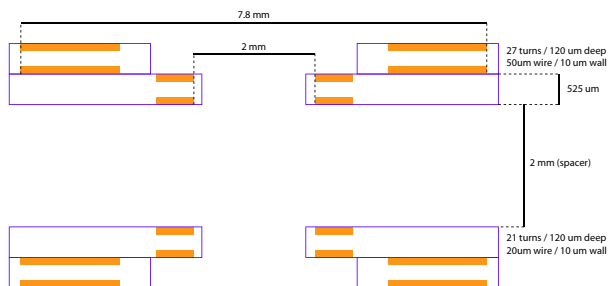



Chip Bonding/Stacking: Each chip coil has its own electrical connections. There is no current flowing between different chips. But for our magnet coil arrangement one Curvature coil chip and one Antibias coil chip need to be bonded together, respectively and the two pairs need to be bonded together with a 2mm thermally conductive spacer for the final assembly. A cross-section of the complete stack is also shown below.



The order of the 4 distinct coils (from top to bottom) is the following: Antibias (+) - Curvature (-) - spacer - Curvature (+) - Antibias (-), where the +/- denotes left/right handed coils. The leads for the Antibias coils face the same way, so do the leads of the Curvature coils, but in the opposite direction. The order is unique and given/labelled by the key-shaped markers next to the soldering pad.

Alignment marks (50 x 50 um squares) for chip-to-chip and stack-to-stack bonding would be available in diagonal corners on our masks. What would be reasonable tolerances for your alignment and bonding processes?



Bonding between chips could either be done by gold thermal compression bonding or gluing with a thermally conductive insulating and ultra-high vacuum compatible epoxy. Differential thermal expansion coefficients will be a critical issue (temperature range: -80C - 200C) and electrical shorting between chips, i.e. coils has to be avoided. Moreover, flatness, i.e. parallelity between the coils is very important.

The Spacer should be thermally conductive, of similar thermal expansion coefficient and should have very close tolerances on the parallelism of its faces. We thought about sapphire, but e.g. Si or any other suited material should also be ok. The bonding should be stable (there will be no mechanical forces on the stack) without inducing too much stress through the temperature cycles.

Wire attachment: Do you have the capabilities to already attach wires/leads to the solder pads? The process and the flying leads (~1") have to be ultra-high vacuum compatible and be able to carry ~4A.

Bibliography

- [1] Ehud Altman et al. “Phase diagram of two-component bosons on an optical lattice”. In: *New J. Phys.* 5.1 (Sept. 2003), pp. 113–113. ISSN: 1367-2630. DOI: [10.1088/1367-2630/5/1/113](https://doi.org/10.1088/1367-2630/5/1/113).
- [2] *amoberkeley / edcam* — *Bitbucket*. URL: <https://bitbucket.org/amoberkeley/edcam> (visited on 06/19/2015).
- [3] W S Bakr et al. “Probing the superfluid-to-Mott insulator transition at the single-atom level”. In: *Science* 329.5991 (July 2010), pp. 547–50. ISSN: 1095-9203. DOI: [10.1126/science.1192368](https://doi.org/10.1126/science.1192368).
- [4] Waseem S Bakr et al. “Orbital excitation blockade and algorithmic cooling in quantum gases”. In: *Nature* 480.7378 (Dec. 2011), pp. 500–3. ISSN: 1476-4687. DOI: [10.1038/nature10668](https://doi.org/10.1038/nature10668).
- [5] A B Bardou et al. “Transverse demagnetization dynamics of a unitary Fermi gas”. In: *Science* 344.6185 (May 2014), pp. 722–4. ISSN: 1095-9203. DOI: [10.1126/science.1247425](https://doi.org/10.1126/science.1247425).
- [6] D Baye and P-H Heenen. “Generalised meshes for quantum mechanical problems”. In: *J. Phys. A. Math. Gen.* 19.11 (Aug. 1986), pp. 2041–2059. ISSN: 0305-4470. DOI: [10.1088/0305-4470/19/11/013](https://doi.org/10.1088/0305-4470/19/11/013).
- [7] Daniel Baye. “Lagrange-mesh method for quantum-mechanical problems”. In: *Phys. status solidi* 243.5 (Apr. 2006), pp. 1095–1109. ISSN: 0370-1972. DOI: [10.1002/pssb.200541305](https://doi.org/10.1002/pssb.200541305).
- [8] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. “Quantum simulations with ultracold quantum gases”. In: *Nat. Phys.* 8.4 (Apr. 2012), pp. 267–276. ISSN: 1745-2473. DOI: [10.1038/nphys2259](https://doi.org/10.1038/nphys2259).
- [9] S. Braun et al. “Negative Absolute Temperature for Motional Degrees of Freedom”. In: *Science* 339.6115 (Jan. 2013), pp. 52–55. ISSN: 0036-8075. DOI: [10.1126/science.1227831](https://doi.org/10.1126/science.1227831).
- [10] SH Brongersma and E Richard. “Two-step room temperature grain growth in electroplated copper”. In: *J. Appl. Phys.* 86.7 (1999), p. 3642.

- [11] B. Capogrosso-Sansone et al. “Critical entropies for magnetic ordering in bosonic mixtures on a lattice”. In: *Phys. Rev. A* 81.5 (May 2010), p. 053622. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.81.053622](https://doi.org/10.1103/PhysRevA.81.053622).
- [12] Ming-Shien Chang et al. “Coherent spinor dynamics in a spin-1 Bose condensate”. In: *Nat. Phys.* 1.2 (Oct. 2005), pp. 111–116. ISSN: 1476-0000. DOI: [10.1038/nphys153](https://doi.org/10.1038/nphys153).
- [13] Franco Dalfovo, Stefano Giorgini, and Sandro Stringari. “Theory of Bose-Einstein condensation in trapped gases”. In: *Rev. Mod. Phys.* 71.3 (Apr. 1999), pp. 463–512. ISSN: 0034-6861. DOI: [10.1103/RevModPhys.71.463](https://doi.org/10.1103/RevModPhys.71.463).
- [14] F Dalfovo et al. “Nonlinear dynamics of a Bose condensed gas”. In: *Phys. Lett. A* 227.3-4 (Mar. 1997), pp. 259–264. ISSN: 03759601. DOI: [10.1016/S0375-9601\(97\)00069-8](https://doi.org/10.1016/S0375-9601(97)00069-8).
- [15] S. De et al. “Quenched binary Bose-Einstein condensates: Spin-domain formation and coarsening”. In: *Phys. Rev. A* 89.3 (Mar. 2014), p. 033631. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.89.033631](https://doi.org/10.1103/PhysRevA.89.033631).
- [16] Jesse Emspak. “Molecules Reach Coldest Temperature Ever”. In: *Sci. Am.* (2015).
- [17] Manuel Endres et al. “The ‘Higgs’ amplitude mode at the two-dimensional superfluid/Mott insulator transition”. In: *Nature* 487.7408 (July 2012), pp. 454–8. ISSN: 1476-4687. DOI: [10.1038/nature11255](https://doi.org/10.1038/nature11255).
- [18] Fang Fang et al. “Condensing magnons in a degenerate ferromagnetic spinor Bose gas”. In: (Nov. 2015), p. 5. arXiv: [1511.05193](https://arxiv.org/abs/1511.05193).
- [19] M. Fattori et al. “Demagnetization cooling of a gas”. In: *Nat. Phys.* 2.11 (Oct. 2006), pp. 765–768. ISSN: 1745-2473. DOI: [10.1038/nphys443](https://doi.org/10.1038/nphys443).
- [20] ME Gehm. *Properties of 6Li*. 2003. URL: <http://www.physics.ncsu.edu/jet/techdocs/pdf/PropertiesOfLi.pdf> (visited on 08/18/2015).
- [21] S. Giorgini, L. Pitaevskii, and S. Stringari. “Condensate fraction and critical temperature of a trapped interacting Bose gas”. In: *Phys. Rev. A* 54.6 (Dec. 1996), R4633–R4636. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.54.R4633](https://doi.org/10.1103/PhysRevA.54.R4633).
- [22] Markus Greiner et al. “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms”. In: *Nature* 415.6867 (Jan. 2002), pp. 39–44. ISSN: 0028-0836. DOI: [10.1038/415039a](https://doi.org/10.1038/415039a).
- [23] S. Gupta et al. “Bose-Einstein Condensation in a Circular Waveguide”. In: *Phys. Rev. Lett.* 95.14 (Sept. 2005), p. 143201. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.95.143201](https://doi.org/10.1103/PhysRevLett.95.143201).
- [24] Jennie Sara Guzman. “Explorations of Magnetic Phases in $F = 1$ ^{87}Rb Spinor Condensates”. PhD thesis. University of California, Berkeley, 2012.
- [25] J. Guzman et al. “Long-time-scale dynamics of spin textures in a degenerate $F = 1$ ^{87}Rb spinor Bose gas”. In: *Phys. Rev. A* 84.6 (Dec. 2011), p. 063625. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.84.063625](https://doi.org/10.1103/PhysRevA.84.063625).

- [26] James Michael Higbie. “First Steps toward Precision Measurements using Multicomponent Bose-Einstein Condensates of 87Rb”. PhD thesis. University of California, Berkeley, 2005.
- [27] Tin-Lun Ho. “Spinor Bose Condensates in Optical Traps”. In: *Phys. Rev. Lett.* 81.4 (July 1998), pp. 742–745. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.81.742](https://doi.org/10.1103/PhysRevLett.81.742).
- [28] Tin-Lun Ho and Michael Ma. “Quasi 1 and 2d Dilute Bose Gas in Magnetic Traps: Existence of Off-Diagonal Order and Anomalous Quantum Fluctuations”. en. In: *J. Low Temp. Phys.* 115.1-2 (Apr. 1999), pp. 61–70. ISSN: 1573-7357. DOI: [10.1023/A:1021894713105](https://doi.org/10.1023/A:1021894713105).
- [29] W. Hofstetter et al. “High-Temperature Superfluidity of Fermionic Atoms in Optical Lattices”. In: *Phys. Rev. Lett.* 89.22 (Nov. 2002), p. 220407. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.89.220407](https://doi.org/10.1103/PhysRevLett.89.220407).
- [30] Gregor Jotzu et al. “Experimental realization of the topological Haldane model with ultracold fermions”. In: *Nature* 515.7526 (Nov. 2014), pp. 237–240. ISSN: 0028-0836. DOI: [10.1038/nature13915](https://doi.org/10.1038/nature13915).
- [31] Yuki Kawaguchi et al. “Topological Excitations in Spinor Bose-Einstein Condensates”. In: *Prog. Theor. Phys. Suppl.* 186 (June 2010), pp. 455–462. ISSN: 0375-9687. DOI: [10.1143/PTPS.186.455](https://doi.org/10.1143/PTPS.186.455). arXiv: [1006.5839](https://arxiv.org/abs/1006.5839).
- [32] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. “Making, probing and understanding Bose-Einstein condensates”. In: (Apr. 1999), p. 90. arXiv: [9904034](https://arxiv.org/abs/9904034) [[cond-mat](#)].
- [33] Wolfgang Ketterle and N.J. Van Druten. “Evaporative Cooling of Trapped Atoms”. In: ed. by Benjamin Bederson and Herbert Walther. Vol. 37. *Advances In Atomic, Molecular, and Optical Physics*. Academic Press, 1996, pp. 181–236. DOI: [http://dx.doi.org/10.1016/S1049-250X\(08\)60101-9](https://doi.org/http://dx.doi.org/10.1016/S1049-250X(08)60101-9).
- [34] In Kwon Kim et al. “Effect of corrosion inhibitor, benzotriazole, in Cu slurry on Cu polishing”. In: *Jpn. J. Appl. Phys.* 47.1 (2008), pp. 108–112. ISSN: 00214922. DOI: [10.1143/JJAP.47.108](https://doi.org/10.1143/JJAP.47.108).
- [35] Charles Kittel and Herbert Kroemer. *Thermal Physics*. W. H. Freeman, 1980, p. 473. ISBN: 0716710889.
- [36] Marco Koschorreck et al. “Universal spin dynamics in two-dimensional Fermi gases”. In: *Nat. Phys.* 9.7 (May 2013), pp. 405–409. ISSN: 1745-2473. DOI: [10.1038/nphys2637](https://doi.org/10.1038/nphys2637).
- [37] J. Kronjäger et al. “Evolution of a spinor condensate: Coherent dynamics, dephasing, and revivals”. In: *Phys. Rev. A* 72.6 (Dec. 2005), p. 063619. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.72.063619](https://doi.org/10.1103/PhysRevA.72.063619).
- [38] A E Leanhardt et al. “Cooling Bose-Einstein condensates below 500 picokelvin”. In: *Science* 301.5639 (Sept. 2003), pp. 1513–5. ISSN: 1095-9203. DOI: [10.1126/science.1088827](https://doi.org/10.1126/science.1088827).

- [39] H. J. Lewandowski et al. “Decoherence-Driven Cooling of a Degenerate Spinor Bose Gas”. In: *Phys. Rev. Lett.* 91.24 (Dec. 2003), p. 240404. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.91.240404](https://doi.org/10.1103/PhysRevLett.91.240404).
- [40] H. J. Lewandowski et al. “Simplified System for Creating a Bose–Einstein Condensate”. In: *J. Low Temp. Phys.* 132.5 (2003), pp. 309–367. ISSN: 0022-2291. DOI: [10.1023/A:1024800600621](https://doi.org/10.1023/A:1024800600621).
- [41] David R. Lide, ed. *CRC Handbook of Chemistry and Physics, 84th Edition, 2003-2004*. Vol. 53. Taylor & Francis, 2003, p. 2616. ISBN: 0849304849. DOI: [10.1136/oem.53.7.504](https://doi.org/10.1136/oem.53.7.504).
- [42] J. C. Light and T. Carrington. “Discrete-Variable Representations and their Utilization”. In: *Adv. Chem. Phys.* Ed. by I. Prigogine and Stuart A. Rice. Vol. 114. Advances in Chemical Physics. Hoboken, NJ, USA: John Wiley & Sons, Inc., Jan. 2000. Chap. 4, pp. 263–310. ISBN: 9780470141731. DOI: [10.1002/9780470141731.ch4](https://doi.org/10.1002/9780470141731.ch4).
- [43] Y. -J. Lin et al. “Rapid production of 87Rb Bose-Einstein condensates in a combined magnetic and optical potential”. In: *Phys. Rev. A* 79.6 (June 2009), p. 063631. DOI: [10.1103/PhysRevA.79.063631](https://doi.org/10.1103/PhysRevA.79.063631). arXiv: [0904.3314](https://arxiv.org/abs/0904.3314).
- [44] T.H. Loftus, A. Vitouchkine, and L. Hollberg. *Permanent magnet axial field zeeman slower*. US Patent 8,710,428. Apr. 2014.
- [45] O. Luiten, M. Reynolds, and J. Walraven. “Kinetic theory of the evaporative cooling of a trapped gas”. In: *Phys. Rev. A* 53.1 (Jan. 1996), pp. 381–389. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.53.381](https://doi.org/10.1103/PhysRevA.53.381).
- [46] G. Edward Marti, Ryan Olf, and Dan M. Stamper-Kurn. “Collective excitation interferometry with a toroidal Bose-Einstein condensate”. In: *Phys. Rev. A* 91.1 (Jan. 2015), p. 013602. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.91.013602](https://doi.org/10.1103/PhysRevA.91.013602). arXiv: [arXiv:1210.0033](https://arxiv.org/abs/1210.0033).
- [47] G. Edward Marti et al. “Coherent Magnon Optics in a Ferromagnetic Spinor Bose-Einstein Condensate”. In: *Phys. Rev. Lett.* 113.15 (Oct. 2014), p. 155302. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.113.155302](https://doi.org/10.1103/PhysRevLett.113.155302).
- [48] G. Edward Marti et al. “Two-element Zeeman slower for rubidium and lithium”. In: *Phys. Rev. A* 81.4 (Apr. 2010), p. 043424. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.81.043424](https://doi.org/10.1103/PhysRevA.81.043424).
- [49] George Edward Marti. “Scalar and Spinor Excitations in a Ferromagnetic Bose-Einstein Condensate”. PhD thesis. University of California, Berkeley, 2014.
- [50] *Marvell Nanofabrication Laboratory Manual*. URL: <https://nanolab.berkeley.edu/labmanual/labmantoc.shtml> (visited on 09/27/2015).
- [51] Abraham H. Maslow. *The Psychology of Science: A Reconnaissance*. 1st. New York: Harper & Row, 1966. ISBN: 0976040239.

- [52] T.B. Massalski et al. *Binary Alloy Phase Diagrams*. American Society for Metals, 2001. ISBN: 9780871704030.
- [53] D C McKay and B DeMarco. “Cooling in strongly correlated optical lattices: prospects and challenges”. In: *Reports Prog. Phys.* 74.5 (May 2011), p. 054401. ISSN: 0034-4885. DOI: [10.1088/0034-4885/74/5/054401](https://doi.org/10.1088/0034-4885/74/5/054401).
- [54] D McKay and B DeMarco. “Thermometry with spin-dependent lattices”. en. In: *New J. Phys.* 12.5 (May 2010), p. 055013. ISSN: 1367-2630. DOI: [10.1088/1367-2630/12/5/055013](https://doi.org/10.1088/1367-2630/12/5/055013).
- [55] Dermot McPeake. “Superfluidity, collective excitations and nonlinear dynamics of Bose-Einstein condensates”. Thesis. Queen’s University, Belfast, 2002.
- [56] Patrick Medley et al. “Spin Gradient Demagnetization Cooling of Ultracold Atoms”. In: *Phys. Rev. Lett.* 106.19 (May 2011), p. 195301. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.106.195301](https://doi.org/10.1103/PhysRevLett.106.195301).
- [57] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Springer Science & Business Media, 1999, p. 323. ISBN: 0387987282.
- [58] John H. Moore et al. *Building Scientific Apparatus*. Cambridge University Press, 2009. ISBN: 1139478656.
- [59] S. Nascimbène et al. “Exploring the thermodynamics of a universal Fermi gas”. In: *Nature* 463.7284 (Feb. 2010), pp. 1057–1060. ISSN: 0028-0836. DOI: [10.1038/nature08814](https://doi.org/10.1038/nature08814).
- [60] UN-Oceans. *UN Atlas of the Oceans*. URL: <http://www.oceansatlas.org/>.
- [61] John F. O’Hanlon. *A User’s Guide to Vacuum Technology*. John Wiley & Sons, 2005, p. 536. ISBN: 0471467154.
- [62] Ryan Olf et al. “Thermometry and cooling of a Bose gas to 0.02 times the condensation temperature”. In: *Nat. Phys.* 11.9 (May 2015), pp. 720–723. arXiv: [1505.06196](https://arxiv.org/abs/1505.06196).
- [63] Abraham J. Olson, Robert J. Niffenegger, and Yong P. Chen. “Optimizing the efficiency of evaporative cooling in optical dipole traps”. In: *Phys. Rev. A* 87.5 (May 2013), p. 053613. ISSN: 1050-2947. DOI: [10.1103/PhysRevA.87.053613](https://doi.org/10.1103/PhysRevA.87.053613).
- [64] Thereza Paiva et al. “Fermions in 3D Optical Lattices: Cooling Protocol to Obtain Antiferromagnetism”. In: *Phys. Rev. Lett.* 107.8 (Aug. 2011), p. 086401. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.107.086401](https://doi.org/10.1103/PhysRevLett.107.086401).
- [65] Belén Paredes et al. “Tonks-Girardeau gas of ultracold atoms in an optical lattice”. In: *Nature* 429.6989 (May 2004), pp. 277–81. ISSN: 1476-4687. DOI: [10.1038/nature02530](https://doi.org/10.1038/nature02530).
- [66] Jee Woo Park, Sebastian A. Will, and Martin W. Zwierlein. “Ultracold Dipolar Gas of Fermionic Na ²³K ⁴⁰ Molecules in Their Absolute Ground State”. In: *Phys. Rev. Lett.* 114.20 (May 2015), p. 205302. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.114.205302](https://doi.org/10.1103/PhysRevLett.114.205302).

- [67] C. J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, 2002, p. 402. ISBN: 0521665809.
- [68] Nguyen Thanh Phuc, Yuki Kawaguchi, and Masahito Ueda. “Beliaev theory of spinor Bose–Einstein condensates”. In: *Ann. Phys.* 328 (Jan. 2013), pp. 158–219. ISSN: 00034916. DOI: [10.1016/j.aop.2012.10.004](https://doi.org/10.1016/j.aop.2012.10.004).
- [69] Lev. P. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. Oxford University Press, 2003, p. 382. ISBN: 0198507194.
- [70] Thomas P. Purdy. “Cavity QED with Ultracold Atoms on an Atom Chip”. PhD thesis. University of California, Berkeley, 2009.
- [71] *Python Learning Resources*. URL: http://python.berkeley.edu/learning%7B%5C_%7Dresources.html (visited on 06/19/2015).
- [72] C.A. Regal and D.S. Jin. “Experimental Realization of the BCS-BEC Crossover with a Fermi Gas of Atoms”. In: *Adv. At. Mol. Opt. Phys.* 54 (2007), pp. 1–79. ISSN: 1049250X. DOI: [10.1016/S1049-250X\(06\)54001-7](https://doi.org/10.1016/S1049-250X(06)54001-7).
- [73] Craig J Sansonetti et al. “Absolute transition frequencies and quantum interference in a frequency comb based measurement of the (6,7)Li D lines”. In: *Phys. Rev. Lett.* 107.2 (July 2011), p. 023001. ISSN: 1079-7114. DOI: [10.1103/PhysRevLett.107.023001](https://doi.org/10.1103/PhysRevLett.107.023001).
- [74] R. Schley et al. “Planck Distribution of Phonons in a Bose-Einstein Condensate”. In: *Phys. Rev. Lett.* 111.5 (July 2013), p. 055301. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.111.055301](https://doi.org/10.1103/PhysRevLett.111.055301).
- [75] S. Seal, S. C. Kuiry, and B. Heinmen. “Effect of glycine and hydrogen peroxide on chemical-mechanical planarization of copper”. In: *Thin Solid Films* 423.2 (2003), pp. 243–251. ISSN: 00406090. DOI: [10.1016/S0040-6090\(02\)00989-6](https://doi.org/10.1016/S0040-6090(02)00989-6).
- [76] I Shvarchuck et al. “Bose-Einstein Condensation into Nonequilibrium States Studied by Condensate Focusing”. In: *Phys. Rev. Lett.* 89.27 (Dec. 2002), p. 270404.
- [77] F. M. Spiegelhalter et al. “Collisional Stability of K 40 Immersed in a Strongly Interacting Fermi Gas of Li 6”. In: *Phys. Rev. Lett.* 103.22 (Nov. 2009), p. 223203. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.103.223203](https://doi.org/10.1103/PhysRevLett.103.223203).
- [78] I. Spielman, W. Phillips, and J. Porto. “Mott-Insulator Transition in a Two-Dimensional Atomic Bose Gas”. In: *Phys. Rev. Lett.* 98.8 (Feb. 2007), p. 080404. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.98.080404](https://doi.org/10.1103/PhysRevLett.98.080404).
- [79] Dan M. Stamper-Kurn and Masahito Ueda. “Spinor Bose gases: Explorations of symmetries, magnetism and quantum dynamics”. In: (May 2012), p. 57. arXiv: [1205.1888](https://arxiv.org/abs/1205.1888).
- [80] Cladiu A. Stan. “Experiments with Interacting Bose and Fermi Gases”. PhD thesis. Massachusetts Institute of Technology, 2005, pp. 1–204.
- [81] Cladiu A. Stan and Wolfgang Ketterle. “Multiple species atom source for laser-cooling experiments”. In: *Rev. Sci. Instrum.* 76 (2005), pp. 1–5. ISSN: 00346748. DOI: [10.1063/1.1935433](https://doi.org/10.1063/1.1935433). arXiv: [0503192](https://arxiv.org/abs/0503192) [physics].

- [82] Daniel A. Steck. *Rubidium 87 D Line Data*. URL: <http://steck.us/alkalidata/>.
- [83] Daniel Adam Steck. *Quantum and Atom Optics*. 2011. URL: <http://steck.us/teaching>.
- [84] Jeff Steinhauer. “Observation of self-amplifying Hawking radiation in an analogue black-hole laser”. In: *Nat. Phys.* 10.11 (Oct. 2014), pp. 864–869. ISSN: 1745-2473. DOI: [10.1038/nphys3104](https://doi.org/10.1038/nphys3104).
- [85] Jian-Jun Sun et al. “High-Aspect-Ratio Copper Via Filling Used for Three-Dimensional Chip Stacking”. en. In: *J. Electrochem. Soc.* 150.6 (June 2003), G355. ISSN: 00134651. DOI: [10.1149/1.1572154](https://doi.org/10.1149/1.1572154).
- [86] S Trotzky et al. “Time-resolved observation and control of superexchange interactions with ultracold atoms in optical lattices”. In: *Science* 319.5861 (Jan. 2008), pp. 295–9. ISSN: 1095-9203. DOI: [10.1126/science.1150841](https://doi.org/10.1126/science.1150841).
- [87] S. Tung et al. “Observation of the Presuperfluid Regime in a Two-Dimensional Bose Gas”. In: *Phys. Rev. Lett.* 105.23 (Dec. 2010), p. 230408. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.105.230408](https://doi.org/10.1103/PhysRevLett.105.230408).
- [88] Ling Wang and Fiona M Doyle. “Mechanisms of Passivation of Copper in CMP Slurries Containing Peroxide and Glycine”. In: *Mater. Res.* 767 (2003), pp. 1–10. ISSN: 02729172. DOI: [10.1149/1.1615611](https://doi.org/10.1149/1.1615611).
- [89] Artur Widera et al. “Precision measurement of spin-dependent interaction strengths for spin-1 and spin-2 87 Rb atoms”. In: *New J. Phys.* 8.8 (Aug. 2006), pp. 152–152. ISSN: 1367-2630. DOI: [10.1088/1367-2630/8/8/152](https://doi.org/10.1088/1367-2630/8/8/152).
- [90] Wikipedia. *Relational database*. 2015. URL: https://en.wikipedia.org/?title=Relational%7B%5C_%7Ddatabase.