

ULTRACOLD ATOMS IN MICROFABRICATED MAGNETIC
TRAPS NEAR SURFACES AND INSIDE OPTICAL RESONATORS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Igor Teper
September 2006

© Copyright by Igor Teper 2006
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Vladan Vuletić) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Mark Kasevich)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Phil Bucksbaum)

Approved for the University Committee on Graduate Studies.

To the memory of my grandfathers, Abram Teper and Yaroslav Zeman, and my
great-grandmother, Elizaveta Zeman

Abstract

Microfabricated magnetic traps (“microtraps”) for ultracold atoms allow the creation of complex, precisely-controlled, and sharply-varying potentials that hold great promise for a variety of new quantum devices for atomic matter waves. We present two experiments that examine the capabilities and also the fundamental limits of microtraps.

In the first experiment, we investigate the stability of magnetically trapped Bose-Einstein condensates and thermal clouds near the transition temperature at distances of 0.5-10 microns from a microfabricated silicon chip. Near a copper film, the trap lifetime is limited by spin flips due to the coupling of the atoms’ electronic spins to magnetic field fluctuations produced by the thermal motion of electrons in the copper (Johnson noise). We present a simple formula that can account for this observed loss process with no free parameters. A dielectric surface has no adverse effect on the trapped atoms until they are brought so close to the surface that the attractive Casimir-Polder potential reduces the trap depth, which leads to loss of atoms.

In the second experiment, we implement sensitive atom detection for atoms in the microtrap through the use of a medium-finesse, macroscopic optical resonator integrated with the chip. As detection methods, we employ both fluorescence into the cavity and atom-induced reduction in cavity transmission. In fluorescence, we register 2.0 photon counts per atom, which allows us to detect single atoms with 75% efficiency in 250 μ s. In absorption, we measure transmission attenuation of 3.3% per atom, which allows us to count small numbers of atoms with a resolution of about 1 atom. We also demonstrate shot-noise-limited atom preparation down to 1 atom.

Acknowledgements

I have been very lucky in my graduate school experiences, both professionally and personally. I spent the first three years of graduate school working at Stanford University, and the last three working at the Massachusetts Institute of Technology, and at both places I have benefited a great deal from the wonderful people around me.

I owe many thanks to the very skilled and supportive members of Steve Chu's group at Stanford, especially Jamie Kerman, Nate Gemelke, Edina Sarajlic, and Sheng-Wey Chiow.

I would also like to thank all the Stanford physics department staff who took care of me in ways small and large throughout my time there. They include Joel Crawford, Stewart Kramer, Rick Pam, Kathleen Guan, Ping Feng, and Maria Frank. During the past three years, Maria in particular has done a lot to make sure all the administrative problems caused by my being three thousand miles away from Stanford were resolved quickly and smoothly.

When we arrived at MIT, we were greeted with open arms at the Center for Ultracold Atoms. Carol Costa especially went out of her way to make sure we felt welcome and were able to quickly adapt to our new surroundings. She is deeply missed.

In the course of my time at MIT, Ellenor Emery Barish, Joanna Keseberg, and Al McGurl have provided much help and support that I'm very grateful for.

I have found the Center for Ultracold Atoms at MIT to be a wonderful place to do physics, due largely to the amazing colleagues I've been surrounded by. Wolfgang Ketterle, Dave Pritchard, and Dan Kleppner, and their many students and postdocs I've overlapped with (whom I will not try to list because I would certainly forget

someone), have helped me on countless occasions and in countless ways.

Over the years, our group has gained much from the work of many talented and dedicated undergraduate students, among them Xinan Wu, Dan Flickinger, Mike Childress, Brendan Shields, and Huanqian Loh.

I am very glad to have gotten to know and appreciate our group's younger graduate students, Jonathan Simon, Andrew Grier, Marko Cetina, and Haruka Tanji. My experiment will soon be left in the very capable hands of Ian Leroux and Monika Schleier-Smith; I am grateful for all their hard work over the past year and I anticipate hearing about the great results they'll achieve in the future.

James Thompson, though a postdoc on another experiment, was extremely helpful to me over the past three years. In addition to being the one in the group who knew how things worked at MIT, he was always willing to share his informed and balanced perspective on various issues.

Adam Black and Hilton Chan welcomed me to the group and helped me negotiate the learning curve when I was a first-year rotation student. Their friendly attitude helped convince me that this was the group where I wanted to settle for the long haul, a decision I have never regretted. Adam in particular has been a source of helpful advice and wisdom, and a good friend, throughout my grad school career.

In the early stages of our experiment, I was lucky to have Cheng Chin as a postdoc for almost two years. He was the driving force behind the experiment's initial assembly, and his great technical skills, energy, and dedication set the tone in the lab. Circumstances forced him to cut short the time he spent on our project, but my accomplishments since his departure are largely a reflection of how much I learned from him.

I have spend the past five years working side-by-side with Yu-Ju Lin, and I could not have asked for a better labmate. Her quiet and unassuming demeanor sometimes hides just how very skilled and knowledgeable she is, but I certainly know the truth. I have always been very impressed by Yu-Ju's perseverance and thoroughness, her always positive and helpful attitude makes her a pleasure to work with, and I have learned more than I can say from our many discussions of various experimental and conceptual issues, as well as from simply working alongside her. On top of everything

else, Yu-Ju has made all of the chips used in our experiments, and has spent countless hours perfecting various aspects of the fabrication recipes. She deserves at least as much credit as I do for the work presented in this thesis.

I have been incredibly fortunate to have Vladan Vuletić as my adviser. His boundless energy and creativity, and his combination of an encyclopedic breadth of technical knowledge with a deep and insightful understanding of physical concepts, make him a great physicist, but it is his patience, enthusiasm, accessibility, and willingness to explain that make him a great mentor. When Vladan talks about physics, one can't help but share his excitement – it was his presentation during my grad school orientation that initially drew me to his group, and he has always been able to get me excited about physics when I have become discouraged. No matter what problem I encountered in the lab, Vladan usually knew the best solution, and if he didn't, he could always quickly come up with many very good ideas. In the way he approaches problems, in the way he sets priorities, in the way he carries himself in the lab and in the classroom, Vladan is a model and inspiration to me.

My parents, Yan Teper and Irina Zeman, have always been amazingly supportive and understanding, and I will certainly never be able to thank them enough for the opportunities I have had because of them. My sister Marina is a constant source of encouragement and advice, which I always appreciate. My grandmothers, Stella and Betya Chollak and Roza Zeman have had more faith in me than I could ever justify.

During my time in graduate school, I was incredibly lucky to meet Beth McKeown, and far luckier still to marry her. She has been my greatest cheerleader, and for that, among many other reasons, I thank her.

Contents

Abstract	vi
Acknowledgements	vii
1 Introduction	1
2 Experimental Setup	5
2.1 Apparatus Overview	5
2.2 Diode Laser System	6
2.2.1 Laser Parameters	6
2.2.2 Laser Frequency Stabilization	8
2.3 Vacuum System and Magneto-Optical Trap	12
2.4 Atom Chip	15
2.4.1 Chip Mount and Auxiliary Coils	15
2.4.2 Chip Current Coils	17
2.5 Imaging	19
3 Microtrap Loading and Evaporation to BEC	23
3.1 Loading the Magnetic Microtrap	23
3.2 Microtrap Characteristics	25
3.3 Evaporation to BEC	28
4 Measurement of Surface Effects	34
4.1 Precise Positioning of Atoms	34

4.1.1	Microtrap Field Countours	34
4.1.2	Calibrating Magnetic Fields at Origin	37
4.1.3	Surface Microscopy	38
4.2	Lifetime Limits due to Surface-Induced Loss	39
4.3	Casimir-Polder Force near Dielectric	44
4.4	Condensate Splitting	49
5	Integrating an Optical Cavity with the Chip	53
5.1	Motivation for Using Cavity	53
5.2	Cavity on a Chip	56
5.2.1	Modified Chip Design	56
5.2.2	Physical Implementation of Cavity	57
5.3	Cavity Parameters	61
5.4	Cavity Length Stabilization	64
6	Cavity-Aided Atom Detection	67
6.1	General Imaging Analysis	67
6.2	Atom Preparation	70
6.3	Measurement of Cavity Cooperativity η_1	71
6.4	Fluorescence Detection	74
6.4.1	Atom Number Measurement	74
6.4.2	Expected Fluorescence Signal	77
6.4.3	Single-Atom Detection	77
6.5	Absorption Detection	79
6.5.1	Atom Number Measurement	79
6.5.2	Expected Absorption Signal	82
6.6	Comparison of Fluorescence and Absorption	86
7	Conclusion and Future Work	89
7.1	Conclusion	89
7.2	Future Work	91

A Derivations of Correlation Relations	94
A.1 Fluorescence	94
A.2 Absorption	95
Bibliography	98

List of Tables

3.1 Parameters for evaporation to BEC 31

List of Figures

2.1	Schematic of grating-feedback diode laser.	6
2.2	^{87}Rb level structure and laser frequencies.	11
2.3	Schematic of main vacuum chamber and MOT beams.	13
2.4	Schematic of atom chip holder.	16
2.5	Layout of atom chip.	18
2.6	Schematic of CCD absorption imaging.	20
3.1	Optical pumping level diagram.	24
3.2	Microtrap transverse vibration frequency measurement.	27
3.3	Phase-space density as function of atom number during evaporation.	32
3.4	Formation of BEC.	33
4.1	Transverse-plane magnetic-field contours for microtrap.	35
4.2	Contours for transverse magnetic field components in region near chip.	36
4.3	Measurement to determine applied magnetic fields at symmetry point.	38
4.4	Surface microscopy and lifetime measurement paths	40
4.5	Trapped atom lifetime as function of distance from dielectric and metallic surfaces	41
4.6	Rapid loss near dielectric	45
4.7	Loss measurement at two locations	48
4.8	Diagram of condensate splitting at symmetry point.	50
4.9	Results of condensate splitting at symmetry point.	51
5.1	Layout of modified atom chip.	58

5.2	Schematic of cavity mount	59
5.3	Diagram of cavity coupling	60
5.4	Cavity lock diagram	66
6.1	Atom detection via fluorescence and absorption.	69
6.2	Positioning atoms in cavity.	72
6.3	Measurement of cavity cooperativity.	73
6.4	Characterization of fluorescence measurement.	76
6.5	Fluorescence measurement results.	78
6.6	Characterization of single-atom detection.	80
6.7	Characterization of absorption measurement.	82
6.8	Absorption measurement results.	83
6.9	Comparison of uncertainties for fluorescence and absorption detection	87

Chapter 1

Introduction

One of the most intensive directions of atomic physics research following the achievement of Bose-Einstein condensation (BEC) in a gas of neutral atoms [1, 2, 3] has been the implementation of quantum atom optics devices that would take advantage of the macroscopic-wave-like nature of BECs. The atom laser [4, 5, 6, 7] is perhaps the most prominent such device, but there are many others that have received experimental and theoretical attention. A research area that has attracted a lot of interest as a promising means for implementing such devices is magnetic traps produced by microfabricated structures (“microtraps”) patterned on chips.

More generally, one way to describe the mission of atomic physics is that it seeks to achieve maximal quantum control over every atomic degree of freedom, both internal and external. To that end, advances in laser sources continue to improve our ability to manipulate atoms’ internal states, while laser and evaporative cooling have made possible near-total control of atomic kinetic energy. Full control over atomic position and motion requires potentials that vary abruptly on the length scale of the atoms’ effective size, given by the de Broglie wavelength, which typically equals around $1 \mu\text{m}$ for ultracold atoms at temperatures around $1 \mu\text{K}$. For magnetically trapped atoms, the sharpness of the potential is given by the magnetic field gradient; since the gradient produced by a conducting wire varies with the inverse square of the distance from the wire, potentials suitable for fine manipulation of atomic position can be created at small distances from miniaturized conductors. Furthermore, in order to

be useful for atom manipulation, these potentials must be not only sharply varying but also precisely known and controlled, which requires that the parameters of the conductors that generate them be likewise precisely controlled. Modern microfabrication techniques allow just this kind of control over the position, size, and shape of conductors deposited on substrates. In addition, the large gradients that can be achieved in microtraps allow high initial trapped atom densities, which corresponds to short thermalization times and therefore allow fast evaporative cooling and BEC production. Microtraps on chips thus offer the capability to produce tight, complex, and precisely controlled magnetic potentials that hold promise for a wide array of quantum atom optics devices.

The first microtrap application to receive intense study was atom waveguides [8, 9, 10]. Magnetic traps generated using straight and parallel current-carrying wires combined with uniform offset fields produce confinement in the plane perpendicular to the wires while the direction along the wires remains unconfined; this is precisely the geometry of a magnetic waveguide. Thus, microtrap waveguides are straightforward to implement, and, furthermore, they are an important component of many more advanced applications, such as guided atoms interferometers (which also require the implementation of beam splitters [11]).

Following the achievement of Bose-Einstein condensation in a microtrap [12, 13, 14], guided atom interferometry in microtraps has received a lot of experimental attention, and, while many serious problems have been encountered, there have recently been several important successes in this area [15, 16, 17, 18].

In addition to trapped atom interferometry, there has been significant work done in BEC Josephson junctions [5, 19] and atomic Fabry-Perot resonators [20, 21] in microtraps. Also, the tight confinement offered by microtraps allows access to one-dimensional physics, such as the creation of a Tonks-Girardeau gas [22, 23, 24, 25, 26], where the interatomic interaction in the presence of confinement leads to a one-dimensional gas of “fermionized” impenetrable bosons [27, 28].

There are, however, both technical and fundamental limits to the miniaturization achievable on a chip. Early experiments observed condensate fragmentation [14, 29, 30], heating [12, 30], and reduced trap lifetime [30]. These effects turned out to

be technical, with the fragmentation being explained by spatial variations of the longitudinal magnetic field near a current-carrying conductor due to [31], while the heating and loss were due to stray radiofrequency fields and could be eliminated by careful electronic design and shielding [14]. But more fundamental limits were also seen, with an observation of loss due to spin flips induced by thermal currents in mesoscopic conductors close to the atoms [32, 33], in good agreement with theory [34, 35].

This dissertation describes experiments that extended the study of surface effects on the stability of ultracold atoms and BECs [36]. Our microtrap design has the advantage of allowing us, through the use of sensitive atom loss processes, to very precisely determine the position of atoms (with respect to the chip surface) that corresponds to a given applied magnetic field. This allows us to measure trapped atom lifetimes down to very short distances (less than a micron) from the surface with excellent distance resolution. Furthermore, we are able to move the trap to position the atoms near either a conducting or a dielectric surface, and measure loss effects from both. We find that the distance dependence of the atom loss as the trap approached a conductor is quantitatively explained by thermal magnetic field fluctuations arising from Johnson-noise-induced currents [34, 35]. Near a dielectric surface, we observe no loss until we approach close enough that the attractive Casimir-Polder potential [37, 38] from the surface begins to limit the trap depth. The loss can be quantitatively modeled by a one-dimensional evaporation model.

Many proposed or existing microtrap experiments would greatly benefit from being able to detect small numbers of atoms or single atoms, either because the miniaturization necessary to conduct them limits the atom number used or because they require measurement of atom statistics and correlations at the single-atom level. The detection of single atoms on a chip is thus an important technical problem in the field [39, 40].

The dissertation discusses our implementation of a sensitive atom detector on the atom chip, which we do by integrating a medium-finesse macroscopic optical resonator around the microtrap [41]. As first discussed by Purcell, atomic emission into and optical depth within a resonator can be greatly enhanced compared to free space

[42], which allows us to observe measurable signals for very small atom numbers. We are able to detect single atoms with high efficiency using cavity-aided fluorescence detection and to count small numbers of atoms with very good resolution using both fluorescence and absorption detection. We are also able to use statistical analysis to show that we achieve shot-noise-limited atom preparation down to one atom.

In Chapter 2, I describe the experimental apparatus used in the surface-effects experiments. Chapter 3 discusses the loading of atoms into the magnetic microtrap on the chip and the evaporation to BEC, and introduces some of the theory of Bose-Einstein condensation. Chapter 4 describes our experiments on the measurement of surface effects, from the precise positioning of atoms near the chip to the measurement and modeling of loss near conducting and dielectric surfaces. Chapter 5 begins with a discussion of resonator physics and a motivation for integrating an optical cavity with our chip, and then goes on to describe the modifications made to our atom chip apparatus in order to accommodate the cavity, as well as the physical implementation of the cavity itself. Finally, in Chapter 6, I discuss our experiments on cavity-aided sensitive atom detection on the atom chip. I begin with a theoretical analysis of the fundamental limits for fluorescence and absorption imaging, then present our results for fluorescence detection for small atom numbers and for single atoms, followed by the results for absorption detection for small atom numbers, and conclude with a comparison between fluorescence and absorption detection in light of both the theoretical analysis and our experimental results. Chapter 7 includes a summary and conclusion, as well as a brief discussion of planned future work.

Chapter 2

Experimental Setup

2.1 Apparatus Overview

Ultracold atom experiments generally require a laser system, a vacuum system, an atom source, coils to generate magnetic fields, and an imaging system. There has been a great deal of technological development in all these areas, so that there now exist a variety of methods to implement each of these experimental subsystems, with different approaches possessing different, and often complementary strengths. The particular implementation chosen can therefore be customized to the needs of the experiment, be it in terms of laser power, vacuum quality, atom number, experiment cycling time, optical access, or many other such parameters.

This chapter describes the experimental apparatus used in our experiments and also discusses some of the considerations and trade-offs involved in our decisions to implement things in a particular way. The apparatus has always been a work-in-progress, and we have made changes and improvements as we have learned more about its capabilities and limitations and have refined our understanding of the various problems we've tried to address. The laser system in particular has undergone several major and minor overhauls over the years, and in its description I discuss some the issues we have encountered.




Figure 2.1: Schematic of initial design for grating-feedback diode lasers used in the experiment. The grating’s first-order diffraction is sent back to the laser while the zeroth-order reflection is transmitted out. The red lines show the laser output and feedback.

2.2 Diode Laser System

2.2.1 Laser Parameters

Our experiment uses external-cavity grating feedback diode lasers (“grating lasers”) [43] operating on the ^{87}Rb D_2 ($5^2S_{1/2} \rightarrow 5^2P_{3/2}$) transition at a wavelength of 780.2 nm. Such lasers provide stability, frequency tunability, and linewidths in the 1-1.5 MHz range, narrow compared to the $\Gamma = 2\pi \times 6.1$ MHz natural linewidth of the ^{87}Rb D_2 line.

The grating laser design is as follows: the laser diode is mounted in a holder which is attached to a temperature-stabilized aluminum plate. A threaded mount for the collimating lens is glued to the laser diode holder so that the collimator is centered on the laser beam. The grating is glued to a piezoelectric transducer (PZT) which is glued to a kinematic lens mount attached to the same aluminum plate several centimeters away from the laser diode. The grating is positioned in a Littrow configuration, with the first-order diffraction going back to the laser diode (see Fig. 2.1).

The kinematic mount allows one to both tune the frequency of the laser and optimize the feedback from the grating in order to maximize laser stability. Fine-tuning of the frequency is achieved by changing the laser current and simultaneously changing the cavity length by applying a voltage to the PZT. The frequency can also be tuned by varying the temperature of the laser diode, which changes the length of the diode chip's internal laser cavity, but this method is much slower and less robust, and therefore is avoided when possible.

The relevant trade-off when considering which grating to use is that between higher feedback, which allows for more stable single-mode operation, and higher output power. We have found that a 1200 lines/mm holographic grating with first-order diffraction of 28% and loss of 17% satisfies both the stability and power requirements of our experiment [44].

An important limitation of the linewidth of grating lasers is mechanical vibrations that modulate the external cavity. To address this issue, we have developed a modified grating laser design, in which the piece that holds the laser diode and collimator is screwed to the piece that holds the grating, with rubber in between to damp out mechanical vibrations.

Another issue that has an impact on the stability of diode lasers is the competition between modes of the external (grating) cavity and the laser diode chip's internal modes. The presence of the internal modes sets a preferred laser frequency region that is significantly narrower than the diode's gain profile, which limits the extent to which the laser frequency can be tuned by the grating. In addition, occasional mode-hopping between internal chip modes limits the stability of the laser, though this can be minimized by careful alignment of the grating feedback combined with optimizing the tuning current so that the internal and external cavities tune together. Problems arising from the presence of internal chip modes can be solved by using diodes that have an anti-reflection (AR) coating applied to the output facet, so that they no longer have internal chip resonant modes. For lasers that use AR-coated diodes, we have been able to use the grating to tune the frequency over a range of 20 nm, compared to 3-4 nm for non-AR-coated lasers.

Our grating lasers put out 15-30 mW, sufficient for all the laser applications in our

experiment with the exception of the magneto-optical trap (MOT), which requires significantly more power. In order to generate sufficient power for the MOT, we use “slave” lasers, which are diode lasers without external gratings that are injected with a small amount of laser light coming from a “master” grating laser. The injection is achieved by coupling several mW of the master laser’s output in through the polarizing beam splitter in the middle of a two-stage optical isolator at the output of the slave laser. The injected light seeds the gain of the slave lasers so that they lase at the exact frequency of the master. To monitor the relative frequencies of the master and slave lasers, small fractions of the beams are picked off and passed through a modulating Fabry-Perot cavity, with a photodiode monitoring the output. Outside the injection region, as the current of the slave is tuned, the frequency tunes with it; when the injection region is reached, the Fabry-Perot transmission peak from the slave overlaps the peak from the master and does not tune with current.

In addition to not having grating-related losses, slave lasers can use lower-quality laser diodes, which are available at higher output power than high-quality diodes. Thus we are able to build slave lasers that put out in excess of 100 mW.

2.2.2 Laser Frequency Stabilization

To stabilize the frequency of our grating lasers we use a two-tier active stabilization chain. A reference laser is locked to a Rb atomic transition via a vapor cell. Beams from the other grating lasers, the “MOT master” and the “repumper,” are overlapped with beams picked off from the reference laser in fibers coupled to fast photodiodes, which measure the beat note between the reference laser and the other lasers. The beat note frequencies are used to generate the feedback signal that is converted into a laser diode current and a corresponding PZT voltage in order to stabilize the MOT master and repumper laser frequencies and tune them in real time with respect to the reference laser.

Thermal motion of atoms at room temperature leads to a broadening of spectral lines via the Doppler effect, and the Doppler-broadened linewidth of about 500 MHz is two orders of magnitude larger than the Rb natural linewidth. Thus, in order to

use room-temperature vapor to lock a laser more tightly than the natural linewidth, it is necessary to implement a Doppler-free locking scheme. In our case, the reference laser is locked to the $|F = 3, m_F = 3\rangle \rightarrow |F' = 4, m'_F = 4\rangle$ transition of the D₂ line in ⁸⁵Rb (⁸⁵Rb and ⁸⁷Rb coexist in our vapor cell and this transition turns out to be the most convenient) using a Doppler-free polarization spectroscopy lock [45]. In this locking scheme, a linearly polarized probe beam is sent through a room-temperature vapor cell with a counterpropagating circularly polarized pump beam derived from the same laser. At the frequency where both beams are resonant with the same set of atoms (the zero-velocity atoms if both beams are resonant with the same transition), the two circularly polarized components of the probe see different amounts of atomic saturation due to the presence of the pump, which will both rotate the probe's polarization axis and make it elliptically polarized in a frequency-dependent way. If the probe light coming out of the sample is measured after a linear polarizer that is close to perpendicular with the initial probe polarization, the result will be a dispersive line shape centered on the transition frequency. This signal can then be directly converted into a voltage used in a negative feedback loop to stabilize the reference laser current and PZT voltage so as to make the laser frequency resonant with the atomic transition. To improve the lock, we mount a permanent magnet near the vapor cell, which defines a quantization axis that optimizes our signal and makes the lock insensitive to stray magnetic fields.

The portion of the reference laser light that is not used for the polarization spectroscopy lock is split up between the beat note locks for the MOT master and repumper lasers. The two locks are very similar – the beat note signal from a fast photodiode is amplified, divided down in frequency to less than 300 kHz, and fed into a phase-locked loop (PLL) integrated circuit. The PLL's internal voltage-controlled oscillator (VCO) locks to the beat note frequency, and the VCO control voltage is then proportional to the difference between the beat note frequency and the free-running VCO frequency. By adding an offset to this voltage and feeding it back to the laser frequency control via the laser current and PZT voltage we can tune the laser to a frequency set by the offset voltage. Changing the offset voltage allows us to tune the lasers in real time by more than a GHz, with the maximum speed of tuning

given by the time constant of the feedback loop.

To illustrate the frequencies involved, the ^{87}Rb D_2 transition level diagram with the frequencies of the MOT and repumper lasers is shown in Fig. 2.2. The reference laser wavelength is 780.244 nm, which puts it 1.1 GHz to the blue of the MOT frequency and 5.5 GHz to the red of the repumper frequency. Our digital frequency dividers do not work well above about 2 GHz, so the repumper lock has an additional step of taking the beat note and mixing it with a 6.8 GHz signal from a stable microwave oscillator, so that the difference frequency of 1.3 GHz is then low enough to be divided down. Note that the frequency difference between the MOT and repumper lasers is slightly less than the 6.8 GHz hyperfine splitting because they couple to different excited states.

In order to minimize mechanical vibrations, the lasers are placed on rubber pads and mounted on the optics table using nylon screws. The entire laser system is also enclosed in a plexiglass box in order to reduce temperature drifts, air currents, and sound pickup.

The lasers are switched using acousto-optical modulators (AOMs) that can turn on and off in less than 1 μs and provide an attenuation of >60 dB in the first order, in combination with mechanical shutters that take ≈ 1 ms to open and close. The AOMs are Isomet 1205C-1 modulators that we drive at a fixed frequency of 80 MHz using Isomet 232A-1 drivers; in order to maximize attenuation, we modify the drivers by putting in an external radiofrequency (RF) switch between the driver's oscillator and its amplifier.

All the lasers that interact with the atoms are coupled into polarization-maintaining single-mode fibers. This both allows us to decouple the alignment of the beams on the fiber input and output sides and gives us clean circular Gaussian beams to work with. The trade-off is that the coupling efficiency for the fibers is around 50%, but we have sufficient laser power available to overcome this limitation.




Figure 2.2: Hyperfine structure of the ^{87}Rb D_2 transition showing the levels used and frequencies for the MOT and repumper lasers. ^{87}Rb has a nuclear spin of $I = 3/2$. The g_F factor refers to the Zeeman shift of the energy levels at low magnetic field [46].

2.3 Vacuum System and Magneto-Optical Trap

The ^{87}Rb atoms used in our experiment are collected and cooled in a MOT inside a vacuum chamber. A schematic of the main vacuum chamber showing the position of the MOT beams and the atom chip is shown in Fig. 2.3. The main vacuum chamber has 10 ports, with the large front and back ports and the six small lateral ports having anti-reflection coated windows, the bottom port used to attach an auxiliary lower chamber on which the vacuum pumps are mounted, and the top port used to attach the assembly that carries the atom chip and the vacuum electrical feedthroughs. The custom-made vacuum chamber and the vacuum components were purchased from MDC and assembled using ConFlat (CF) flanges and single-use copper gaskets.

The vacuum chamber is continuously pumped by an ion pump (Varian VacIon *Plus* 75, run at 3 kV) and, when necessary, by a titanium sublimation pump (Varian 916-0017), which maintain a base pressure of a few $\times 10^{-10}$ mbar. The lower vacuum chamber also has a valved-off Kwik-Flange connector, which allows the attachment of a turbomolecular pump that is used to pump the vacuum system down from atmospheric pressure to the 10^{-5} mbar threshold where the ion pump can be used.

The atoms source used in our experiment is Rb getters mounted several centimeters away from the center of the main vacuum chamber, which emit Rb gas when they are heated by a current. We typically use getter currents in the range of 5-5.8 A, and each getter lasts for many hundreds of hours of continuous use.

The single-chamber design of the experiment and also the presence of the Rb getters fairly close to the trapping region limit the pressure we can achieve, and therefore the lifetime of the atoms. We have been able to observe magnetic trap lifetimes of up to 10 s in the center of the vacuum chamber and up to 4 s in the region close to the atom chip. While these lifetimes are short compared to some other ultracold atom experiments, they are sufficiently long that they are not a serious limitation for the kinds of experiments we have done and are interested in doing.

The MOT is located in the center of the vacuum chamber. There are three laser beams employed, two diagonal and nearly perpendicular to each other, in the y - z




Figure 2.3: Schematic of the main 10-port vacuum chamber and the MOT beams. A third MOT beam comes out of the page, along x , centered on the intersection of the other two. The vacuum chamber is 30 cm high with view ports ranging from 4 cm to 12.5 cm in diameter. The chamber is pumped from below. The chip is located above the MOT, mounted upside-down on a hollow cylinder attached to the chamber's top flange and water-cooled.

plane, and one horizontal, along x , all with diameter of 2.5 cm and each retroreflected to produce a total of six beams incident on the MOT region. The MOT light is circularly polarized, and the retroreflection mirrors have $\frac{\lambda}{4}$ wave plates mounted on them, so that the retroreflected beams have the same circular polarization with respect to their direction of propagation as the original ones. Each of the three MOT beams has approximately 15 mW of laser light coming from the MOT slave lasers in it, and one of them (the one along x) also has about 2 mW of repumper light.

The magnetic field gradients necessary for the MOT are generated by a pair of 93-turn solenoid coils wound around the front and back ports of the main vacuum chamber and used in the anti-Helmholtz configuration, with currents flowing in opposite directions. The coils are wound with hollow metal tubing to allow for cooling water to flow through the center of the wire.

The zero of the field generated by the MOT coils, which is where atoms are collected and cooled, is located in the center of the region of overlap of the MOT beams. The MOT laser is tuned on the order of one natural linewidth Γ to the red of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition in ^{87}Rb , and the circularly polarized MOT light incident on the atoms both cools them through polarization gradient cooling and, in combination with the magnetic field gradient produced by the MOT coils, provides the restoring force that traps them at the magnetic field zero. The Doppler shifts, light shifts, and Zeeman shifts in a MOT tune the atom closer to resonance with the laser beam that will cool it or push it toward the magnetic field zero.

While the MOT laser is resonant with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition, it is powerful enough to cause occasional (once about every 10^4 scattering events) off-resonant excitations on the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition, which can cause the atoms to decay to the $|F = 1\rangle$ ground hyperfine manifold, which is not trapped by the MOT. To counter this effect, the repumper laser tuned to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition is used. Any atoms that decay to the $|F = 1\rangle$ ground state are thus quickly pumped back to the $|F = 2\rangle$ ground state via the $|F' = 2\rangle$ excited state, so they remain trapped.

In addition to the MOT coils, there is another pair of 53-turn water-cooled coils wound around the top and bottom ports of the main vacuum chamber used in the

Helmholtz configuration, with currents flowing in the same direction, to generate a spatially constant magnetic field in the vertical direction. These coils, called the “transfer” coils, can be used to displace the magnetic field zero generated by the MOT coils by several centimeters, which is necessary for loading the magnetic microtrap.

Finally, there is a set of six non-water-cooled coils mounted in a cube around the main vacuum chamber that is used to create constant offset fields in all three directions. These coils, referred to in pairs as the X -, Y -, and Z -bias coils, can both cancel stray magnetic fields and displace the magnetic field zero by several millimeters in any direction.

2.4 Atom Chip

2.4.1 Chip Mount and Auxiliary Coils

A diagram of the chip mounting system can be seen in Fig. 2.4. The atom chip is glued to a copper chip carrier using vacuum-compatible epoxy. The chip carrier is screwed to the “top piece”, a copper cylinder that is welded to a steel CF flange so that it can be attached to the main vacuum chamber. The top piece is hollow, with a threaded opening at the top used to attach tubing carrying cooling water.

The top piece also incorporates an auxiliary magnetic coil, the “intermediate” coil, which consists of 77 circular turns of Kapton-coated copper wire. The chip carrier itself incorporates another auxiliary coil, called “ $Q5$,” a 6-turn coil located 2 mm above the chip’s surface. $Q5$ is elongated along y , which is the weak confinement direction of the microtrap, with an aspect ratio of about 3.5:1. Together, the intermediate and $Q5$ coils allow us to adiabatically load the atoms from a large magnetic trap generated by the MOT coils into a microtrap generated by the microfabricated wires on the chip.

The chip used in our initial series of experiments is patterned on a 300 μm -thick silicon substrate coated with a 1 μm -thick electrically insulating Si_3N_4 layer. The Cu conductors are 2.15 ± 0.20 μm thick, fabricated by photolithography and wet etching and coated with a 100 nm Au layer in order to allow wire bonding to gold-coated




Figure 2.4: Vacuum chamber top piece and atom chip carrier. The top piece, which is attached to a CF flange, is hollow and has pipe thread at the top to allow for water cooling. The atom chip is glued to the chip carrier, which is then screwed to the top piece. The intermediate and Q5 coils assist with loading atoms into the microtrap. The holder for the microtrap imaging lens is mounted on the chip carrier.

contact pads on the chip carrier. Kapton-coated copper wires are soldered to the contact pads with flux-free vacuum-compatible solder and then soldered and crimped to connectors attached to a Macor piece mounted near the chip. Another set of Kapton-coated copper wires goes from the connectors on the Macor piece to vacuum feedthroughs.

2.4.2 Chip Current Coils

The layout of the chip is shown in Fig. 2.5. The radial (also referred to as “transverse” or “quadrupole”) confinement is generated using a 2-wire geometry by one of several nested rectangular conducting loops in conjunction with a bias field perpendicular to the chip surface. The primary loop used in this experiment is $Q2$, which has conductors that are $2\ \mu\text{m}$ thick and $20\ \mu\text{m}$ wide with centers separated by $100\ \mu\text{m}$. The axial (also referred to as “longitudinal” or “Ioffe-Pritchard”) confinement is generated by a $25\ \mu\text{m}$ -thick, $150\ \mu\text{m}$ -wide gold ribbon wire-bonded between two pads located on opposite sides of the radial-confinement-generating loops and running $155\ \mu\text{m}$ above the chip surface, parallel to it. There are several such ribbons, and our initial intention was to trap atoms between two of them, but it turned out that they were spaced too closely for us to be able to load between two of them, and we ended up trapping the atoms outside of the outermost one (in the y -direction), labeled “ I_{+3} ”, and using the external y gradient generated by the MOT coils to confine them on the other side.

In addition to $Q2$, there are two larger chip quadrupole loops, which, while utilized during preparatory stages, are not used in the experiments described below. $Q3$ has conductors that are $300\ \mu\text{m}$ apart (center-to-center) and $150\ \mu\text{m}$ wide, while $Q4$ has conductors that are $1\ \text{mm}$ apart and $500\ \mu\text{m}$ wide. The loops are nested so that $Q2$ is inside $Q3$, which is inside $Q4$. Also, there is an even smaller loop, $Q1$, which has conductors that are $30\ \mu\text{m}$ apart and $10\ \mu\text{m}$ wide, which was damaged so that the current path was cut; $Q1$ carries no current, but is used as a conducting surface for measurements described in Chapter 4.

The maximum current limits of the chip wires were determined outside of the vacuum chamber by monitoring the wires’ resistance when current flowed through them.

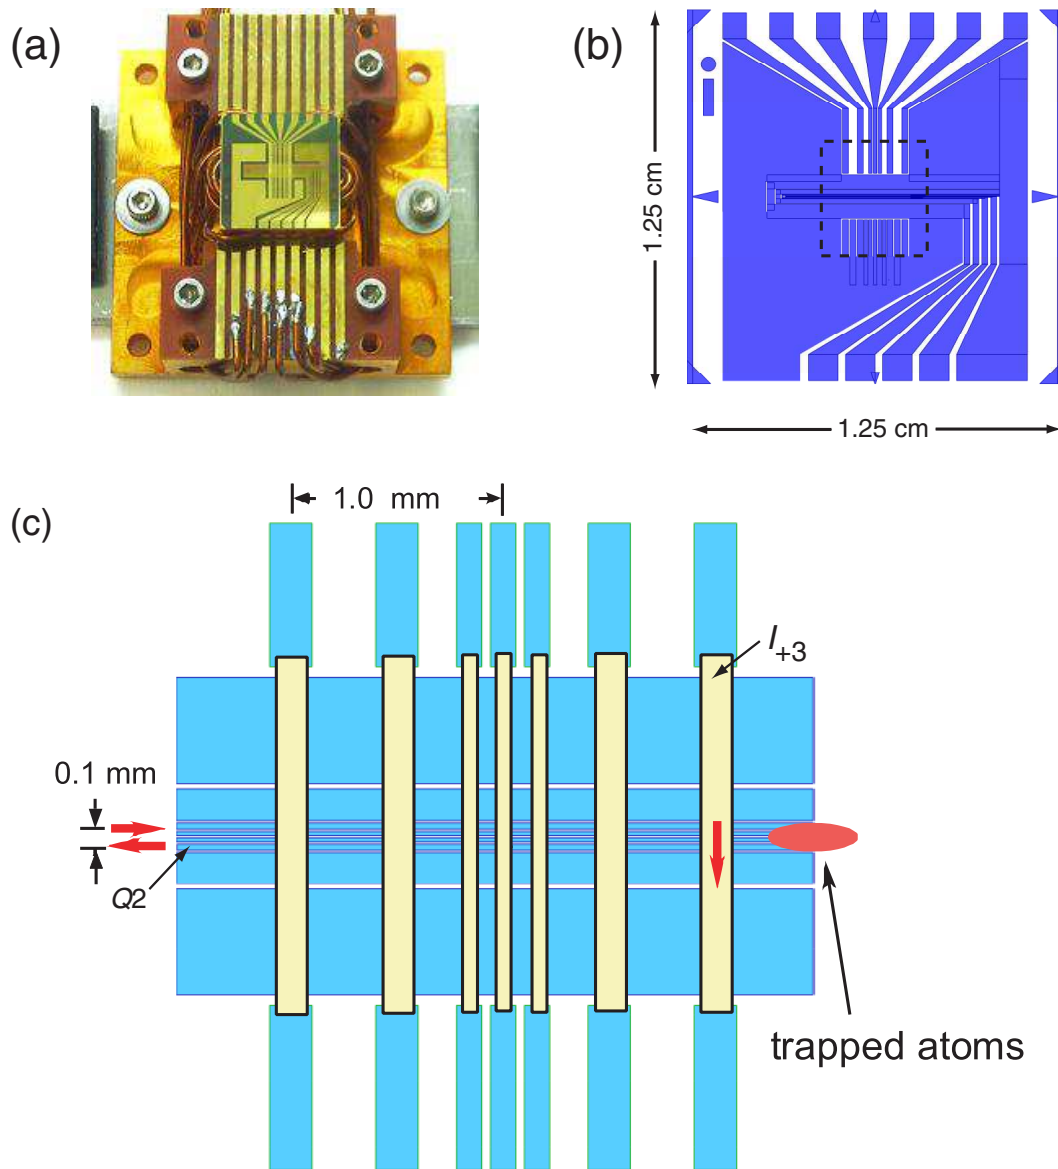


Figure 2.5: (a) Photograph of the atom chip mounted on the chip carrier. (b) The mask used to create the chip wires; conductors are shown in blue. (c) Schematic of the central section of the chip, marked by the dashed box in (b), showing the trapping wires and the location of the chip microtrap. The conductors on the chip's surface, which include the quadrupole wires and the pads for bonding the Ioffe-Pritchard wires are in blue; the Ioffe-Pritchard ribbons are in yellow. The quadrupole loop and Ioffe-Pritchard ribbon primarily used in the experiments are labeled $Q2$ and I_{+3} , respectively, and the directions of current flow through them are indicated.

The resistivity of copper has an approximately linear dependence on temperature in the range of temperatures encountered in our experiment, with a 40% increase in resistivity when the temperature rises from 300 K to 400 K [47]. During preliminary tests outside of the vacuum chamber, $Q2$, the primary quadrupole wire used, heated up by 100 K in steady state when a current of 1 A was flowed through it. Since $Q2$'s duty cycle in our experiment was always less than 50%, 1 A was determined to be a safe maximum current, even when taking into account worse heat dissipation in vacuum compared with air. The maximum safe current for the Ioffe-Pritchard ribbon was determined to be in excess of 2 A.

The 1 A current in $Q2$ corresponds to a current density of 2.5×10^6 A/cm², an extremely large current density in ordinary contexts. The ability of the chip conductors to sustain such high current densities is due to their excellent thermal coupling to the substrate, which is aided by having wires with a low and wide cross section.

It should be noted that the heat dissipation on this chip decreased significantly over the three-year period of the chip's use. We attribute this to the deterioration of the epoxy used to glue the chip to the chip carrier, possibly accelerated by excessive heating of the chip carrier on one occasion when the current-regulating circuit malfunctioned. When we removed the chip carrier from our vacuum chamber at the end of its use cycle, the chip came away from the carrier with very little effort, which confirms the deterioration of the epoxy.

2.5 Imaging

The imaging in our experiment is performed by two complementary methods – an amplified photodiode for imaging the atom cloud far away from the chip, and a CCD camera for imaging atoms trapped near the chip surface.

The photodiode, which has two gain settings for output voltage as a function of incident optical power, 1 kV/W and 10 kV/W and a time resolution of 700 μ s, measures light collected by a large lens placed directly outside one of the vacuum chamber's side ports, with a mirror between the lens and the photodiode to optimize imaging at different trap positions. The photodiode signal is calibrated using atomic




Figure 2.6: CCD absorption imaging schematic. Dashed lines indicate the focusing of the dark-field image. The final lens before the camera is on a 1D translation stage to optimize focusing. A band-pass filter in front of the CCD and a black plastic tube that encloses the imaging path between the CCD and the final focusing lens are used to minimize stray light.

fluorescence on resonance at intensities much higher than saturation intensity combined with a calculation of the photodiode’s effective solid angle. Specifically, given an effective solid angle $4\pi\alpha$ and a transition with wavelength λ and natural linewidth Γ , the optical power incident on the photodiode when the transition is fully saturated will be $P = N\alpha(\Gamma/2)(hc/\lambda)$, where N is the atom number.

The camera we use is an Apogee AP260E cooled CCD camera, with an imaging area of 10.2×10.2 mm containing 512^2 pixels and quantum efficiency of 37% at 780 nm. A schematic of the CCD absorption imaging can be seen in Fig. 2.6. A collimated imaging beam, whose size can be varied by changing the collimation, comes in parallel to the chip surface and is partially absorbed by the atoms. The dark-field image is focused and magnified by 3 lenses, the first of which is in vacuum right next to the chip, and the last of which is on a 1D translation stage oriented along the beam path.

The in-vacuum lens is a low-aberration graded-index glass lens (Gradium from

LightPath Technologies, $f/\# = 2.5$) with a focal length of 2.2 cm mounted on the chip carrier, ~ 2.4 cm from the center of the chip, vertically centered on the chip surface.

We have used several different configurations for the later two focusing lenses and camera, depending on what imaging magnification we wanted.

In order to minimize the amount of stray light that enters the CCD, we put a band-pass filter in front of the CCD and enclose the imaging path between the CCD and the final focusing lens in a tube made from black plastic cloth. For many of the measurements, we also turn off the room lights.

The camera absorption imaging is performed by taking two images, a signal image that contains the atoms and a background image taken several hundred ms later, when the atoms are gone from the imaging region. The absorption signal extracted from comparing the two images is converted into an atom number after taking into account magnification, background dark counts, and fluctuations in imaging beam intensity between the signal and background shots. Magnification is measured by looking at positions of wire bonds in the camera image and using the known distances between them, which were precisely determined by the fabrication process. The dark count rate, which accounts for both the intrinsic dark count of the camera and any stray room light that enters the camera, is determined by looking at the camera signal with the imaging beam turned off. The shot-to-shot fluctuations in imaging beam intensity are measured by looking at a region of the image which is inside the imaging beam but has no atoms in it and comparing the measured intensity in that region between the signal and background shots.

The attenuation of a beam due to the presence of N atoms is given by $I = I_0 e^{-Nd}$, where I_0 is the intensity without atoms and d is the optical depth of a single atom. Solving for the number of atoms gives $N = -\frac{1}{d} \ln \frac{I}{I_0} = \frac{1}{d} \ln \frac{I_0}{I}$. The atom signal in a pixel for our camera imaging is thus calculated as:

$$N = \frac{A}{\sigma} \ln \left(\frac{I_{B,1} - I_{D,1}}{I_{B,2} - I_{D,2}} \times \frac{I_2 - I_{D,2}}{I_1 - I_{D,1}} \right), \quad (2.1)$$

where I_D is the measured average dark count rate per pixel, I_B is the average non-dark counts per pixel in the region without atoms, I is the number of counts in

the pixel in question, and the subscripts 1 and 2 refer to the signal and background images, respectively. The conversion constant $\frac{A}{\sigma}$ is the ratio between the area imaged onto each pixel, which can be determined from the magnification, and the effective scattering cross section for the atoms, given by $\frac{1}{3} \times \sigma_{max}$, where $\sigma_{max} = \frac{3}{2\pi} \lambda^2$ is the maximum scattering cross section for the transition and the factor of $\frac{1}{3}$ accounts for the fact that we're imaging with mostly π light. The uncertainty in the calibration is due mostly to uncertainty in this Clebsch-Gordan/polarization reduction factor, which may be wrong by as much as 50% but is only an absolute calibration effect, while we are primarily interested in relative atom number measurements.

Due to diffraction from the chip surface, we were not able to obtain high-quality CCD absorption images closer than about $100 \mu\text{m}$ from the chip, though we were able to see distorted images down to less than $50 \mu\text{m}$ away.

Chapter 3

Microtrap Loading and Evaporation to BEC

3.1 Loading the Magnetic Microtrap

After collecting atoms in the MOT for several seconds (typically 5-10 s), we quickly increase the MOT current from 9 A, which corresponds to gradients of 5 G/cm in the x and z directions and 10 G/cm along y , to 50 A, which increases the gradients by the same factor; at the same time, the MOT laser frequency is detuned from 12 MHz to the red of the transition to 36 MHz to the red. This combination of increasing magnetic field gradient and laser detuning produces a “compressed MOT” [48], which allows higher densities in the MOT. The compressed MOT has a much lower loading rate, so loading into a normal MOT for a long time, and then compressing it for a very short time (30-50 ms) allows us to achieve both large numbers and high densities. In addition, we increase the transfer coil current when we compress the MOT, which moves the MOT vertically from the center of the MOT beams to their upper edge, closest to the chip.

Once the atoms are compressed, we extinguish the MOT light and pump the atoms into the magnetically trappable (low-field seeking) $|F = 2, m_F = 2\rangle$ ground state (in the MOT, the atoms occupy a mixture of magnetic sublevels). The optical pumping is done by coupling in from below a collimated, σ^+ -polarized beam which is derived




Figure 3.1: Diagram of the optical pumping levels and transitions. The σ^+ -polarized optical pumping beam (red) tuned to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition drives atoms into the $|F = 2, m_F = 2\rangle$ magnetically trappable ground state. Once the atoms are in $|F = 2, m_F = 2\rangle$, they no longer have a resonant interaction with the optical pumping beam (pink) since there is no corresponding excited state. The repumper (orange) drives atoms that decay to the $|F = 1\rangle$ hyperfine ground state manifold back into the $|F = 2\rangle$ hyperfine ground state manifold.

from the MOT master laser and tuned to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition. At the same time, we ramp up the current in $Q5$, the coil embedded in the chip carrier, to displace the magnetic field zero downward too quickly for the atoms to follow, so that the atoms experience a uniform nonzero magnetic field that helps define the quantization axis for the pumping beam polarization. The advantage of pumping on the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition instead of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition, which would also transfer the atoms into the $|F = 2, m_F = 2\rangle$ ground state, is that $|F = 2, m_F = 2\rangle$ is a dark state – once the atoms are pumped there, they no longer have an excited state to make a σ^+ transition to, so they stop scattering, and being heated by, pump photons. However, using this transition also means that, while being optically pumped, the atoms can decay into the $|F = 1\rangle$ hyperfine ground state manifold, in which case they need to be repumped back. Thus the repumper laser needs to be on during optical pumping; we use the repumper light mixed in one of the MOT beams. The optical pumping transitions are shown in Fig. 3.1.

The optical pumping sequence takes 200-400 μs , and we end up loading as much as 65% of the atoms in the MOT into the initial magnetic quadrupole trap formed by the MOT coils (compared to about 15% without optical pumping). Over the next 400 ms, the magnetic trap is adiabatically compressed and deformed into a Ioffe-Pritchard microtrap, which combines a quadrupole trap in the x and z dimensions with an offset field and two endcap gradients in the y direction. The trap is moved toward the chip by simultaneously changing currents in the transfer coil, the intermediate coil, and $Q5$. At the same time, currents in the Ioffe-Pritchard ribbon, I_{+3} , and the Y -bias coils, move the magnetic trap minimum in the y direction. When the trap gets close enough to the surface, it is trapped in the x - z plane by the $Q2$ gradient, and the Ioffe-Pritchard trap is formed, with gradients from I_{+3} and the MOT coils serving as the endcaps in the y direction and the current in the Y -bias coils determining the offset field in the trap.

The microtrap loading procedure is self-centering – the atoms are forced to follow the magnetic field zero/minimum from the macroscopic magnetic trap into the microtrap. The aspect that needs to be optimized is having sufficient confinement and trap depth along the way to keep all the atoms trapped.

We begin with around 3×10^7 atoms collected in the MOT after 8 s of loading, which gives us around 2×10^7 $|F = 2, m_F = 2\rangle$ atoms in the initial magnetic trap. Of these, we end up loading about 15%, or 3×10^6 atoms, into the microtrap, with a temperature of 300 μK and a density of 2×10^{12} cm^{-3} , which correspond to a phase-space density, given by $n\lambda_{th}^3$, where n is the 3D number density and $\lambda_{th} = \sqrt{\hbar^2/(2\pi m k_B T)}$ is the thermal de Broglie wavelength, of 3×10^{-6} .

3.2 Microtrap Characteristics

The final magnetic trap is located 50 μm from the chip surface, where the gradient from $Q2$ is largest, and approximately 400 μm on the outside of I_{+3} in the y direction. For the experimental $Q2$ current of 0.6 A, the x and z gradients are 4.9 kG/cm.

The vibration frequency for atoms in a harmonic magnetic trap is given by balancing the magnetic potential energy with the mechanical energy of the oscillation:

$$\mu_B \Delta B (g_F m_F) = \frac{1}{2} m \omega^2 A^2 \quad (3.1)$$

where $\mu_B = h \times 1.4$ MHz/G is the Bohr magneton, ΔB is the difference between the maximum and the minimum magnetic field experienced by the atom, g_F is the Landé g -factor for the level (for the $|F = 2, m_F = 2\rangle$ ground state in ^{87}Rb , $g_F m_F = 1$), m is the mass of the atom, ω is the vibration frequency, and A is the amplitude of the motion. For motion perpendicular to the trap offset field with a trap gradient b_z and an offset field b_0 , $\Delta B = \frac{1}{2} \frac{b_z^2}{b_0} A^2$, so the transverse vibration frequency is given by

$$\omega_z = \sqrt{\frac{\mu_B b_z^2}{b_0 m}}. \quad (3.2)$$

For typical offset fields of 1.5 G, the transverse vibration frequencies in our microtrap are $\omega_x = \omega_z = 2\pi \times 5.1$ kHz. For motion along the direction of the trap offset field, $\Delta B = b_y'' A^2$, so the axial vibration frequency is given by,

$$\omega_y = \sqrt{\frac{2\mu_B b_y''}{m}}, \quad (3.3)$$

where b_y'' is the magnetic field curvature (which is actually twice the second derivative of the magnetic field). Using a numerical simulation of the magnetic fields produced by our coils, we obtain $b_y'' = 1200$ G/cm², which corresponds to an axial vibration frequency of $\omega_y = 2\pi \times 63$ Hz. The aspect ratio of the atom cloud is inversely proportional to the ratio of the corresponding vibration frequencies; in our situation, this leads to a cigar-shaped cloud elongated along y , with an aspect ratio of about 80:1.

Once we evaporatively cool down to a small, cold cloud of atoms, we are able to measure the transverse vibration frequency with high sensitivity by applying a radio-frequency (RF) modulation to Q_5 and observing which RF value induces resonant loss in the trap. A result of this measurement is shown in Fig. 3.2. There are two clear resonance dips, a broad one at a little over 5 kHz and a narrower one around 3.5 kHz. We attribute the broad dip to the vibration frequency of $|F = 2, m_F = 2\rangle$ atoms and the narrow one to the vibration frequency of residual $|F = 2, m_F = 1\rangle$




Figure 3.2: Measurement of the transverse vibration frequency in the microtrap. Atoms evaporatively cooled to around $5 \mu\text{K}$ were excited by RF modulation coupled in through Q_5 , and atom number was measured after 500 ms of applied RF.

atoms in our trap, which should be a factor of $\sqrt{2}$ lower.

Given the microtrap vibration frequencies, the initial-phase space density, and the trap depth of around 10 G, it is possible to estimate the maximum number of atoms that can be loaded into the microtrap. The total number energy states available in the trap can be approximated by

$$N_E \approx \frac{(\mu_B B_{max})^3}{(\hbar\omega_x)(\hbar\omega_y)(\hbar\omega_z)} D \quad (3.4)$$

where B_{max} is the trap depth and D is the phase-space density. Using the numbers for our microtrap, we get $N_E = 5 \times 10^6$, which is within a factor of two of our loaded atom number. Thus we believe that the number of atoms loaded into the microtrap is limited by the number of available states.

3.3 Evaporation to BEC

To achieve Bose-Einstein condensation, we employ microwave forced evaporative cooling. Evaporative cooling reduces the temperature of a sample by removing the atoms from the high-temperature tail of the Maxwell-Boltzmann distribution, thus reducing the average kinetic energy per atom for the remaining sample. As the sample cools, the removal threshold is reduced to keep pace, so that the removal of the hot atoms continues until the desired temperature is reached. Ref. [49] provides an excellent overview of the various aspects of evaporative cooling of trapped atoms.

The evaporation process is exponential – atom number, temperature, density, and phase-space density go down exponentially with time, where time is measured in units of the number of elastic collisions an atom experiences (since the density, and therefore the collision rate, increase during the evaporation process, the exponential timescale of evaporation becomes shorter as evaporation goes on). An important dimensionless parameter that characterizes the efficiency of evaporation is

$$\gamma = -\frac{d(\ln D)}{d(\ln N)} \quad (3.5)$$

which is the ratio between the (positive) exponential rate of change of the phase-space density and the (negative) exponential rate of change of the atom number.

Another parameter that characterizes the evaporation process is η , which is a measure of how far out on the Maxwell-Boltzmann tail the forced removal of atoms takes place (i.e., for a sample of atoms with average kinetic energy $k_B T$, atoms with energy above $\eta k_B T$ are removed). Since the tail of the thermal distribution is a decaying exponential, at high η evaporation removes linearly hotter atoms but exponentially fewer of them per unit time than evaporation at lower η . Thus, if there are no other time scales in the experiment, evaporation at high η is always more efficient; in practice, the finite lifetime of atoms in the trap sets an upper limit on the evaporation time, and thus on η . At the same time, evaporative cooling relies on thermalization to continually refill the hot tail of the thermal distribution, and the time scale for thermalization is set by the elastic collision rate, $\Gamma_{el} = n\sigma v\sqrt{2}$, where n is the 3D number density, σ is the collision cross section, and v is the atom thermal velocity.

In this way, the collision rate sets a lower limit on the evaporation time.

The compression of the atom cloud due to the microtrap's high transverse confinement provides us with a large initial elastic collision rate, around 1300 s^{-1} , while the microtrap's background-limited lifetime of about 3 s means we have to evaporate relatively quickly. Given these limits, we choose to evaporate at $\eta \approx 5\text{-}6$.

The mechanism we use for the removal of high-energy atoms is resonant microwave transitions. Atoms trapped in the $|F = 2, m_F = 2\rangle$ hyperfine ground state can make a microwave transition to the $|F = 1, m_F = 1\rangle$ hyperfine ground state, which is high-field seeking, or anti-trapped; these atoms are then lost from the microtrap. By tuning the microwave frequency to the blue of the zero-field $|F = 2, m_F = 2\rangle \rightarrow |F = 1, m_F = 1\rangle$ transition frequency by an amount $\Delta\nu = 1.5\eta k_B T/h$, we can selectively remove only those atoms that can climb the magnetic potential high enough to experience a Zeeman shift that brings them into resonance with the microwaves, which are exactly the atoms with energy $\eta k_B T$ (the factor of 1.5 in the detuning comes from the fact that the $|F = 1, m_F = 1\rangle$ state sees a shift of $\Delta E = -0.5\eta k_B T$ at the same field at which the $|F = 2, m_F = 2\rangle$ state sees a shift of $\Delta E = \eta k_B T$).

We apply the microwaves via a horn placed against one of the vacuum ports, about 10 cm away from the atoms. The frequency is derived from a 100-180 MHz voltage-controlled oscillator mixed with a 100 MHz stable oscillator and then mixed again with a 6.8 GHz microwave source, giving us a tuning range of 6834-6880 MHz (the zero-field microwave transition is at 6834 MHz). The microwave power out of the horn is around 1 W, which corresponds to a microwave Rabi frequency in the microtrap of around 1 kHz, much slower than the decay rate of coherences between the $|F = 2, m_F = 2\rangle$ and $|F = 1, m_F = 1\rangle$ states, which means that once the atoms make a transition to $|F = 1, m_F = 1\rangle$, they leave the trap before they can oscillate back to $|F = 2, m_F = 2\rangle$.

We also use microwave-induced loss to measure the temperature of our atom cloud. We scan the microwave frequency over the energy profile of the cloud over many otherwise identical iterations of the experiment, and use short microwave pulses to selectively remove atoms with a certain Zeeman shift, after which we measure the atom number. By converting microwave frequency to atom energy, we can then make

a spectrum of the energy distribution of the atoms in the cloud, and fit it to extract the atom number and temperature and the value of the offset field.

For a cloud with a temperature small compared to Zeeman shift from the trap offset field, the magnetic field from a Ioffe-Pritchard trap can be treated as harmonic [50]. The potential energy density of states is then given by

$$\rho_{U'} \equiv \int \delta(U' - U(\mathbf{r})) d^3r \sim \int \delta(U' - |\mathbf{r}|^2) d^3r \sim \sqrt{U'}, \quad (3.6)$$

and the fit function for the microwave spectra, which is the Maxwell-Boltzmann distribution weighted by the density of states, is

$$N(U) = N_0 - A\sqrt{U - U_0}e^{-\frac{U-U_0}{k_B T}}, \quad (3.7)$$

where N_0 is the total atom number, U_0 is the Zeeman shift at trap bottom, and A is a fit parameter that depends on the strength and duration of the microwave pulse.

When the temperature is comparable to or larger than the Zeeman shift from the trap offset field at the bottom of the trap, the trap is linear in the two transverse dimensions and harmonic in the axial dimension. We find, however, that the microwave spectrum can be fit well by a function that assumes linear confinement in all dimensions, which corresponds to

$$\begin{aligned} \rho_{U'} &\sim \int \delta(U' - |\mathbf{r}|) d^3r \sim U'^2 \\ N(U) &= N_0 - A(U - U_0)^2 e^{-\frac{U-U_0}{k_B T}}. \end{aligned} \quad (3.8)$$

The evaporation process is optimized by dividing it into 300 ms steps during each of which the microwave frequency is ramped linearly, with the endpoint chosen to yield $\eta \approx 5-6$.

The parameters obtained for each step of the evaporation to BEC are shown in Table 3.1, and a plot of the phase-space density versus the atom number during evaporation is shown in Fig. 3.3. Diffraction from the chip surface attenuates and distorts absorption images taken with the CCD camera of clouds in the microtrap

where we perform the evaporation, 50 μm from the chip surface, so we move the trap to 100 μm away for imaging and measuring atom numbers.

Atom Number	T [μK]	Trap Depth [μK]	Density [cm^{-3}]	Phase-Space Density
2.9×10^6	270	1400	4.7×10^{12}	7.1×10^{-6}
9.2×10^5	140	770	6.4×10^{12}	2.6×10^{-5}
4.85×10^5	78	430	1.1×10^{13}	1.0×10^{-4}
2.5×10^5	52	260	1.3×10^{13}	2.3×10^{-4}
1.7×10^5	38	190	1.6×10^{13}	4.6×10^{-4}
8.5×10^4	26	150	1.5×10^{13}	7.3×10^{-4}
5.0×10^4	18	87	1.7×10^{13}	1.5×10^{-3}
4.7×10^4	9.5	63	4.4×10^{13}	9.8×10^{-3}
4.0×10^4	5.0	32	1.0×10^{14}	0.06
1.45×10^4	3.0	21	8.5×10^{13}	0.10
1.05×10^4	2.3	14	8.7×10^{13}	0.16
6.0×10^3	1.1	7.7	1.7×10^{14}	1.0
4.3×10^3	0.52	2.0	3.4×10^{14}	5.9

Table 3.1: Relevant parameters during the evaporation process. The exact values can drift significantly day-to-day; the above numbers were compiled over a period of several days when the evaporation ramp was initially optimized. The critical phase-space density for Bose-Einstein condensation is 2.61.

In order to show that we had achieved Bose-Einstein condensation directly (as opposed to indirectly by calculating the phase-space density), we take one-dimensional time-of-flight measurements of our cold atom clouds. To do this, we move the trap to 100 μm from the chip surface and then quickly turn off the current in the Ioffe-Pritchard ribbon, I_{+3} , releasing the confinement in the axial direction, and causing the atom cloud to propagate along y under the influence of the gradient from the MOT coils while still trapped in the x and z directions. By taking absorption images of the cloud after several ms of axial expansion, we are able to observe the axial momentum distribution of the cloud and see the formation of the condensate peak, which manifests itself as an inverted parabola on top of the Gaussian thermal distribution. One-dimensional time-of-flight images showing the final stages of the evaporation and the formation of the BEC peak, along with the corresponding fits to the linear density along the axial direction, are shown in Fig. 3.4. The evaporation




Figure 3.3: Phase-space density as a function of atom number during evaporation. The solid line is a fit to γ , the logarithmic derivative of phase-space density with respect to atom number; we obtain $\gamma = 2.1 \pm 0.2$ for the entire evaporation process.




Figure 3.4: 1D time-of-flight images of the final stages of evaporation to BEC, along with fits to density profiles along center of cloud, Gaussian for thermal distributions (blue) and parabolic for the condensate peak (red). (a) Just before condensation; trap depth= $9.5 \mu\text{K}$, $T \approx 1.1 \mu\text{K}$. (b) Condensate peaks begins to form; trap depth= $3.3 \mu\text{K}$. (c) Nearly pure BEC; trap depth= $2.1 \mu\text{K}$.

process characterized here takes 3.7 s to achieve Bose-Einstein condensation; we were later able to reduce that time to just under 3 s.

Chapter 4

Measurement of Surface Effects

4.1 Precise Positioning of Atoms

4.1.1 Microtrap Field Contours

Our magnetic microtrap is a two-wire trap, where the transverse confinement is generated by counterpropagating currents in two parallel chip wires along with an offset field applied perpendicular to the chip surface. The contours of the magnitude of the magnetic field for such a trap are shown in Fig. 4.1. (It should be noted that the real microtrap in our experiment is not quite this simple, since we have other gradients present than those due to the $Q2$ wires, and these additional gradients slightly distort the magnetic field contour map compared to that shown in Fig. 4.1. We numerically model all of our magnetic-field-generating coils in order to account for this distortion.)

To move the atoms in the transverse (x - z) plane, bias fields are applied in the x and z directions. In the region between the wires and far away from the chip, these have the intuitively expected effect – an applied magnetic field in the x direction moves the trap in the x direction, and similarly for z . Near the chip surface, however, the situation changes, and applying a magnetic field in x can, in some regions, move the trap in z , and vice versa.

Fig. 4.2 shows the contours for the two transverse components of the magnetic field, B_x and B_z , in the region between the wires and close to the chip. The contours




Figure 4.1: Contours of the magnitude of the magnetic field for the microtrap in the x - z plane. The trap at $(x = 0 \mu\text{m}, z = 50 \mu\text{m})$ is generated by two counterpropagating currents along y , flowing through the $Q2$ wires located at $(x = \pm 50 \mu\text{m}, z = 0 \mu\text{m})$, and a uniform bias field along z .

are symmetric about the point $(x = 0 \mu\text{m}, z = 0 \mu\text{m})$, referred to as the “symmetry point” or “origin” (note that due to the finite height of the chip wires, the origin is located about $1 \mu\text{m}$ away from the chip surface). At $(x \approx 0 \mu\text{m}, z \approx 15 \mu\text{m})$, the B_x contours are nearly vertical and the B_z contours are nearly horizontal, but near the chip both sets of contours curve so that at $(x \approx \pm 10 \mu\text{m}, z \approx 0 \mu\text{m})$, the B_x contours are nearly horizontal (so that a magnetic field along x will actually move the trap along z) and the B_z contours are nearly vertical (so that applying a magnetic field along z will move the trap along x).

In addition, since the magnetic field is symmetric about the origin, the magnetic field minimum that produces the microtrap will have a corresponding magnetic field minimum, called the “image trap,” located on the opposite side of the origin, as shown in Fig. 4.2. Usually, the image trap will be located inside the chip; however, if the microtrap is sufficiently close to the origin along z , the image trap can emerge out of the surface.




Figure 4.2: Contours of B_x (dashed) and B_z (solid) produced by the $Q2$ wires near the chip surface in the region between the two $Q2$ wires in the absence of any other transverse magnetic field gradients; contour spacing is ≈ 0.4 G. Note that the symmetry point for the magnetic field contours, $(x = 0, z = 0)$, is located not at the chip surface but $1 \mu\text{m}$ above it, since that corresponds to the centers of the $2\text{-}\mu\text{m}$ -high $Q2$ wires. The $Q1$ wires, which have no current flowing through them, are $10 \mu\text{m}$ wide and $2 \mu\text{m}$ high and centered at $x = \pm 15 \mu\text{m}$. Due to the symmetry of the magnetic field configuration, when the microtrap (red circle) is located above the symmetry point, there is a corresponding magnetic field minimum, or image trap, exactly opposite it (blue circle). The red dotted and dashed arcs represent the measurement trajectories used to determine the magnetic fields needed to put the microtrap on the $\Delta B_x = 0$ and $\Delta B_z = 0$ contours, respectively.

4.1.2 Calibrating Magnetic Fields at Origin

The existence of the image trap can be used to determine extremely precisely the position of the microtrap produced by a given applied magnetic field. This essentially boils down to determining what values for the homogeneous applied bias fields B_x^0 and B_z^0 put the trap at the origin, or, equivalently, what value of B_x^0 puts the trap along the contour labeled $\Delta B_x = 0$ in Fig. 4.2 and what value of B_z^0 puts the trap along the contour labeled $\Delta B_z = 0$.

The measurement takes advantage of the fact that as the microtrap is brought closer to the origin, and therefore to the image trap, atoms can overcome the barrier between the microtrap and the image trap and end up colliding with the surface. For a transverse gradient b_t and an offset field B_0 , the magnetic-barrier height's dependence on the distance d between the microtrap and the image trap is given by $\sqrt{(b_t \frac{d}{2})^2 + B_0^2} - B_0$, which can be approximated by $\frac{1}{2}(b_t \frac{d}{2})^2 / B_0$ for $b_t d \ll B_0$, as is the case here. Combined with the exponential dependence of the tail of the Maxwell-Boltzmann distribution on energy (see Eq. 3.7), this means that the rate of atom loss due to the image trap depends very strongly on d . Along a given B_x or B_z contour, d is minimized exactly when that contour crosses the line $\Delta B_z = 0$ or $\Delta B_x = 0$, respectively. Thus, B_x^0 and B_z^0 , which depend on, and thus account for, both stray fields and unknown offsets due to imperfect alignment or positioning of non-chip magnetic coils, can be determined by maximizing atom loss while scanning B_x and B_z near the origin.

To measure the atom loss as we vary B_x (B_z), an atom cloud is prepared with the desired atom number and temperature and the desired value for B_x (B_z) is applied when the microtrap is still far from the origin. Then the not-scanned transverse bias field B_z (B_x) is ramped to $\Delta B_z = 120$ mG ($\Delta B_x = 110$ mG), a value that places the trap very close to the origin and gives good contrast for the scan given the holding time. The microtrap is then held at that final position for several tens of ms, after which B_z (B_x) is ramped back, and the trap is moved far away from the chip and imaged to determine the atom number. The above procedure is then repeated for a slightly different value of B_x (B_z), and the process continues until a sufficient number of data points is collected. In this and other such situations, B_x is applied via the

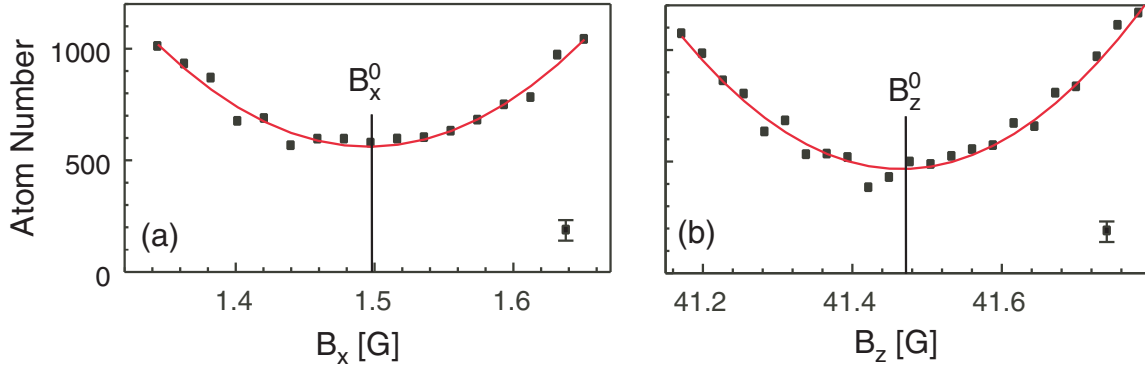


Figure 4.3: Plots of atom number remaining in the microtrap after the trap is held near the origin as a function of (a) B_x to determine B_x^0 (corresponding to the dotted red segment in Fig. 4.2) and (b) B_z to determine B_z^0 (corresponding to the dashed red segment in Fig. 4.2). Quadratic fits used to extract the points of maximum loss are plotted in red. Sample points with error bars showing the uncertainty of our atom-number measurement are plotted for reference.

X -bias coils, while B_z is applied via $Q5$, which, compared to the bias coils, can be switched faster (in ms instead of tens of ms) and can produce much larger fields near the chip, necessary since the required range of microtrap motion is over $100 \mu\text{m}$ in z and only around $20 \mu\text{m}$ in x (a disadvantage to using $Q5$ in this way, as effectively a bias coil, is that it produces non-negligible gradients that we must carefully keep track of).

The results of the field-to-position calibration measurement are shown in Fig. 4.3. B_x^0 is determined with a precision $\delta B_x^0 = 4 \text{ mG}$, and B_z^0 is determined with a precision $\delta B_z^0 = 10 \text{ mG}$. In the spatial region near the origin, where the fields applied in addition to B_x^0 and B_z^0 are small, these uncertainties correspond to a trap position uncertainty of $<50 \text{ nm}$ in the x - z plane, much smaller than the condensate transverse size of 300 nm .

4.1.3 Surface Microscopy

To confirm the accuracy of our ability to infer the position of the atoms from the applied magnetic fields, we measure a line of constant half-life of $\tau = 22 \text{ ms}$ near one of the unused conductors located near the chip center, labeled $Q1$ in Fig. 4.2.

Specifically, we measure, for a given value of B_z , what value of B_x puts the atoms at a distance to the wire that leads to surface-induced loss that reduces their lifetime in the trap to a certain value, in this case 22 ms. We apply a certain B_x far away from the chip, then quickly ramp B_z to the desired value that puts the atoms close to the surface, hold the atoms there for 22 ms, then ramp B_z back and image the atoms. For each value of B_z , we iterate this process many times, varying the applied B_x to determine the value that leaves us with half the atom number we start with.

The resulting 22-ms-half-life contour is shown in Fig. 4.4. Using our calibration of B_x^0 and B_z^0 in combination with a numerical simulation of our applied magnetic field configuration, we are able to map the contour measured in terms of applied transverse magnetic fields onto a set of spatial positions, which we find follows very closely the cross-section of the chip wire at a distance of around $1.5 \mu\text{m}$.

4.2 Lifetime Limits due to Surface-Induced Loss

The ability to know very accurately the position of the atom cloud is used to measure the lifetime τ of atoms in the microtrap as a function of their distance from a surface, for both dielectric and metallic surfaces. The spatial paths used to perform these measurements are shown in Fig. 4.4. The lifetime measurement is performed by bringing the trap to the desired position and holding it there for a set time, then bringing it back away from the surface and observing the atom number. This procedure is repeated for different holding times, and the resulting atom-number-vs.-holding-time graph is fitted with a decaying exponential to extract the lifetime at that trap position. The combined results of all these lifetime measurements, obtained for samples with a temperature $T = 1 \mu\text{K}$ and a trap offset field of $B_0 = 0.57 \text{ G}$, are shown in Fig. 4.5.

As magnetically trapped atoms are brought closer to a dielectric surface, their lifetime remains constant, at around 3.5 s, which is the lifetime limit due to background gas collisions, down to a distance of around $2.5 \mu\text{m}$, independent of cloud temperature in the range between BEC and $3 \mu\text{K}$. For shorter distances, there is a sharp drop in lifetime, the nature of which is examined in more detail in the next




Figure 4.4: Diagram of surface effect measurements. The 22-ms-half-life contour (red) calculated from the magnetic-field-to-position map measured previously very closely follows the outline of the current-free $Q1$ wire cross-section (gray). The inset shows the corresponding trajectory in applied-transverse-magnetic-field space. The measurement paths for lifetime near a dielectric (black) and near a metal (blue) are also shown. The dielectric-approaching path avoids the origin to prevent coupling to the image trap. Approximate B_x (dotted) and B_z (dashed) contour lines are shown in light gray.




Figure 4.5: Trapped atom lifetime as a function of distance from a dielectric (solid squares) and a metal (open circles) surface, for $T=1 \mu\text{K}$ and $B_0=0.57 \text{ G}$. The distance error bars for the measurement are due to uncertainty in conductor thickness and possible miscalibration of applied magnetic fields away from the origin. The dotted line is the calculated lifetime above the metal due to thermal B fields only; the solid line includes the one-body lifetime of 2.5 s. The open triangles are measurements for a pure condensate above the dielectric.

section.

As atoms are brought closer to a metal surface, they experience a distance-dependent lifetime reduction. This is due to coupling of the electron spin of the trapped atoms to magnetic field fluctuations produced by the thermal motion of electrons in the conductor [34]. The AC component of these fluctuations at the frequency that corresponds to the magnetic sublevel splitting due to the trap's offset magnetic field, $\omega_t = g_F \mu_B B_0 / \hbar$, can drive resonant transitions (“spin flips”) between magnetic sublevels in the atoms, which can cause them to end up in untrapped magnetic sublevels, at which point they are lost from the magnetic trap.

In the limit that the metal film thickness t is much smaller than the skin depth δ at the transition frequency, $\delta = \sqrt{2\rho/(\omega_t \mu_0)}$, where ρ is the resistivity of the metal, Johnson noise throughout the entire conductor thickness contributes to the magnetic field fluctuations, and they are frequency independent. For us, $t = 2.15 \pm 20 \mu\text{m}$ and $\omega_t = 400 \text{ kHz}$, so $\delta \approx 120 \mu\text{m} \gg t$, which means the analysis of the spin flip process can be done in the magnetostatic approximation.

The spin flip rate $|F, m\rangle \rightarrow |F, m - 1\rangle$ due to magnetic field fluctuations from a metallic layer of thickness t , temperature T , and resistivity ρ as a function of the distance d from the metal is derived in Ref. [35] to be

$$\Gamma_{Fm} = C_{Fm}^2 \frac{3\mu_B^2 k_B T}{32\pi\epsilon_0^2 \hbar^2 c^4 \rho} [d(1 + d/t)]^{-1}, \quad (4.1)$$

where $C_{Fm}^2 = |\langle F, m - 1 | S_- | F, m \rangle|^2$ and S_- is the electron spin lowering operator. For a wire of finite width w (in this case, $w = 10 \mu\text{m}$), we interpolate this formula so that it gives the above expression for $d \ll w$, while for $d \gg w$, it gives the result for Γ_{Fm} that we derive from Refs. [34, 35] assuming that only thermal currents along the wire contribute to the spin flips. The result is that, compared to the large-distance limit of the thickness-dependent term $(1 + d/t)^{-1} \rightarrow (d/t)^{-1}$, the limit for the width-dependent term is reduced by a factor of 2, so that the interpolated width-dependent term that incorporates both the large-distance and small-distance limits is $(1 + 2d/w)^{-1}$.

By measuring the increase in the resistance of Q2 as the chip heats up during

the experiment we determine $T = 400$ K, and from that we get $\rho = \rho_{\text{Cu}}(400\text{K}) = 2.4 \times 10^{-8} \Omega \text{ m}$. We assume the atoms are lost in a simple 2-step cascade process, $|2, 2\rangle \rightarrow |2, 1\rangle \rightarrow |2, 0\rangle$ and ignore transitions made in the opposite direction, which allows us to replace C_{Fm}^2 by $(C_{22}^{-2} + C_{21}^{-2})^{-1} = (4 + \frac{8}{3})^{-1} = 0.15$.

We can then substitute these results as well as the values of physical constants into Eq. 4.1 to get the following expression for the loss rate due to spin flips:

$$\Gamma_{Fm} = 13 \mu\text{m s}^{-1} \times [d(1 + d/t)(1 + 2d/w)]^{-1}. \quad (4.2)$$

In order to compare this result with our data, we combine the above surface-induced loss rate with the distance-independent one-body loss rate of $\gamma = 0.4 \text{ s}^{-1}$ observed at $d \geq 10 \mu\text{m}$. The plots for the expected lifetime due to this loss both with and without including the one-body lifetime and with no fitting parameters are shown in Fig. 4.5, as a solid and dotted line, respectively. The agreement with experimental results is excellent for distances $d > 3 \mu\text{m}$, while for $d \leq 3 \mu\text{m}$ the interpolated formula predicts a longer lifetime than we observe. Initially, we had thought that this discrepancy was due to increased loss at short distances due to patch electric potentials from ^{87}Rb atoms adsorbed on the surface, as studied in Ref. [51], but a more detailed calculation in Ref. [52] concluded that the presence of a conducting surface was all that was necessary to explain the loss we measured.

Except for the point closest to the metal surface, the trapped atom lifetime τ is independent of atom temperature, and a measurement of τ as a function of distance performed with an offset field $B_0 = 1.5 \text{ G}$, so that the transition frequency is nearly three times larger, yielded a similar distance-dependent loss, consistent with the frequency-independent loss model.

Note that for $d \gg t$ and $d \gg w$ (the limit of a thin, narrow wire), the spin-flip rate derived in Eq. 4.2 drops as d^{-3} and increases linearly with t and w , while, for $d \ll t$ and $d \ll w$ (the limit of an infinite conducting plane), the spin flip rate drops as d^{-1} and is independent of t and w . This suggests that, to avoid spin flips, one should try to work at distances large compared to the wire thickness and width; these distances can still be very small on an absolute scale because the excellent heat

dissipation obtainable on atom chips allows one to achieve very large current densities (see Section 2.4.2), and thus relatively large currents can flow through very thin wires.

4.3 Casimir-Polder Force near Dielectric

While, as mentioned above, the atom lifetime τ near a dielectric is constant and independent of temperature for $d > 2.5 \mu\text{m}$, at shorter distances, where τ experiences a sharp drop, it does depend on temperature, with longer lifetimes for colder clouds. The drop in lifetime is very sharp as a function of distance, however, so the lifetime measurement is not the best way to investigate this dependence. Instead, we use a single short holding time of 15 ms and measure the fraction χ of the initial atom number that survives after that holding time as a function of the distance from the metal surface at which the trap is held. Such a measurement allows us to see the distance dependence of the loss, and we perform the measurement for a BEC and for thermal clouds at $2.1 \mu\text{K}$ and $4.6 \mu\text{K}$ in order to investigate the temperature dependence of the loss process. The results of this measurement are shown in Fig. 4.6.

The colder thermal cloud exhibits loss closer the surface than the hotter one, and a condensate survives even closer in; this result is consistent with surface-induced evaporation [53, 33]. In addition, for all sets of atom parameters all atoms are lost for $d < 1 \mu\text{m}$.

If there were no atom-surface interactions, the trap depth would be given by the difference between the value of the magnetic field at the surface and its value at the trap minimum, which is the offset field, and we would expect that some fraction of the condensate, at least, would survive until much closer to the surface than $1 \mu\text{m}$ (recall that the condensate transverse size in z is $\approx 300 \text{ nm}$). However, there is an attractive Casimir-Polder interaction present [37], which reduces the trap depth and causes the trap to disappear at finite d (see inset in Fig. 4.6).

The Casimir-Polder interaction is the retarded form of the van der Waals interaction. For an atom and a surface, the standard $1/d^3$ van der Waals force is present at small separation, when the fluctuations in dipole moment can travel between the atom and the surface quickly compared to the atom's response time $\tau_{atom} = 1/\omega$,

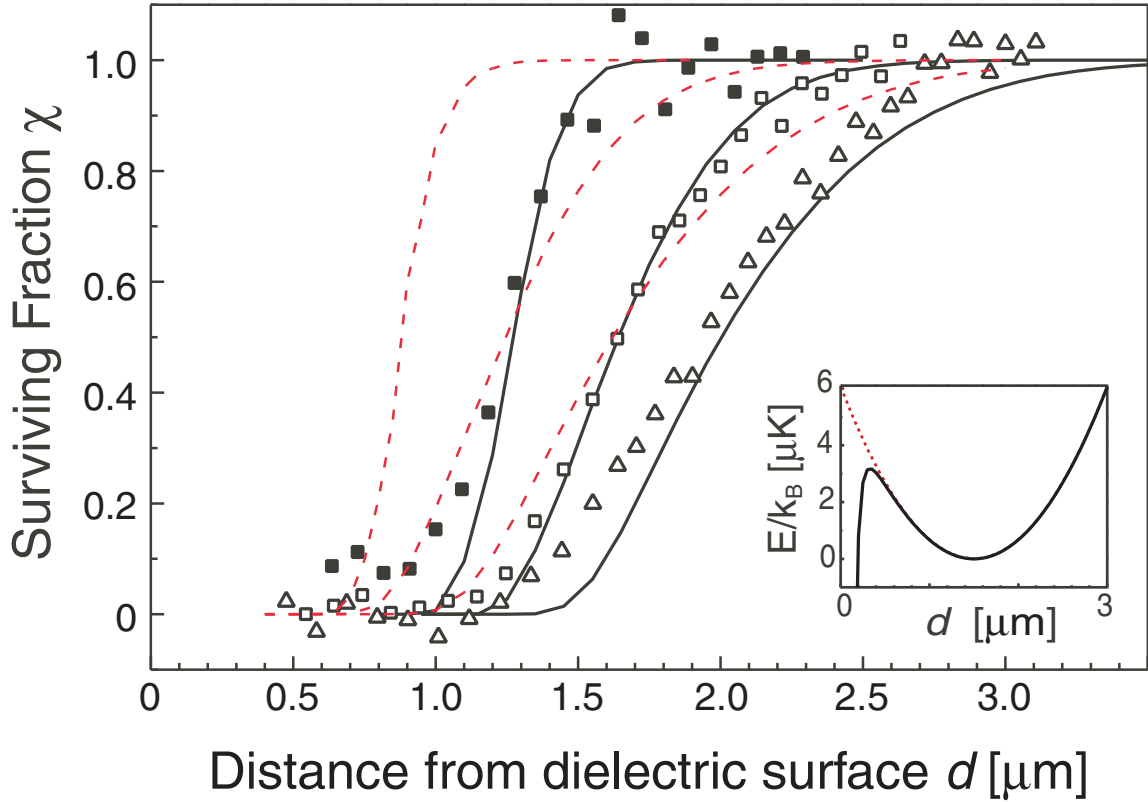


Figure 4.6: Surviving fraction after a 15 ms holding time as a function of distance d to a dielectric surface for a BEC (normalized to 900 atoms; solid squares) and for thermal clouds at $2.1 \mu\text{K}$ (normalized to 8100 atoms; open squares) and $4.6 \mu\text{K}$ (normalized to 11500 atoms; open triangles). The solid lines are loss model calculations using the nominal value for the Casimir-Polder coefficient $C_4 = 8.2 \times 10^{-56} \text{ J m}^4$ while the red dashed lines assume no atom-surface interactions ($C_4 = 0$). The inset shows trapping potentials as a function of d for a trap distance $d = 1.5 \mu\text{m}$ for the nominal C_4 value (solid line) and for $C_4 = 0$ (dotted line).

where ω is the angular frequency for the atom's strongest transition. However, when the separation d becomes so large that the transit time $t = d/c$ between the atom and the surface becomes comparable to τ_{atom} , which happens for $d \geq \lambda/(2\pi) = c/\omega$, the Casimir-Polder potential, given by $V_{CP} = -C_4/d^4$, comes into effect. In the case of ^{87}Rb , $\lambda/(2\pi) \approx 200$ nm, so the Casimir-Polder potential describes the atom-surface interaction in the region observed.

Quantitatively,

$$C_4 = \frac{3\hbar c\alpha}{32\pi^2\epsilon_0}\psi(\epsilon), \quad (4.3)$$

where $\alpha = 5.25 \times 10^{-39}$ F m² is the ^{87}Rb ground-state polarizability and $\psi(\epsilon) = 0.46 \pm 0.05$ is a dimensionless factor computed in Ref. [54] given the Si_3N_4 dielectric constant of $\epsilon = 4.0 \pm 0.8$. Putting in numerical values into the above equation gives a nominal value of $C_4 = 8.2 \pm 0.9$ J m⁴.

We model the measured loss due to surface interactions near a dielectric as a two-step process: an immediate loss of the Boltzmann tail of the atom energy distribution as atoms with energies above the potential barrier lowered by the surface potential are lost from the trap, followed by one-dimensional evaporation that takes place over the $t_0 = 15$ ms during which the atoms are held near the surface in a trap with a transverse frequency of 3.6 kHz. The fraction of atoms remaining after the immediate loss is given by $F = 1 - e^{-\eta}$, where $\eta = U/(k_B T)$ is the ratio between the Casimir-force-limited trap depth $U(d)$ and the atomic thermal energy $k_B T$.

The loss rate for 1D evaporation is

$$\Gamma_l = f(\eta)e^{-\eta}\Gamma_{el}, \quad (4.4)$$

where Γ_{el} is the elastic collision rate and $f(\eta)$ is derived in Ref. [55] as

$$f(\eta) = \frac{1}{2\sqrt{2}}(1 - 1/\eta + 3/(2\eta^2)), \quad (4.5)$$

accurate within 5% for $\eta \geq 4$. We account for atom tunneling through the barrier as a correction to Γ_l due to additional loss from trapped energy levels during the

holding time, computed using the WKB approximation. For the condensate, we assume the the loss is due to collisions with a residual thermal cloud at $T_c/2$. Given elastic collision times $1/\Gamma_{el}$ of 0.2, 0.9, and 1.5 ms for the condensate, 2.1 μK , and 4.6 μK clouds, respectively, we can calculate the remaining atom fraction $\chi_{\text{CP}} = Fe^{-\Gamma_{el}t}$. The results for the calculation of atom loss due to the Casimir-Polder force are plotted in Fig. 4.6, along with a set of loss curves calculated in the absence of any finite-range atom-surface interaction ($C_4 = 0$).

The data plotted in Fig. 4.6 can be interpreted as a measurement of the Casimir-Polder coefficient C_4 . The dominant source of error in this measurement is the uncertainty of the distance calibration. To make sure that there was no systematic offset in the magnetic field calibration, we performed the same measurement on the opposite side of the symmetry point in the x direction. The results, shown in Fig. 4.7, show that there was no systematic shift between the two sides. Note also that the deviation from zero in the atom fraction observed for points closest to the surface for the BEC in Fig. 4.6 is not duplicated on the other side, where the surviving fraction goes to zero. We thus interpret the nonzero data points close to the surface as due to noise and not a real physical phenomenon like quantum reflection [56].

Given that there is not systematic shift between the two sides, the distance uncertainty is dominated by a ± 100 nm uncertainty due to the ± 200 nm uncertainty in conductor thickness t (see Fig. 4.2). An additional 10% scaling error about the symmetry point distance $d_0 = 1.6 \mu\text{m}$ is due to an estimated 10% calibration uncertainty in the applied magnetic field ΔB_x away from the origin.

Furthermore, the C_4 coefficient used in our calculations is accurate for a Si_3N_4 infinite half-space, while our chip actually has a $1 \mu\text{m}$ Si_3N_4 layer on a Si substrate, which forms a dielectric waveguide. The correction due to this geometry is estimated using Ref. [57] to be less than 20%. The model calculation is also affected by $\sim 10\%$ uncertainties in atom temperature and trap vibration frequency. Given all these sources of error, if we use the loss model to extract the value of C_4 from our measurements, we get a 66% confidence range of between $1.2 \times 10^{-56} \text{ J m}^4$ and $41 \times 10^{-56} \text{ J m}^4$, which includes the nominal value of $8.2 \pm 0.9 \text{ J m}^4$.

The good agreement between our measurement and the loss model with no free

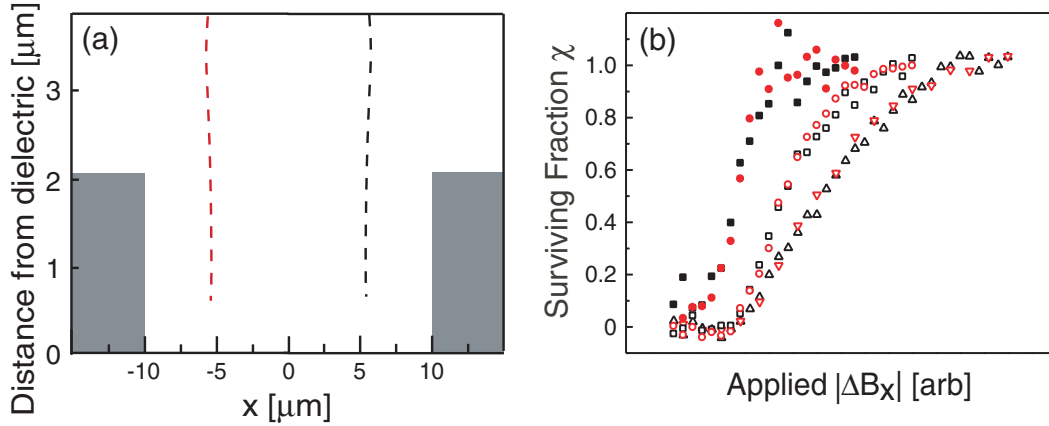


Figure 4.7: Calibration loss measurement. (a) Approximate paths for the original loss measurement near dielectric (black dashed line) and the calibration measurement on the other side of the origin in x (red dashed line). (b) Results for the surviving atom fraction after 15 ms as a function of absolute applied ΔB_x for the two measurements and three sets of atom parameters: the original measurement (same as Fig. 4.6) for a BEC (solid black squares) and for thermal clouds at $2.1 \mu\text{K}$ (open black squares) and $4.6 \mu\text{K}$ (open black triangles) and the calibration measurement (solid red circles, open red circles, and open red triangles, respectively).

parameters suggests that the reduction of trap depth due to the Casimir-Polder potential is what limits the lifetime of atoms at small distances $d \leq 2.5 \mu\text{m}$ from the dielectric chip surface. The model fits the data better at large χ than at small χ , which is probably due to the fact that our simple 1D evaporation model breaks down for $\eta \leq 1$ and also does not take into account the reduction in atom cloud temperature during evaporation, which is more significant at small χ . Furthermore, the interaction between a condensate and a background thermal cloud is more complicated than the model accounts for. Nevertheless, our data exclude $C_4 = 0$, even if we allow for the largest possible systematic error.

This measurement of the Casimir-Polder potential and the surface microscopy measurement described above show the potential of ultracold atoms to be used as extremely sensitive probes of surface properties. Ultracold atoms are very sensitive to magnetic and electric fields, and since the size of a BEC in a tight trap can be much smaller than an optical wavelength, trapped atom probes can achieve excellent spatial resolution. Indeed, condensates have recently been used as sensitive probes

of surface properties such as short-range electrical potentials arising from polarized adsorbates [51], and to achieve a precise measurement of the Casimir-Polder potential out to distances of several μm [58].

4.4 Condensate Splitting

The two-wire trap geometry can be used to magnetically split a condensate in two, which, if it can be done coherently and reproducibly, could have great potential for atom interferometry on a chip. The process by which this splitting is achieved is shown in Fig. 4.8. When the condensate is brought near the surface along the symmetry axis (shown in red and pink in the figure), the transverse quadrupole trap containing the condensate joins with the image trap at the symmetry point to form a transverse hexapole trap, which then splits laterally in the plane of the current-carrying chip wires (in x) into two transverse traps, and, if everything is done symmetrically, each of those traps will have a condensate with half of the atoms. If, however, the condensate is brought near the surface off-center, as shown in dark and light blue in Fig. 4.8, it will miss the symmetry point and the original trap will not join with the image trap, so the atoms will remain in one trap. If the original trap comes close to the symmetry point without quite hitting it, so that the separation from the image trap is small, some atoms can couple to the image trap without the two traps overlapping, which will result in asymmetric splitting.

To illustrate the splitting process, we measure the number of atoms that end up in each of the two lateral traps, also called “output traps,” after the incoming trap passes through or near the symmetry point as a function of the distance by which the trap is misaligned from the center. It should be noted that, in a given shot, we can only look at the atoms in one of the two output traps, either the left or the right. This is due to the fact that these two traps are mirror images of each other, and so exist on opposite sides of the symmetry point. Thus, if we ramp B_x so that the trap on the left of center moves away from the chip, the trap on the right will move into the chip and the atoms in it will be lost, and vice versa.

The results obtained for such a measurement, done with an initial condensate of




Figure 4.8: Contour map of B_x (dashed) and B_z (solid) near the chip (same as Fig. 4.2) with paths for condensate splitting near the symmetry point as B_z is changed to bring the trap near the surface. If the trap starts out centered along x (solid red circle) and comes in along the symmetry axis (solid red arrow), then the image trap will join with it at the symmetry point (pink dashed arrow) to form a transverse hexapole trap which will then separate into two traps laterally (solid pink arrows), each of which will have half the initial atom number (solid pink circles). If the trap starts out off-center (solid dark blue circle), it will come in off the symmetry axis (solid dark blue arrow) while the image trap passes on the other side of the symmetry point (dashed light blue arrow) and will never join with it. Thus the initial trap ends up with all the atoms (solid dark blue circle) while the image trap ends up with no atoms (open light blue circle).

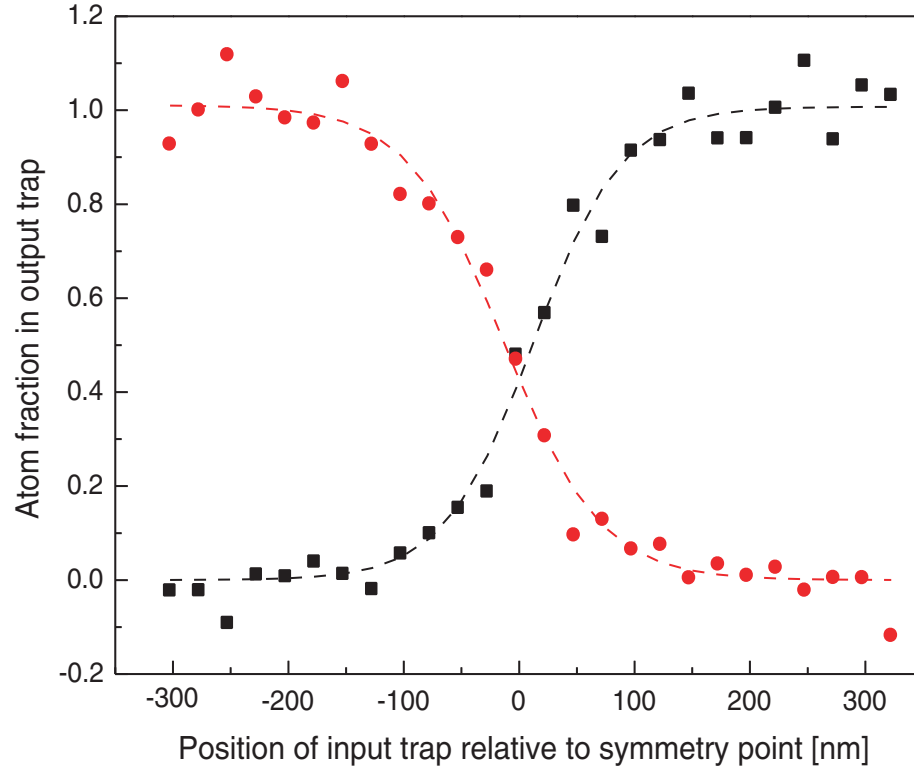


Figure 4.9: Fraction of atoms from the initial microtrap that end up in the left output trap (red circles) and the right output trap (black squares) as a function of the initial trap’s x alignment to the symmetry point. The two dashed curves show smooth-step fits to the atom fraction measurements.

900 atoms and a splitting time of 10 ms, are shown in Fig. 4.9. The splitting is extremely sensitive to alignment – if the incoming microtrap is misaligned to the left (right) by as little as 150 nm, all the atoms end up in the left (right) trap, and the trap has to be centered to within about 20 nm to get equal splitting.

The timescale for the splitting is an important experimental parameter. If the splitting is done too quickly, the condensate will be excited, while doing it too slowly causes atoms to be lost in the vicinity of the symmetry point (this loss was exploited in the calibration measurement plotted in Fig. 4.3). The splitting time of 10 ms was chosen so that the loss was small, about 10% at the point of equal splitting (the graph shows that, for equal splitting, each trap ends up with about 45% of the initial total

atom number).

This splitting time, however, is not sufficiently long to avoid creating excitations in the condensate that can be seen in the absorption images. This is due to the fact that, when the two transverse quadrupole traps join to form a hexapole trap at the origin, the vibration frequency temporarily becomes very low, which, in combination with the anharmonicity of the hexapole trap, makes it very easy to excite the condensate. One way to address the excitation problem is to perform the splitting farther away from the surface, so that the loss is much slower, which should allow much slower splitting times. A tighter initial trap would also help, by making it harder for atoms to couple to the image trap, which is the primary mechanism for the loss.

In addition to loss during splitting, another limitation is the effect, after splitting, of the Casimir-Polder force due to the proximity of the surface. The two output traps are located more-or-less in the plane of the chip wires, although they are moved about $0.5 \mu\text{m}$ farther away from the surface by the gradient from the small but macroscopic $Q5$ coil. Even with the presence of this additional gradient, the output traps are located only about $1.5 \mu\text{m}$ from the surface, which, as shown in Fig. 4.6, is still within the range of the Casimir-Polder potential (though, for a condensate, it is at the outer limit of that range). Thus, the proximity of the surface could have a disruptive effect on any interference-type experiments performed using the double-well potential that creates the two lateral traps.

While loss due to the proximity of the dielectric chip surface is a limitation of this splitting technique, the splitting can be done far enough from the chip wires themselves to eliminate any Johnson-noise-induced loss. In the splitting example presented here, the nearest edges of the current-carrying $Q2$ wires are about $35 \mu\text{m}$ away from the symmetry point, completely outside the range of the Johnson-noise-induced loss (recall that the closer wires whose cross-section can be seen in Fig. 4.8 carry no current and are not used for trapping the atoms).

Chapter 5

Integrating an Optical Cavity with the Chip

5.1 Motivation for Using Cavity

It was first pointed out by Purcell 60 years ago, in regard to radio frequencies, that the spontaneous emission of radiation by an atom is altered when the atom is placed inside a resonator or cavity [42]. In particular, the emission rate Γ_c into a cavity of quality factor Q at frequencies resonant with it is enhanced compared to the free-space emission rate Γ_{fs} by a factor, known as the “single-atom cooperativity,” given by

$$\eta_c = \frac{\Gamma_c}{\Gamma_{fs}} \frac{3Q\lambda^3}{4\pi^2 V}, \quad (5.1)$$

where V is the volume of the resonator mode under consideration. This can be interpreted as being due either to the increased density of electromagnetic modes inside the resonator or to the superradiance-like interaction of the oscillating dipole with its image oscillators.

This enhancement is not limited to radio frequencies, and has been observed in both the microwave [59] and optical [60] domains. In the optical case, for a symmetrical confocal resonator of length L with a Gaussian mode of waist w_0 , the mode

volume is given by $V = \pi L w_0^2 / 2$ [59], and the cavity finesse \mathcal{F} , which is the ratio of the cavity free-spectral range, $\nu_{FSR} = c/(2L)$, to the cavity linewidth, is given by π/q^2 , where $q^2 \ll 1$ is the single-mirror transmission (and mirror losses are negligibly small). Then, Eq. 5.1 simplifies to

$$\eta_c = \frac{12\mathcal{F}}{\pi(w_0k)^2}, \quad (5.2)$$

where $k = 2\pi/\lambda$. Note that, in the absence of the cavity (if, for example, the cavity mirrors were replaced with collection lenses), the emission into the cavity volume would be given by $6/(w_0k)^2$, so the presence of the cavity mirrors enhances emission by a factor of $2\mathcal{F}/\pi$. η_c describes the combined emission into both directions of the cavity; for our experiments, we are often interested in the emission into only one direction of the cavity, which, for a symmetric cavity, is given by

$$\eta_1 = \frac{\eta_c}{2} = \frac{6\mathcal{F}}{\pi(w_0k)^2}. \quad (5.3)$$

From the point of view of the atoms, this enhancement of the resonant scattering rate inside a cavity looks like an increase in the cavity's effective solid angle, while from the point of view of light in the cavity, it looks like an increased optical depth for the atoms, since the light interacts with each atom many times as it circulates around the cavity.

In the case of more than one atom in a cavity, the enhancement factor for forward scattering into the cavity is increased by a factor equal to the atom number N , for a total scattering rate of $N^2\eta_c$, due to the constructive interference of photons scattered by many atoms into the same cavity mode.

This increase in the interaction between atoms and intracavity light can be exploited to achieve strong coupling between atoms and light. Strong coupling is achieved when the Rabi flopping rate between atomic excitations and intracavity electric field excitations (photons), given by $g_c = \sqrt{N\eta_c\Gamma\kappa}$, where Γ is the natural linewidth of the atomic transition and κ is the linewidth of the cavity, becomes greater than either Γ or κ . In the regime where $\Gamma \approx \kappa$, this reduces to $N\eta_c > 1$. Generally, high-finesse cavities are required to achieve the strong coupling limit for a single atom,

but having many atoms in the cavity can allow one to achieve strong coupling with relatively low-finesse cavities.

Integrating cavities with chip-based magnetic microtraps is promising for several reasons. The tight confinement achieved in microtraps produces narrow, cigar-shaped atom clouds that are spatially very well suited to being overlapped with narrow cavity modes, which should permit the loading of large numbers of atoms into the volume of a cavity mode with a small waist. Thus, microtraps should facilitate cavity QED experiments with many atoms and quantum communication experiments using atomic ensembles that rely on large optical depths [61].

At the same time, there are many proposed atom chip experiments, such as the implementation of a Tonks-Girardeau gas [22, 23, 62] or an atomic Fabry-Perot interferometer [21] in a magnetic trap, that may greatly benefit from measuring atom statistics and correlations at the single-atom level, which has been achieved for a free-space trap via the use of a cavity [63]. In addition, the preparation and detection of single atoms, which has been achieved in free-space cavity experiments without magnetic traps [64, 65], if implemented in a microtrap setup would constitute an important step toward quantum information processing with neutral atoms, which could take advantage of the tight, complex, precisely controlled, and scalable magnetic traps available on microchips [66]. And, in general, one of the great advantages of atom chip experiments is that they offer tight confinements and control of atoms at small length scales; experiments that push the technological limits of atom chips in these respects may have to work with small atom numbers, which means it is important to be able to detect small atom numbers in magnetic microtraps near surfaces [39, 40]. Since the presence of a cavity can magnify the optical depth, and thus the effective atom number, cavities seem ideally suited to these small-atom-number atom chip applications.

5.2 Cavity on a Chip

5.2.1 Modified Chip Design

The optical cavity is integrated into a setup that uses a new chip, modified from the version used for the surface effects experiments and incorporating new ideas which derived from that work. Measurements done with our original chip show that a two-wire trap could be used to split a Bose-Einstein condensate far away from the current-carrying chip wires, so that Johnson noise from the conductors was not a problem (see Figs. 4.8 and 4.9). However, this splitting occurs in the plane of the trapping wires, close to the chip surface, and thus it is susceptible to Casimir-Polder forces from the dielectric chip surface. To solve this problem, we modified our new chip by etching away the dielectric surface around the metal chip wires after they are deposited, so that the chip wires end up being raised above the rest of the chip surface on ridges of dielectric. The plane of the chip wire centers is thus moved farther from the chip surface.

The removal of dielectric is accomplished by using a slightly different chip substrate. Whereas our first chip used a substrate with 1 μm of Si_3N_4 insulator on top of Si, the new chip uses a substrate with a 3- μm Si layer on top of a 0.2- μm SiO_2 insulating layer on top of bulk Si. After the metal wires are deposited, they are used as a mask when the Si is etched down to the SiO_2 layer around them, which leaves the wires raised on 3- μm ridges of Si, so that the plane of the chip wires is far enough from the surface to be immune to the disruptive effects from the Casimir-Polder force.

The layout of the chip wires is also altered in the new design, which is shown in Fig. 5.1. The second-largest chip coil, $Q3$, is eliminated, while $Q2$ is modified so that it has a larger loop on the end, for possible use as an additional trap. Furthermore, in order to accommodate the cavity mode near the chip, the Ioffe-Pritchard ribbons are raised to pass 500 μm away from the chip, compared to 155 μm away for the old chip. Atoms are once again confined axially by a combination of the gradient from a Ioffe-Pritchard ribbon, I_{+1} and the MOT gradient; however, in order to couple to the cavity, atoms are trapped about 200 μm away from the chip surface, and because of this larger distance and also due to its lower resistance and therefore lower heating,

we use $Q4$, not $Q2$, to provide the transverse confinement.

5.2.2 Physical Implementation of Cavity

The cavity is formed by two identical high-reflectivity, low-loss mirrors mounted on opposite sides of the chip. The mirrors are curved, with radius of curvature $R_c = 2.5$ cm, and the cavity is designed to be slightly longer than confocal so that different transverse modes can be resolved. The mirrors are manufactured by Research Electro-Optics and have maximum reflectivity/minimum transmission at a wavelength of 850 nm. Single-mirror transmission at 780 nm is measured to be about $5 \pm 2 \times 10^{-4}$, with loss expected to be much smaller than that.

One of the mirrors is glued with Torr-Seal vacuum epoxy to a long, thin-walled, cylindrical ceramic piezoelectric transducer (PZT) to allow the cavity length to be tuned with an applied voltage range of several hundred volts. The other mirror is glued to a short PZT, which is not used in the experiment because we were not able to solder leads to it. The PZTs with mirrors attached are glued to stainless steel mounts which are then screwed to the chip carrier with vented screws in such a way that the cavity waist lies about $200 \mu\text{m}$ above the chip surface. Fig. 5.2 shows a schematic of how the cavity was mounted on the chip carrier. Note that since one of the mirrors is very close to the chip while the other is separated from it by the length of the long PZT, the center of the cavity is located close to the chip's edge.

The cavity is aligned by hand, through an iterative procedure of coupling light into the cavity and seeing how good the alignment is, slightly loosening one of the screws that hold the mirror mounts in place, adjusting that mount to improve the alignment, tightening the screws, checking the alignment again, etc. The cavity mode threads the space between the Ioffe-Pritchard ribbons and the chip surface, which allows us to use scatter light from either the chip's surface or the Ioffe-Pritchard ribbons to align the cavity in the vertical (z) direction. A razor blade attached to a micrometer stage is used during cavity alignment to block part of the cavity mode in order to characterize the spatial profile of the cavity mode and determine from it the distance of the mode from the chip surface, and also to see the horizontal (x) centering of

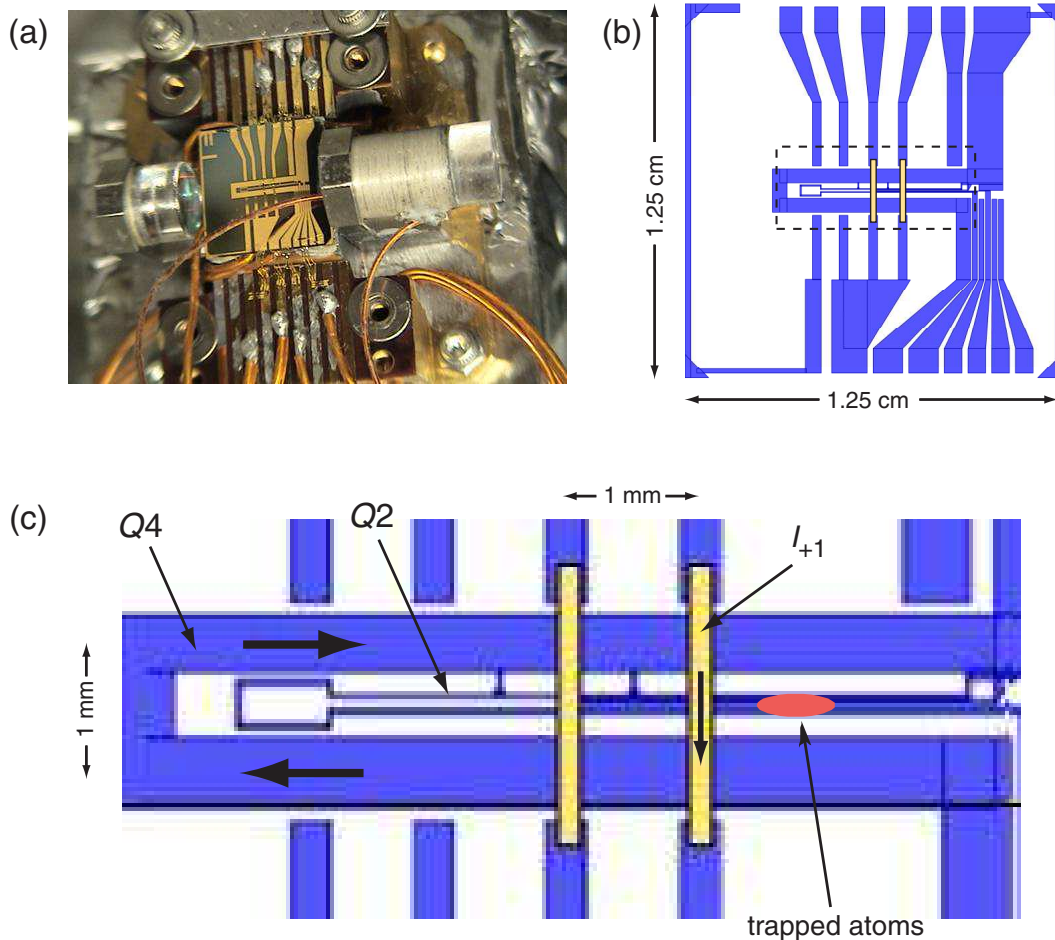


Figure 5.1: (a) Photograph of the new atom chip and integrated optical cavity mounted on the chip carrier. (b) The mask used to create the new chip's wires; conductors are shown in blue. (c) Schematic of the central section of the new chip, marked by the dashed box in (b), showing the trapping wires and the location of the chip microtrap. The conductors on the chip's surface, which include the quadrupole wires and the pads for bonding the Ioffe-Pritchard ribbons are in blue; the Ioffe-Pritchard ribbons are in yellow. The quadrupole loop and Ioffe-Pritchard ribbon primarily used in the cavity experiments are labeled $Q4$ and I_{+1} , respectively, and the directions of current flow through them are indicated.

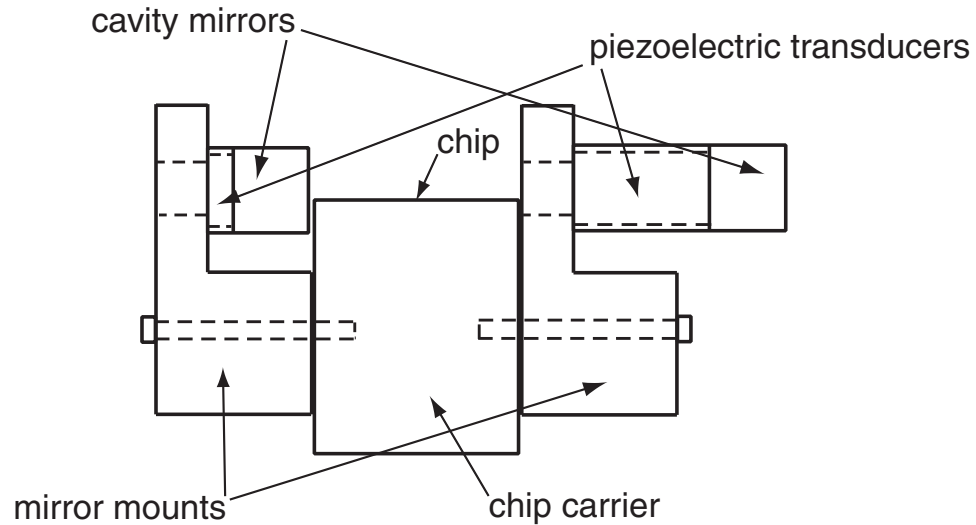


Figure 5.2: Schematic of the way the cavity is mounted on the chip carrier. The mirrors are glued to ceramic cylindrical-shell PZTs, which are glued to steel mounts that are then screwed to the chip carrier. Details of the chip carrier are omitted for simplicity.

the cavity mode. The near-confocal nature of the cavity also makes it easier to align, since, for a confocal cavity, small amounts of misalignment, both in angle and position, only displace the mode along the mirror surfaces without degrading it.

The method for coupling light into and out of the cavity is shown in Fig. 5.3. The output of a collimated fiber is passed through a two-lens telescope to expand the beam for optimal mode-matching, then, after a beam-steering mirror, is focused into the cavity's TEM_{00} mode by a lens mounted on a 3D translation stage just outside the vacuum chamber. The output of the cavity is coupled into another collimated fiber in an exact mirror-image of the coupling-in setup, except the fiber has a band-pass filter before its input, and the output of that fiber goes to either an avalanche photodiode or to a single-photon counting module. We also use a large-area unamplified photodiode right after the cavity out-coupling lens to quickly measure large, slow cavity output signals.

By adjusting the 3D translation stage that holds the final focusing mirror, the mode-matching can be optimized so that 95% of the cavity transmission comes out in the TEM_{00} mode.




Figure 5.3: Diagram of the method for coupling laser light into and out of the cavity. A beam comes out of a collimated fiber, is expanded by a telescope to optimize mode-matching, and focused into the cavity by a lens on a 3D translation stage. The output coupling is the mirror image of the input coupling.

The output of the cavity is mode-matched and coupled into a single-mode fiber, which acts as an excellent spatial filter for background light, since only light that comes out of the cavity mode is well-coupled to the fiber mode. A bandpass filter located right in front of the fiber input helps exclude room light. We achieve a coupling efficiency of the cavity output into the single-mode fiber of about 70%.

5.3 Cavity Parameters

While it is advantageous to have the cavity mode as close to the chip surface as possible to maximize the microtrap gradient for atoms within the cavity mode, the distance is limited by the need to have the mode far enough away so that none of it is cut off by the chip's surface. To compute how close to the surface we can get, we can integrate the intensity over the cross-section of the TEM₀₀ mode minus the portion cut off by the chip surface. The normalized intensity profile of a Gaussian mode is given by

$$I = \frac{2e^{-2r^2/w^2}}{\pi w^2} = \frac{2e^{-2(x^2+y^2)/w^2}}{\pi w^2}, \quad (5.4)$$

where w is the mode spot size, so the fractional loss f_l from having the mode cut off by a plane located a distance d from the center is given by:

$$f_l = 1 - \int_d^\infty \int_{-\infty}^\infty \frac{2e^{-2(x^2+y^2)/w^2}}{\pi w^2} dx dy. \quad (5.5)$$

While our cavity is not exactly confocal, it is close enough to confocal to be treated as such for the purposes of this calculation. For a symmetric confocal cavity of length $L = 2.5$ cm, the waist size w_0 is given by

$$w_0 = \sqrt{\frac{L\lambda}{2\pi}} = 57 \mu\text{m}, \quad (5.6)$$

and the Rayleigh range by $z_R = L/2 = 1.25$ cm. Thus, the mode size at the mirror, which is 1.25 cm from the waist, is given by $w = w_0\sqrt{2} = 81 \mu\text{m}$. Since one of the mirrors is very close to the edge of the chip, we can use the value of $w = 81 \mu\text{m}$

to calculate how close to the surface the cavity mode can be. Assuming we want the loss fraction f_l to be at least an order of magnitude smaller than the single-mirror transmission of 5×10^{-4} , so that the cavity is still transmission-dominated, the mode has to be located at least $160 \mu\text{m}$ from the chip. Adding in a healthy margin of error due to the difficulty of determining the distance between the mode and the surface to better than a few tens of μm gives the $200 \mu\text{m}$ distance used.

Once the cavity is inside the vacuum chamber, its length can be tuned by applying a high voltage through a vacuum feedthrough connected across the long PZT that holds one of the cavity mirrors. By applying a voltage range of 0-415 V, the cavity can be scanned nearly linearly over slightly more than two free spectral ranges.

The cavity free-spectral range (ν_{FSR}), which is the frequency difference between successive cavity peaks, is determined by using the MOT master and repumper lasers together. Beams derived from both lasers are coupled into the cavity while it is being scanned by applying a high-voltage triangle wave to the PZT and the lasers are tuned until the transmission peaks from the MOT overlap those of the repumper. Note that the peaks occur when the cavity length L corresponds to an integer number of half-wavelengths:

$$L = \frac{n\lambda}{2} = \frac{nc}{2\nu} \quad (5.7)$$

where n is an integer, c is the speed of light, and ν is the laser frequency. When the same cavity length produces peaks for two lasers with different frequencies ν_1 and ν_2 , we get

$$\begin{aligned} \nu_1 &= \frac{n_1 c}{2L} \\ \nu_2 &= \frac{n_2 c}{2L} \\ \nu_2 - \nu_1 &= (n_2 - n_1) \frac{c}{2L} \quad , \end{aligned} \quad (5.8)$$

which means that the two laser frequencies differ by an integral multiple of $\nu_{FSR} = c/(2L)$. By calculating back from the VCO frequencies for the MOT and repumper

beat-note locks, which give $1/8192$ of the beat-note frequency, when the MOT and repumper peaks are overlapped on the cavity transmission spectrum, we obtain a frequency difference between the two lasers of 5632 ± 3 MHz. Since we know that L for our near-confocal cavity is slightly larger than 2.5 cm, ν_{FSR} must be slightly less than 6 GHz, so our measured frequency difference corresponds to one free spectral range. Given ν_{FSR} of 5632 ± 3 MHz, we can compute the length of the cavity: $L = c/(2FSR) = 26.62 \pm 0.02$ mm.

We can use the same two-laser method to determine the transverse mode spacing, by overlapping the TEM_{00} mode of the repumper cavity transmission with the TEM_{02} mode of the MOT cavity transmission. The frequency difference should then differ from ν_{FSR} by one transverse mode spacing. From this measurement, we get a transverse mode spacing of 229 ± 3 MHz. The transverse mode spacing can also be calculated directly from the cavity length, using formulas found in [67]; this calculation yields 232 ± 3 MHz, consistent with our measurement.

Given the length of the cavity, we can also determine more precisely the waist size w_0 and Rayleigh range z_R of the TEM_{00} mode. The deviation from confocality is very small, so the values of these parameters are very close to their values for the confocal case. Using the cavity length and results found in [67], we get $w_0 = 56 \mu\text{m}$ and $z_R = 1.26$ cm.

The finesse \mathcal{F} of the cavity can be calculated from the free-spectral range and the single-mirror transmission, $q^2 = 5 \pm 2 \times 10^{-4}$, if we assume cavity losses are negligible, as they should be for these high-quality mirrors: $\mathcal{F} = \pi/q^2 = 7000 \pm 3000$. However, the large uncertainty in q^2 leads to a large uncertainty in \mathcal{F} . Thus, it makes sense to try to measure the finesse instead, which, given our very precise knowledge of ν_{FSR} , can be done by measuring the cavity linewidth.

The linewidth of the cavity transmission, γ_c is measured by first calibrating the PZT voltage scan in terms of frequency by using either the transverse mode spacing or 14.6 MHz sidebands modulated onto the laser by an electro-optical modulator (EOM). This calibration allows us to convert between the time it takes for the cavity to scan over a feature and the frequency width of that feature. From this, we determine the width of the cavity resonance.

The cavity linewidth measurement yields values for $\gamma_c/(2\pi)$ in the range of 0.65-1.2 MHz, depending on the laser used. This is due to the fact that the width of the cavity transmission peak is a convolution of the laser and cavity linewidths. The value of the convoluted linewidth depends on the shape of the laser and cavity resonances; the cavity resonance should be a Lorentzian, but the laser line is some combination of a Lorentzian and a Gaussian. When convoluted, Lorentzian linewidths add linearly, while Gaussian linewidths add in quadrature.

Linewidths for diode grating lasers depend on both the diode used and the grating used, as well as the alignment and stability of the external cavity. It is by using our narrowest laser (the MOT master) and optimizing the feedback by minimizing the lasing threshold that we are able to measure the narrowest cavity transmission resonance linewidth of $2\pi \times 650 \pm 50$ kHz. Given our result from the mirror transmission measurement, it is reasonable to think that this measured linewidth is dominated by the cavity's intrinsic linewidth, which gives us $\kappa = 2\pi \times 650$ kHz and $\mathcal{F} = 2\pi\nu_{FSR}/\kappa = 8600 \pm 700$.

Once we know \mathcal{F} , we can calculate the single-atom cooperativity η_c using Eq. 5.2. An important detail here is that the magnetic microtrap is not located at the cavity waist (which is near the edge of the chip), but is displaced from it by about 5 mm. The mode size at the position of the atoms is then given by

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (5.9)$$

which gives $w = 60 \mu\text{m}$ for $z = 5$ mm. Combined with the values of \mathcal{F} and $k = 2\pi/780$ nm, this yields $\eta_c = 0.14 \pm 0.01$, which means the single-direction cooperativity is $\eta_1 = 0.07$.

5.4 Cavity Length Stabilization

The length, and thus the resonance frequency, of the cavity is stabilized to a locked laser using a Pound-Drever-Hall scheme [68, 69]. A simplified layout of how the lock is set up is shown in Fig. 5.4. An electro-optical modulator (EOM) adds 14.6 MHz

sidebands to the laser beam. If the carrier is nearly resonant with the cavity, the reflected portion at that sideband frequency experiences a large phase shift, while the carrier, not resonant with the cavity, is unshifted. A low-noise avalanche photodiode monitors the reflected beam coming back from the cavity and registers the beat note between the carrier and the sidebands, which is then demodulated by the sideband modulation frequency to give a DC signal proportional to the relative phase between the carrier and the sidebands. Near cavity resonance, this phase shift is proportional to the detuning between the laser carrier frequency and the cavity resonance, and it is used as the error signal for an electronic feedback loop that produces a voltage applied to the cavity PZT to keep the cavity resonant with the laser. This technique can be used to lock the cavity either to the laser carrier frequency, or to one of the sidebands; in both cases, the frequency that is far from cavity resonance serves as a stable phase reference against which the near-resonant frequency is beated in order to extract its phase shift. Locking to a sideband instead of the carrier allows the cavity to be stabilized with much lower intracavity power, since the locking signal depends only on the product of the powers in the two frequencies and thus does not depend on which of the frequencies is resonant with the cavity. Typical optical powers used in the locking beam are about $1 \mu\text{W}$ in the carrier and about 1 nW in the sidebands.

When optimized, the cavity lock achieves a stability of $140 \text{ kHz}/\sqrt{10\text{kHz}}$, measured by monitoring cavity transmission and converting from intensity noise back to frequency noise. Most of our measurements, however, cannot be performed in the presence of the cavity light, since its transmission through the cavity can obscure our experimental signals. An offset can be applied to the cavity tuning signal when the cavity locking light is turned off in order to temporarily stabilize the cavity length. By doing this, we can achieve $<400 \text{ kHz}$ frequency jitter in the several ms after the locking light is turned off, which is when we perform our measurements.

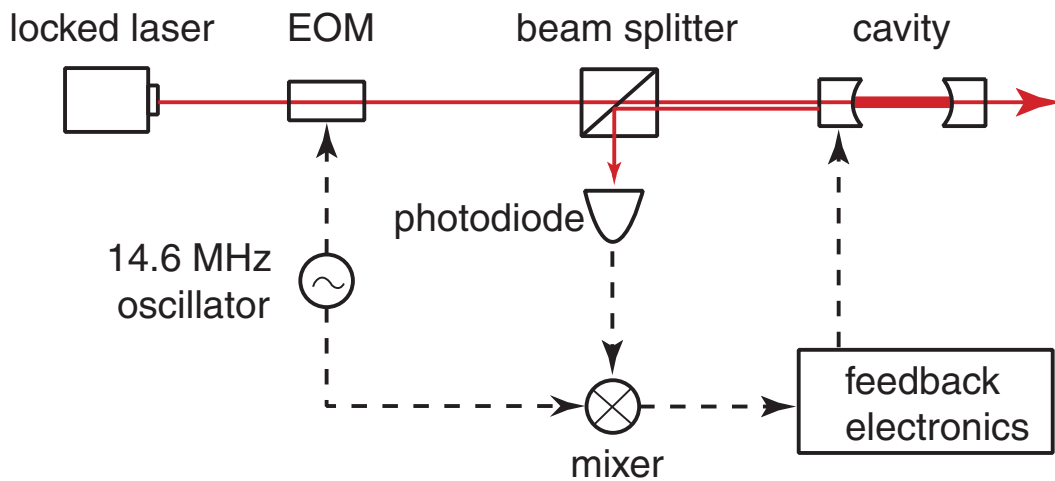


Figure 5.4: Simplified layout for locking the cavity to a laser using the Pound-Drever-Hall method. The EOM adds sidebands to the laser beam, while the photodiode monitors the beat note between the carrier and sideband beams reflected from the cavity. The beat note is demodulated to produce an error signal proportional to the detuning between the laser and the cavity resonance, which is fed back to the cavity PZT to keep the cavity locked to the laser. Solid red lines are optical paths and dashed black lines are electronic signal paths. Fibers, coupling optics, etc., are omitted for simplicity.

Chapter 6

Cavity-Aided Atom Detection

6.1 General Imaging Analysis

There are, in general, two methods of detecting atoms on resonance, fluorescence and absorption. In fluorescence, the atom scatters incoming laser light from a pump beam, a fraction of which is captured by the imaging optics. In absorption, on the other hand, it is the incoming laser beam itself that is collected and detected by imaging optics, and when atoms scatter light out of it, that shows up as photons missing from the beam; the signal in absorption is a negative one, with less light collected in the presence of atoms than without them.

It is useful to compare the fundamental limits of these two detection methods; in particular, let us for the moment ignore the issue of spatial resolution and analyze only atom-number resolution. For this, let us consider a sample of a atoms, in a situation where each of them scatters an average of m photons, with an imaging system that detects a fraction α of all the photons scattered by the atoms (here we ignore saturation effects, so that doubling the laser power and observing for half the time produces the same observed signal, which allows the analysis to be independent of time). For fluorescence, let us assume that the background, i.e. the signal detected in the absence of atoms, can be neglected, since it is a technical and not a fundamental limitation. The detected signal S_f will be given by $S_f = a\alpha m$, with a shot noise given by $N_f = \sqrt{a\alpha m}$, while the signal that corresponds to one atom, $s_f^{(1)}$ will be

$s_f^{(1)} = S_f/a = \alpha m$. The shot-noise-limited atom number uncertainty for fluorescence imaging will then be given by

$$\Delta_f a = \frac{N_f}{s_f^{(1)}} = \sqrt{\frac{a}{\alpha m}}. \quad (6.1)$$

Note that since the signal per atom is constant while the shot noise of the measurement scales as \sqrt{a} , the atom number uncertainty also grows with \sqrt{a} . Thus, in a fluorescence measurement with a negligible background, the atom-number resolution gets worse as the square root of the atom number.

To try to analyze the equivalent absorption measurement, with the same number m of absorbed and rescattered photons and an absorption beam well-matched to the collection optics, let us consider the imaging scheme in Fig. 6.1, which uses a diffraction-limited imaging system through which is coupled a probe beam with the atoms located at its waist. Note that, in this situation, $\alpha = 3/(w_0 k)^2 = \sigma/(2\pi w_0^2)$, where $\sigma = 3\lambda^2/(2\pi)$ is the maximal atomic scattering cross-section, λ is the wavelength of the atomic transition used, $k = 2\pi/\lambda$ is the corresponding wave number, and w_0 is the beam waist. The fraction of photons absorbed by an atom located at the waist is given by $I_0\sigma/P$, where I_0 is the peak probe beam intensity, P is the probe beam power, and we have assumed $\sigma \ll w_0$. For a Gaussian beam profile, defined by $I(r) = I_0 e^{-2r^2/w_0^2}$, $I_0 = 2P/(\pi w_0^2)$, so the fraction of photons absorbed by an atom is given by $2\sigma/(\pi w_0^2) = 12/(w_0 k)^2 = 4\alpha$. If we consider the case where the total absorption is small, so that the absorption beam is not significantly depleted and the absorption signal can be treated as being linear in atom number, then the requirement that each atom absorb and rescatter m photons means that the total number of incoming probe beam photons must be $m/(4\alpha)$. The corresponding shot noise is given by $N_a = \sqrt{m/(4\alpha)}$, while the signal that corresponds to one atom is just $s_a^{(1)} = m$. Therefore, the shot-noise-limited atom number uncertainty for absorption imaging will be given by

$$\Delta_a a = \frac{N_a}{s_a^{(1)}} = \sqrt{\frac{1}{4\alpha m}}. \quad (6.2)$$




Figure 6.1: Diagram of detection of atoms via fluorescence and absorption. Atoms are either pumped by a resonant pump beam from which they scatter photons into the imaging system (fluorescence), or they block a resonant probe beam coupled through the imaging system and focused at their position and scatter photons out of it into free space (absorption).

This is only a factor of 2 better than the uncertainty for fluorescence imaging for $a = 1$, but, unlike the case for fluorescence, the atom number uncertainty for absorption does not depend on a , as long as the low-absorption limit holds.

The above analysis is performed for free-space imaging, but it is not hard to extend it to the case of cavity-aided detection. The effects of a cavity can be modeled by taking the lenses in Fig. 6.1 and replacing them with high-reflectivity mirrors. In the presence of such a cavity, both the emission rate into the cavity, which governs fluorescence imaging, and the effective optical depth of an atom in the cavity, which governs absorption imaging, are enhanced by a factor of $2\mathcal{F}/\pi$, where \mathcal{F} is the cavity finesse. Another way to think of it is that the collection efficiency of the optics, α , is effectively increased by this factor, which affects both fluorescence and absorption. If, instead of a confocal cavity, where the light passes through the atoms twice on every round trip, we instead considered a ring cavity, where the light passes through the atoms only once on every round trip, this enhancement factor due to the presence of the cavity would be halved, to \mathcal{F}/π .

6.2 Atom Preparation

To perform atom detection measurements, we initially load 10^5 atoms into a Ioffe-Pritchard microtrap located $200 \mu\text{m}$ from the surface and displaced by about $150 \mu\text{m}$ to the side of the cavity axis (in x) in order to prevent heating and loss produced by the cavity-length stabilization light. To minimize the effect of the cavity locking light on the atoms while keeping the cavity locked to atomic resonance, we lock the cavity using light one cavity free spectral range, $\nu_{FSR} = 5.6 \text{ GHz}$, detuned to the blue of atomic resonance. This beam is derived from the repumper, which is ordinarily tuned 6.6 GHz to the blue of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition used for measurement (see Fig. 2.2), but is tuned by 1.0 GHz in a few ms when the MOT is turned off and the atoms are loaded into the magnetic trap. The power in the sideband, which is resonant with the cavity, is only about 1 nW , but that is enhanced by a factor of $\mathcal{F}/\pi = 2700$ inside the cavity. The enhanced power of a few μW is sufficiently large that, even at a detuning of 5.6 GHz , the cavity locking light has a measurable destructive effect on the atoms over hundreds of ms.

The microtrap located $200 \mu\text{m}$ away from the surface has a transverse gradient of 230 G/cm , which corresponds to transverse vibration frequencies around 300 Hz for typical offset fields, and an axial vibration frequency of 50 Hz . The transverse confinement in the trap is generated by the $Q4$ chip coil (see Fig. 5.1). At a distance of $200 \mu\text{m}$ from the surface, the $Q2$ chip coil can produce gradients about twice those of $Q4$ for typical currents, but we do not use it for these experiments due to the large heat load it puts out when used.

When current is run through the chip's wires, it heats up the chip carrier, which causes it to expand. While the water cooling of the top piece stabilizes the chip carrier's temperature on long time scales, it is too slow to prevent this heating and expansion during an experimental cycle. Since the cavity mirrors are mounted on the chip carrier, its expansion causes the cavity length to increase, which means that the cavity lock circuit needs to apply a voltage to the cavity PZT to counter this effect and keep the cavity length constant. Because of its smaller cross-section, the resistance of $Q2$ is about an order of magnitude higher than that of $Q4$, so it puts

out significantly more heat for similar currents, which means that the cavity drifts faster and by a larger total amount. It turns out to be much more difficult to keep the cavity locked while there is current flowing in $Q2$ than it is for a current in $Q4$, so we end up only using $Q4$ to create the microtrap (the Ioffe-Pritchard ribbon used to create the axial confinement, I_{+1} , has a much lower resistance than $Q4$ and so does not produce much heat).

Once the atoms are loaded into the microtrap, we use a 400 ms microwave evaporation to remove all but a small number of cold atoms, at a typical temperature of 15 μK . Part of the reason we use this quick and inefficient cut and not the kind of efficient evaporation that we use to achieve BEC is that we want to reduce the heating due to $Q4$ current, but it is also the case that the microtrap is so weak that the low initial density sets a very long timescale for initial evaporation, significantly longer than the one-body lifetime in the microtrap.

6.3 Measurement of Cavity Cooperativity η_1

The main parameter that governs both absorption and fluorescence in the cavity is the single-direction single-atom cooperativity, $\eta_1 = \frac{6\mathcal{F}}{\pi(w_0k)^2}$. While we can calculate the expected cooperativity, $\eta_1 = 0.07$, from basic cavity parameters (see Sec. 5.3), it is important to confirm it experimentally. To do this, we measure the tuning of the cavity transmission resonance by the atomic index of refraction for atomic samples large enough that the atom number can be determined by our standard CCD absorption imaging. The off-resonant tuning of the cavity by N atoms well-localized at the cavity axis is given by

$$\delta\nu = \frac{\kappa}{2} \frac{\Gamma}{\Delta} N \eta_1, \quad (6.3)$$

where κ is the cavity linewidth, Γ is the linewidth of the atomic transition, and $\Delta \gg \Gamma$ is the detuning between the laser and the atomic transition. For us, $\kappa = 2\pi \times 650$ kHz and $\Gamma = 2\pi \times 6.1$ MHz.

Before measuring η_1 , it is necessary to determine what applied magnetic fields




Figure 6.2: Plots of the single-direction cavity cooperativity, η_1 , extracted from a measurement of cavity tuning, $\delta\nu$, as the atoms are moved through the cavity by applied bias fields in the two transverse directions. The red curves show quadratic fits to the measured data. The B_x measurement was performed with a cloud of 5500 atoms, while the B_z measurement was performed with a cloud of 6000 atoms, both at temperatures around 15-20 μK . The detuning between the laser and the atomic resonance is $\Delta = 15\Gamma$ in both cases.

put the atoms on the cavity axis, where their coupling to the cavity is maximized. This is also measured using the cavity tuning $\delta\nu$, by varying the applied bias fields B_x and B_z and looking at the cavity tuning for a given atom number. The cavity tuning is measured by scanning over a predetermined range of cavity PZT voltages with and without atoms in the cavity, noting how much the cavity resonance shifts within that scan between the two situations, and converting that time difference into a frequency shift by calibrating the speed of the scan in units of cavity frequency. Once we know $\delta\nu$, we can extract η_1 using Eq. 6.3. The values for η_1 obtained from these measurements given the knowledge of the atom number N and the detuning Δ are shown in Fig. 6.2.

Once we can position the microtrap to optimize the atom-cavity coupling, we can measure η_1 for clouds of various size that are large enough for us to determine the atom number with absorption imaging. The results are shown in Fig. 6.3.

While the calculated η_1 for an atom localized on the cavity axis is 0.07, the value we measure is significantly smaller. This is due to the fact that the finite temperature,

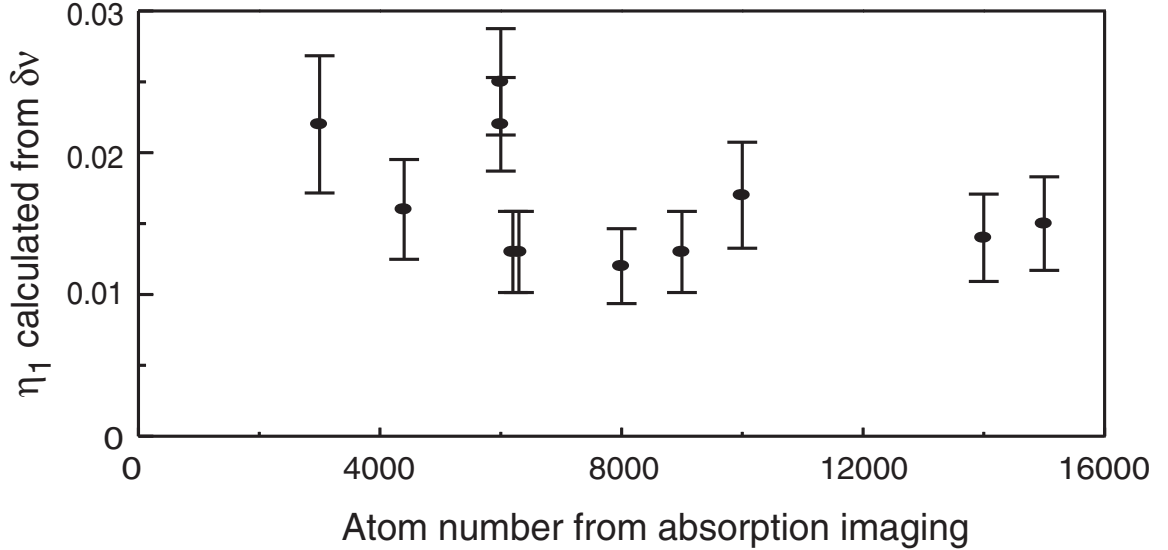


Figure 6.3: Plots of the single-direction cavity cooperativity, η_1 , extracted from a measurement of cavity tuning, $\delta\nu$, for samples with different atom numbers that are large enough to be measured via standard absorption imaging. The detuning between the laser and the atomic resonance is $\Delta = 15\Gamma$ in both cases.

and the resulting finite size, of the atom cloud reduce the coupling of the atoms to the cavity. There are two main reduction factors: a geometric factor g_c , which reduces the spatial overlap between the atom cloud and the cavity mode because many of the atoms are not located on the cavity axis; and an atom-light coupling reduction factor C_+ , which reduces the Clebsch-Gordan coefficient for the atoms' coupling to intracavity σ^+ light (which is the polarization that produces the maximal η_1 and is therefore used for this measurement) by smearing out the magnetic quantization axis, since atoms that have enough energy to move away from the trap bottom will experience transverse components of the magnetic field.

It is possible to calculate g_c given the known cavity mode diameter w and a measured transverse root-mean-square size of the atom cloud, r_{rms} . Integrating over the normalized transverse atom density distribution weighted by the relative intensity of intracavity light (normalized to 1 on the cavity axis) yields g_c :

$$g_c = \int_0^\infty \frac{1}{2\pi r_{rms}^2} e^{-r^2/(2r_{rms}^2)} \times e^{-2r^2/w^2} \times 2\pi r dr = \frac{w^2}{w^2 + 4r_{rms}^2}. \quad (6.4)$$

For the smaller clouds used in Fig. 6.3, we calculate g_c to be 0.4 after extracting the rms size of the cloud from a measurement similar to that shown in Fig. 6.2 and calculate C_+ to be 0.8 by numerically integrating the effective Clebsch-Gordan coefficient over the atom cloud. Given these two factors, we'd expect to measure $\eta_1 \approx 0.022$, which is consistent with what we observe for the smaller atom clouds. For the larger clouds, with $N > 8000$, we observe a lower value of η_1 , which makes sense since those clouds are hotter and therefore larger spatially, and so would have larger reduction factors due to finite cloud size.

6.4 Fluorescence Detection

6.4.1 Atom Number Measurement

For fluorescence detection, we illuminate the atoms from the side with a retroreflected pump beam at an angle of 70° to the cavity axis. In order to simplify retroreflection alignment, the beam is focused down at the planar retroreflection mirror (if it was instead focused at the atoms, we'd have to use a curved mirror or a planar mirror and a lens for retroreflection), which produces a spot size of $250 \mu\text{m}$ at the atoms. This is comparable to the distance of the atoms to the chip surface, so some of the beam is cut off by the chip. It is preferable for the fluorescence pump beam to be as close as possible to parallel to the cavity axis, which is also the magnetic field quantization axis, in order to maximize σ^+ polarization purity. However, the space through which the fluorescence beam can pass unobstructed is very limited, which restricts the angle at which it can come in and also possibly by the edge of the cavity mount and/or the mount for the absorption imaging lens. The retroreflected beam is thus attenuated in intensity compared to the original beam.

To optimize the coupling of the fluorescence pump beam to atoms in the cavity, we use loss of atoms due to scattering of pump light. We put the microtrap at the position where the atoms couple strongest to the cavity, and then shine on a pulse of pump light and observe, via CCD absorption imaging, how many atoms remain in the trap after the pulse as a function of fluorescence beam alignment. In order to

get good signal-to-noise, we use large and hot atom clouds ($N \sim 100,000$) for this measurement, but that should not affect the result for optimal alignment.

To characterize both the atom number preparation and the number of photon counts detected per atom in fluorescence detection, we illuminate the atoms with a pump beam nearly resonant ($\Delta/\Gamma = 1-3$) with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition and with a peak intensity slightly above saturation intensity ($I_s = 2 \text{ mW/cm}^2$), and count the photons emerging from the cavity within a window of $750 \mu\text{s}$ using a fiber-coupled single-photon counting module (Perkin-Elmer SPCM-AQR-12-FC counter based on a silicon avalanche photodiode). We do not measure exactly on resonance so that absorption from the atoms doesn't spoil the cavity finesse. We compile histograms of counts containing 100-250 shots for different RF final settings, which correspond to different average numbers of prepared atoms.

Assuming that the photons emitted by each atom obey Poisson statistics, the following relation can be derived:

$$\frac{\langle n^2 \rangle}{\langle n \rangle} - 1 = g_{aa} \langle n \rangle + \langle p \rangle, \quad (6.5)$$

where $\langle a \rangle$ is the mean atom number for the given histogram, $\langle p \rangle$ is the mean number of photon counts per atom, $n = ap$ is the total number of signal photon counts in a given shot, and $g_{aa} \equiv \frac{\langle a^2 \rangle - \langle a \rangle^2}{\langle a \rangle^2}$ is the atom-atom correlation function (see Appendix A.1 for the derivation). g_{aa} should be equal to 1 if the atoms obey Poisson statistics and equal to $(1 + f^2)$ in the presence of (technical) fractional atom number noise of magnitude f [70]. The values of $(\langle n^2 \rangle / \langle n \rangle) - 1$ can be computed from each histogram independently without any knowledge about $\langle a \rangle$ or $\langle p \rangle$, given that we can measure the background count rate independently, and assuming that the background is uncorrelated with the signal. The results plotted against $\langle n \rangle$, along with a linear fit, are shown in Fig. 6.4. The fit gives a slope of $g_{aa} = 1.05 \pm 0.02$, which implies that the fractional noise on our signal is 0.25 ± 0.10 , and therefore Poissonian fluctuations dominate for the atom numbers we measure; and an intercept of $\langle p \rangle = 1.9 \pm 0.3$ photon counts per atom. The fact that our atom preparation is shot-noise-limited down to 1 atom is itself a notable achievement, and it means we can perform various




Figure 6.4: Characterization of atom number fluctuations and detected photons per atom for fluorescence detection. Each data point corresponds to the statistics for one histogram. n is the total number of signal photons detected. The slope of the fit gives the atom-atom correlation function, g_{aa} , and the y-axis intercept gives the average number of photon counts per atom, $\langle p \rangle$.

small-atom-number experiments while being confident of our atom statistics.

Having confirmed the Poisson statistics of our atom number preparation, we can fit $\langle a \rangle$ and $\langle p \rangle$ for each histogram individually, assuming a Poisson distribution of atoms with mean $\langle a \rangle$, each of which emits a Poisson distribution of photons of which we detect a mean of $\langle p \rangle$. A typical histogram with fit, and a plot of the combined results of all histogram fits are shown in Fig. 6.5. To a good approximation, the average number of photon counts per atom is independent of atom number, with $\langle p \rangle = 2.0 \pm 0.2$ counts/atom, with 0.3 background counts in the 750 μs measurement window, dominated by the dark count rate of the single photon counter (SPCM). An average of the signal time traces yields a $1/e$ signal decay time of $\tau = 150 \mu\text{s}$, likely limited by the atoms' being pushed out of the cavity mode due to an imbalance

between the intensities of the original and retroreflected pump beams because of partial obstruction of the retroreflected beam by various mounts around the chip.

6.4.2 Expected Fluorescence Signal

To see if our result for $\langle p \rangle$ makes sense, we can estimate what we would expect $\langle p \rangle$ to be based on known experimental parameters. Since our cavity resonance is much narrower than the atomic line, the cavity collects predominantly the coherently scattered photons, the center peak of the Mollow triplet [71]. The number of photon counts we would expect to detect per atom is thus given by $\langle p \rangle = \Gamma_{coh} \tau \eta_1 (\kappa / \gamma_c) \times C_+ \times g_c \times f \times q \times l$, where $\Gamma_{coh} = \Gamma / 8 = 2\pi \times 760$ kHz is the maximum coherent scattering rate for the transition, $\eta_1 = 0.07$ is the cavity's intrinsic single-atom cooperativity, $\gamma_c = 2\pi \times 1$ MHz is the average linewidth of the cavity transmission, which is a convolution of the cavity and laser linewidths, $C_+ = 0.3$ is the calculated averaged Clebsch-Gordan coefficient for σ^+ intracavity light coming from the scattering process (the other polarizations are not resonant with the cavity; C_+ is much smaller here than in the measurement of η_1 described in Section 6.3 because of the large angle between the fluorescence beam and the magnetic quantization axis), $g_c = 0.6$ accounts for the imperfect overlap between the atoms and the cavity mode due to the finite size of the atomic cloud (this is larger than the $g_c = 0.4$ factor obtained in the measurement of η_1 because the much smaller atom numbers used for the fluorescence measurement are obtained in colder, and therefore spatially smaller, samples), $f = 0.7$ is the coupling efficiency into the single-mode fiber, $q = 0.58$ is the quantum efficiency of the SPCM at 780 nm, and $l = 0.7$ is the signal reduction due to mechanical cavity vibrations, measured independently via cavity transmission. The combination of the above factors predicts $\langle p \rangle = 1.7 \pm 0.9$, close to our measured value.

6.4.3 Single-Atom Detection

Since the scattering time of $150 \mu\text{s}$ is significantly shorter than the $750 \mu\text{s}$ collection window used above, reducing the measurement time improves the signal-to-noise ratio for small atom numbers, where the noise is dominated by the background and

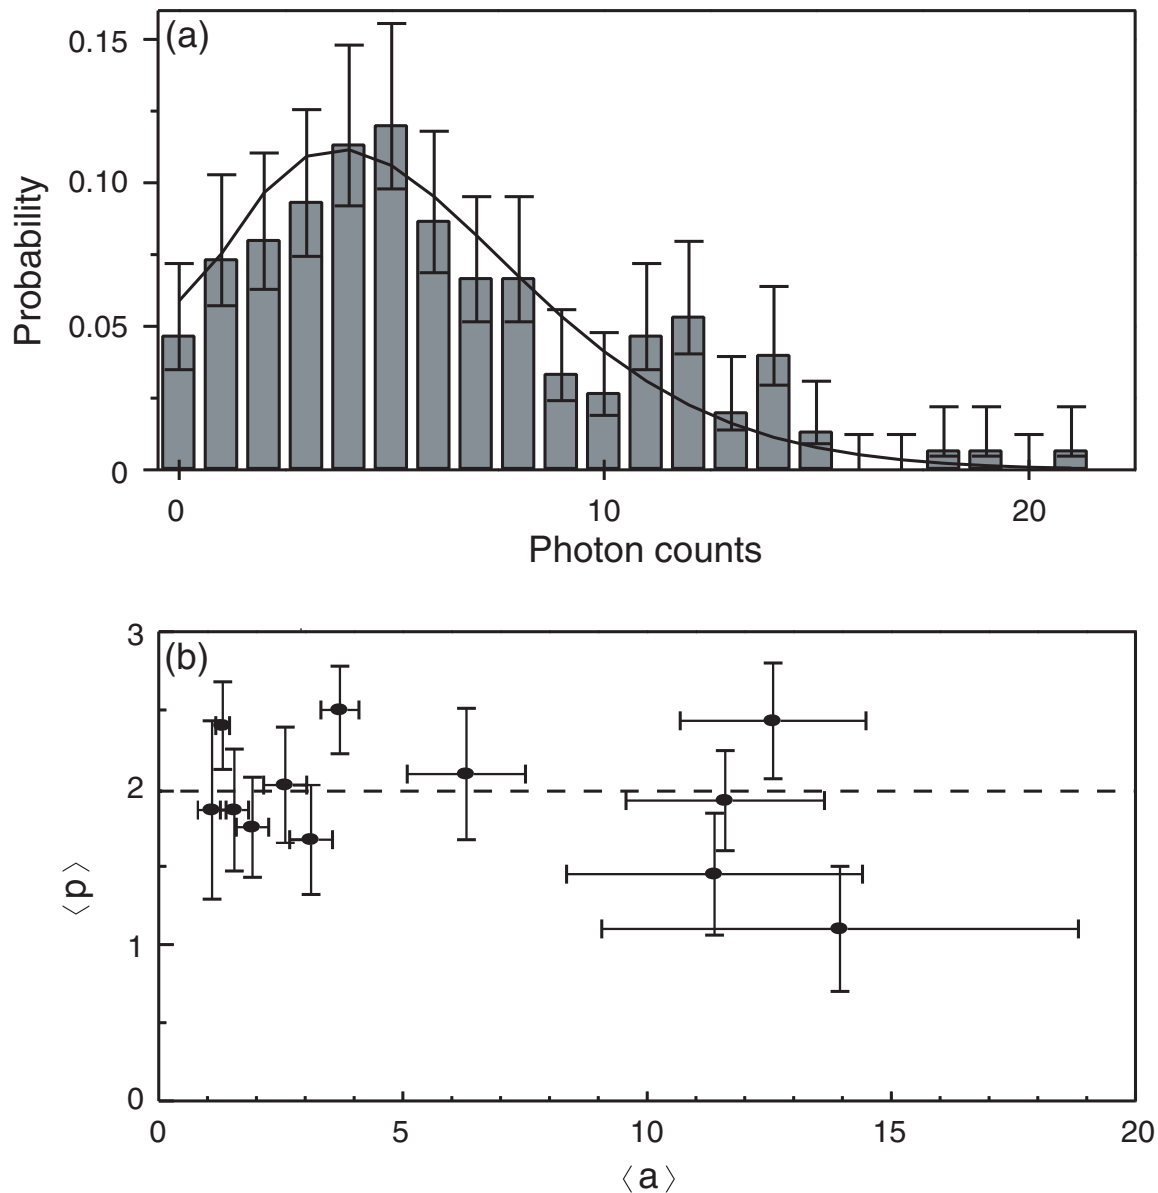


Figure 6.5: (a) Typical normalized histogram of 150 fluorescence measurements, with Poisson fit to $\langle a \rangle$, the mean number of atoms, and $\langle p \rangle$, the mean number of photon counts per atom (here, $\langle a \rangle = 3.1 \pm 0.4$ and $\langle p \rangle = 1.7 \pm 0.3$). Error bars are Poisson uncertainties given the measurement statistics. (b) Results of such fits to 12 different histograms; error bars correspond to 1 standard deviation in the fits for $\langle a \rangle$ and $\langle p \rangle$. The dashed line shows the combined fit for $\langle p \rangle$ of 2.0 counts/atom.

not the shot noise of the signal counts. Thus, in order to quantify our fluorescence measurement as a single-atom detector, we reduce the measurement window to 250 μs , and use a better SPCM (Perkin-Elmer SPCM-AQR-14-FC), which reduces the background count rate from 400 Hz to 280 Hz. We then take a fluorescence histogram of 200 shots with the minimal atom number that is still easy to distinguish from zero, with the result shown in Fig. 6.6. The Poisson fit to the resulting histogram gives $\langle a \rangle = 0.85 \pm 0.08$ and $\langle p \rangle = 1.4 \pm 0.1$. The fit is shown in Fig. 6.6 along with two comparison fits to $\langle a \rangle$ that assume a value for $\langle p \rangle$ that is either half or twice the best-fit value. Fig 6.6 includes a comparison of a computed normalized histogram produced by background counts to one produced by the signal from one atom.

The result of $\langle p \rangle = 1.4$, combined with a measured background of 0.07 counts in 250 μs , means that, if we set our detection threshold to ≥ 1 count, our single-atom detector is characterized by an atom quantum efficiency of 75% and a false detection rate of 7%, at a maximum single-atom count rate of 4 kHz. If we instead choose a threshold of ≥ 2 counts, the atom quantum efficiency drops to 40%, and the false detection rate to 0.2%, at the same fast maximum count rate.

6.5 Absorption Detection

6.5.1 Atom Number Measurement

While the fluorescence measurement makes a good single-atom detector, we expect an absorption measurement to provide better atom number resolution for $a > 1$, as derived in Section 6.1. For absorption detection, we couple the probe laser beam into the cavity TEM_{00} mode on atomic resonance and with an intracavity saturation parameter equal to 0.2, and monitor the resonant transmission through the cavity in with and without atoms trapped in the cavity mode volume. Since the unattenuated absorption beam is much more intense than typical observed fluorescence signals, the fractional shot noise is significantly smaller, which makes intensity noise due to cavity mechanical vibrations the dominant noise source. To make absorption measurements shot-noise limited, the laser linewidth is broadened, by frequency modulation coupled

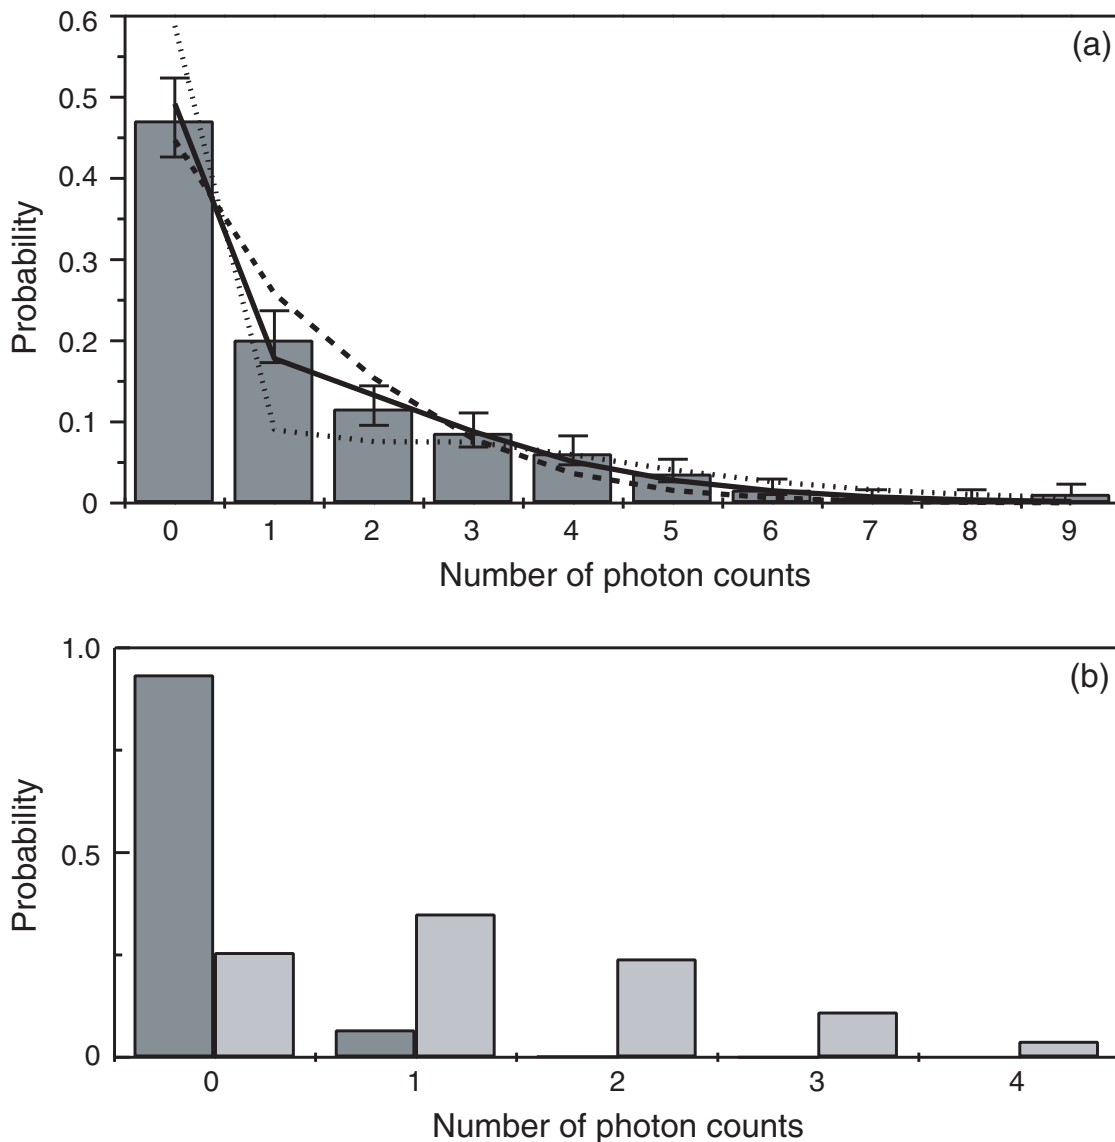


Figure 6.6: (a) Normalized histogram of 200 fluorescence measurements with small atom number. Three Poisson fits are included: best two-parameter fit of $\langle a \rangle = 0.85 \pm 0.08$, $\langle p \rangle = 1.41 \pm 0.14$ (solid line); and two single-parameter fits for $\langle a \rangle$ assuming a particular $\langle p \rangle$: $\langle a \rangle = 1.46$ given $\langle p \rangle = 0.7$ (dashed line); $\langle a \rangle = 0.49$ given $\langle p \rangle = 2.8$ (dotted line). The reduced χ^2 for these fits are 0.4, 1.8, and 3.7, respectively. Error bars are Poisson uncertainties given the measurement statistics. (b) Computed normalized histogram due to background photon counts (dark gray) and due to the signal from one atom (light gray).

in directly across the laser diode, to 30 MHz, much larger than the cavity linewidth κ , so that the intensity noise on the cavity transmission due to cavity vibrations is negligible compared to the photon shot noise.

Similarly to fluorescence detection, we compile histograms collected in 1 ms for different atom preparation parameters. However, since the range of possible count rates is much larger, we bin the counts in bins of 25 in order to achieve reasonable measurement statistics. The measured absorption “background,” i.e. the signal collected without any atoms, is $\langle n_0 \rangle = 3680$ counts, and the absorption signal lasts for a $1/e$ time of $560 \mu\text{s}$. We can characterize our atom preparation statistics by looking at the statistics of the histograms. In this case, we have to linearize the transmission dependence in order to be able to derive a useful statistical relation:

$$n = \frac{n_0}{1 + as} \approx n_0(1 - as), \quad (6.6)$$

where n is the detected photon number, n_0 is the background detected photon number, a is the atom number, and s is the fractional absorption per atom. Given the above approximation, we can derive the following:

$$\frac{\langle n_0 \rangle^2 (\langle n^2 \rangle - \langle n \rangle^2 - \langle n \rangle)}{\langle n_0 \rangle^2 \langle a \rangle \langle s \rangle} + \langle a \rangle \langle s \rangle = g_{aa} \langle a \rangle \langle s \rangle + \langle s \rangle, \quad (6.7)$$

where g_{aa} is the atom-atom correlation function and the expectation values are taken for the given histogram (see Appendix A.2 for the derivation). The values of the left-hand side in the above equation can be computed from each histogram independently without any knowledge about $\langle a \rangle$ or $\langle s \rangle$ (only their product, which can be determined from $\langle n \rangle$ and $\langle n_0 \rangle$ occurs on the left-hand side). The results plotted against the fractional absorption $\langle a \rangle \langle s \rangle$, along with a linear fit, are shown in Fig. 6.7. The fit gives a slope of $g_{aa} = 0.88 \pm 0.06$ and an intercept of $\langle s \rangle = 3.5 \pm 0.5\%$ /atom. The low value of g_{aa} is probably due to the fact that the linear approximation is not an entirely accurate description of all our absorption measurements, which is reasonable given that the histogram with the largest absorption has a fractional absorption of close to 20%. The linearized calculation may thus systematically underestimate g_{aa} . Nevertheless, it appears that this fit confirms that our atom number preparation for




Figure 6.7: Characterization of atom number fluctuations and fractional absorption per atom for absorption detection. Each data point corresponds to the statistics for one histogram. n is the total number of photons detected, n_0 is the number of photons expected without any atoms, and as is the fractional absorption due to the presence of atoms. The slope of the fit gives the atom-atom correlation function, g_{aa} , and the y-axis intercept gives the mean fractional absorption per atom, $\langle s \rangle$.

absorption has Poisson statistics.

Given Poisson atom preparation and photon absorption statistics, we can fit the histograms to extract values for the mean atom number $\langle a \rangle$ and the mean fractional transmission reduction, or absorption, per atom, $\langle s \rangle$ from each histogram individually. A typical histogram with fit, and a plot of the combined results of all histogram fits are shown in Fig. 6.8. From these measurements, we obtain $\langle s \rangle = 3.3 \pm 0.3\%$ /atom.

6.5.2 Expected Absorption Signal

The lineshape of the intensity of cavity transmission for an infinitely narrow laser is given by the square of the transmitted electric field amplitude:

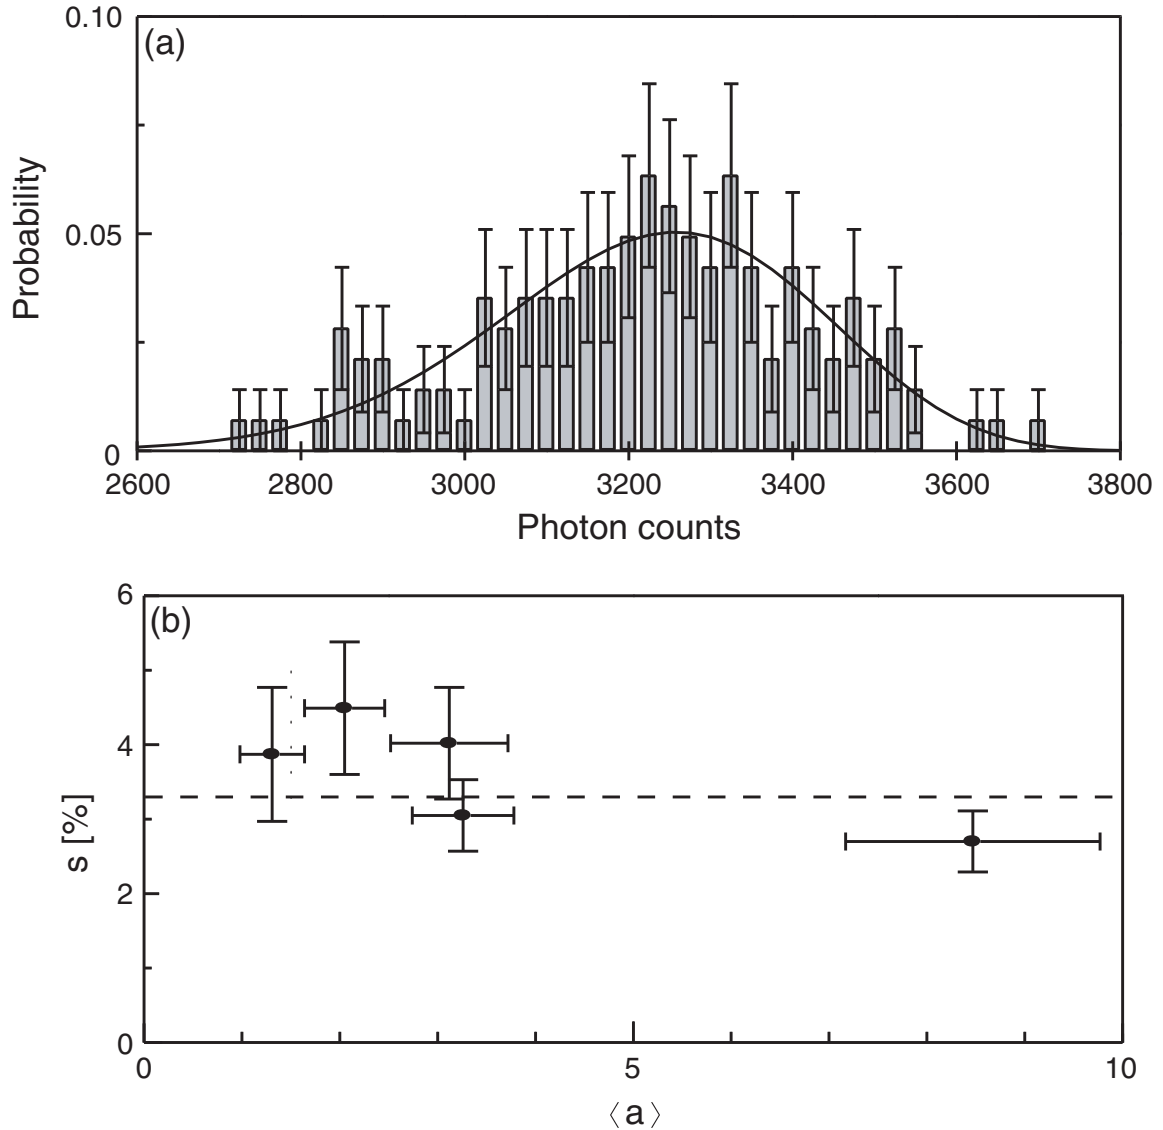


Figure 6.8: (a) Typical normalized histogram of 142 absorption measurements, with Poisson fit to $\langle a \rangle$, the mean number of atoms, and $\langle s \rangle$, the mean fractional transmission reduction per atom. The signal collected without any atoms is $\langle n_0 \rangle = 3680$ counts. Error bars are Poisson uncertainties given the measurement statistics. (b) Results of such fits to 5 different histograms; error bars correspond to 1 standard deviation in the fits for $\langle a \rangle$ and $\langle s \rangle$. The dashed line shows the combined fit for $\langle s \rangle$ of 3.3 %/atom.

$$\left| \frac{q^2}{q^2 + i(2\pi\Delta_c/\nu_{FSR})} \right|^2 = \frac{(\kappa/2)^2}{(\kappa/2)^2 + \Delta_c^2}, \quad (6.8)$$

where $q^2 = \pi/\mathcal{F}$ is the single-mirror transmission with \mathcal{F} as the cavity finesse, κ is the cavity linewidth, Δ_c is the detuning from the cavity resonance, and $2\pi\Delta_c/\nu_{FSR}$ is the phase accumulated by light detuned by Δ_c from the cavity resonance over one intracavity round trip. This equation corresponds to a Lorentzian lineshape with linewidth κ .

In the presence of an atom, the transmission lineshape for a cavity is given by

$$\left| \frac{q^2}{q^2 + l_{at}^2 + i(2\pi\Delta_c/\nu_{FSR})} \right|^2 = \left| \frac{q^2}{q^2 + l_{at}^2} \right|^2 \times \left| \frac{q^2 + l_{at}^2}{q^2 + l_{at}^2 + i(2\pi\Delta_c/\nu_{FSR})} \right|^2 \quad (6.9)$$

where l_{at}^2 is the round-trip electric field amplitude reduction due to absorption by the atom, and we have assumed that the atomic and cavity resonances overlap and the atomic linewidth is much greater than κ , so that the contribution to the round-trip phase due to atomic detuning is negligible compared to that due to cavity detuning. The round-trip electric field amplitude reduction, l_{at}^2 , is equal to the one-direction intensity reduction, which is just the absorption fraction calculated in Sec. 6.1 to be equal to $12/(w_0k)^2 = 2q^2\eta_1$. The cavity linewidth in the presence of an atom is given by

$$\kappa' = \frac{q^2 + l_{at}^2}{q^2} \kappa = (1 + 2\eta_1)\kappa, \quad (6.10)$$

and the cavity transmission lineshape is given by

$$\left(\frac{1}{1 + 2\eta_1} \right)^2 \times \frac{(\kappa'/2)^2}{(\kappa'/2)^2 + \Delta_c^2} \approx \frac{1}{1 + 4\eta_1} \times \frac{(\kappa'/2)^2}{(\kappa'/2)^2 + \Delta_c^2}, \quad (6.11)$$

for $\eta_1 \ll 1$. In the low-total-absorption linear regime, the mean fractional transmission reduction per atom, $\langle s \rangle$, is then equal to $4\eta_1$.

The above analysis is correct for a probe laser which is very narrow compared to cavity (and atomic) resonance. However, in our absorption measurement, we

modulate the probe laser to make it much broader than both cavity and atomic resonances. To incorporate the effects of this broadening, let us consider the convolution of the above cavity transmission line with a laser line that has an intrinsic Lorentzian linewidth Γ_l and is broadened to a much larger linewidth Γ_{br} .

The convolution of a normalized unbroadened laser with frequency ω_l and linewidth Γ_L with a cavity line with resonance frequency ω_c and linewidth κ and a normalization factor that makes the cavity transmission unity on resonance is given by

$$\int_{-\infty}^{+\infty} \frac{1}{\pi} \cdot \frac{\Gamma_l/2}{(\Gamma_l/2)^2 + (\omega - \omega_l)^2} \cdot \frac{(\kappa/2)^2}{(\kappa/2)^2 + (\omega - \omega_c)^2} d\omega = (\kappa/2) \frac{\Gamma_l/2 + \kappa/2}{(\Gamma_l/2 + \kappa/2)^2 + (\omega_c - \omega_l)^2} \quad (6.12)$$

To incorporate the broadening, let us integrate this convolution over a distribution of laser lines centered in the range between $\omega_c - \Gamma_{br}$ and $\omega_c + \Gamma_{br}$, and, since $\Gamma_{br} \gg \Gamma_l, \kappa$, take the limit $\Gamma_{br} \rightarrow \infty$:

$$\lim_{\Gamma_{br} \rightarrow \infty} \int_{\omega_c - \Gamma_{br}}^{\omega_c + \Gamma_{br}} (\kappa/2) \frac{\Gamma_l/2 + \kappa/2}{(\Gamma_l/2 + \kappa/2)^2 + (\omega_c - \omega_l)^2} d\omega_l = \frac{\pi\kappa}{2}. \quad (6.13)$$

If we incorporate the cavity line broadening and transmission attenuation due to the presence of one atom, this becomes

$$\begin{aligned} \lim_{\Gamma_{br} \rightarrow \infty} \int_{\omega_c - \Gamma_{br}}^{\omega_c + \Gamma_{br}} (\kappa'/2) \frac{\Gamma_l/2 + \kappa'/2}{(\Gamma_l/2 + \kappa'/2)^2 + (\omega_c - \omega_l)^2} \cdot \frac{1}{(1 + 2\eta_1)^2} &= \\ &= \frac{\pi\kappa'}{2} \cdot \frac{1}{(1 + 2\eta_1)^2} = \frac{\pi\kappa}{2} \cdot \frac{1}{(1 + 2\eta_1)}. \end{aligned} \quad (6.14)$$

Thus, the ratio between the peak transmission in the presence of an atom to that without an atom is then given by $1/(1 + 2\eta_1)$, which means that, if we include the effects of the laser broadening, the linearized mean fractional transmission reduction per atom, $\langle s \rangle$, is equal to $2\eta_1$. This means that when we broaden the probe laser in order for our absorption measurement to be limited by photon shot noise, we lose half our signal.

The expected fractional transmission reduction per atom is then given by $\langle s \rangle =$

$2\eta_1 \times C_+ \times g_c \times r_\tau$, where $C_+ = 0.8$ is the calculated reduction due to smearing out of the magnetic quantization axis, $g_c = 0.6$ is the geometric overlap reduction factor calculated from the measured finite cloud size, and $r_\tau = 0.47$ is the ratio between the transmission reduction averaged over the 1 ms measurement window and the peak transmission reduction, given the scattering lifetime of 560 μs . If we put in all the reduction factors, we expect $\langle s \rangle = 3.2 \pm 0.7 \%$ /atom, which agrees very well with our observed value.

6.6 Comparison of Fluorescence and Absorption

Using the measured values of $\langle p \rangle = 2.0 \pm 0.2$ counts/atom for fluorescence and $\langle s \rangle = 3.3 \pm 0.3\%$ /atom for absorption, we can evaluate how well these two methods can determine the atom number in a single measurement. The expected fluorescence detection atom number uncertainty $\Delta_f a$ and the expected absorption detection atom number uncertainty $\Delta_a a$ are both plotted as a function of atom number in Fig. 6.9.

The atom number confidence intervals plotted in Fig. 6.9 include both the fundamental uncertainties discussed in Sec. 6.1 and additional technical uncertainties. For fluorescence, $\Delta_f a$ is limited by both the shot noise of the detected single photons and the statistical uncertainty in the mean number of photons per atom $\langle p \rangle$ given by our combined histogram fit, as well as by the background photon counts (for the detector with larger dark count rate used for the original fluorescence measurement, which gives a background of $n_{bg} = 0.3$ counts). Quantitatively,

$$(\Delta_f a)^2 = \left(\frac{\partial a}{\partial n_f} \right)^2 (\delta_P n_f)^2 + \left(\frac{\partial a}{\partial \langle p \rangle} \right)^2 (\delta \langle p \rangle)^2 + \left(\frac{\partial a}{\partial n_{bg}} \right)^2 (\delta n_{bg})^2, \quad (6.15)$$

where $n_f = a \langle p \rangle$ is the total number of signal photons detected, so $\partial a / \partial n_f = 1 / \langle p \rangle$, and $\delta_P n_f$ is the Poisson uncertainty (shot noise) in n_f ; $\partial a / \partial \langle p \rangle = -a / \langle p \rangle$ and $\delta \langle p \rangle = 0.2$ is the statistical uncertainty in our measured value of $\langle p \rangle$; n_{bg} is the number of background photon counts detected, $\partial a / \partial n_{bg} = \partial a / \partial n_f = 1 / \langle p \rangle$ because fluctuations in background photon counts are indistinguishable from fluctuations in




Figure 6.9: Single-shot atom-number measurement $1\text{-}\sigma$ confidence intervals for fluorescence ($\langle p \rangle = 2.0 \pm 0.2$ counts/atom, measurement window $t = 750 \mu\text{s}$, background of 0.3 counts; dark gray) and absorption ($\langle s \rangle = 3.3 \pm 0.3$ %/atom, $t = 1$ ms; light gray), assuming proper calibration (so that the mean inferred atom number is the same as the actual atom number).

signal photon counts, and δn_{bg} is the uncertainty in the number of background photons detected; assuming the background is independent of our signal and Poisson-distributed, $(\delta n_{bg})^2 = n_{bg} = 0.3$. Of the three terms that contribute to $\Delta_f a$, only the first is fundamental; the uncertainty in $\langle p \rangle$ and the background are technical limits.

For absorption, $\Delta_a a$ is limited by both the shot noise of the detected signal photons and the statistical uncertainty in the mean absorption per atom $\langle s \rangle$ given by our combined histogram fit. Quantitatively,

$$(\Delta_a a)^2 = \left(\frac{\partial a}{\partial n_a} \right)^2 (\delta_P n_a)^2 + \left(\frac{\partial a}{\partial \langle s \rangle} \right)^2 (\delta \langle s \rangle)^2, \quad (6.16)$$

where $n_a = n_0/(1 + a\langle s \rangle)$ is the total number of photons detected, so $\partial a/\partial n_a = (1 + a\langle s \rangle)^2/(n_0\langle s \rangle)$, $n_0 = 3680$, and $\delta_P n_a$ is the Poisson uncertainty (shot noise) in n_a , which, for the large values of n_a we observe, means $(\delta_P n_a)^2 = n_a$; $\partial a/\partial \langle s \rangle = -a/\langle s \rangle$ and $\delta \langle s \rangle = 0.3\%$ is the statistical uncertainty in our measured value of $\langle s \rangle$. Only the first of the two terms that contribute to $\Delta_a a$ is fundamental; the uncertainty in $\langle s \rangle$ is a technical limit.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

The work described in this thesis investigates some of the important technical capabilities and fundamental limits of experiments that use magnetic microtraps on atom chips to trap and manipulate ultracold atoms.

The first question addressed in this work is, “How close we can get to atom chips?” The kind of miniaturization that can be achieved in microtraps has several advantages, but there are also disruptive effects from the chips themselves that set in at the small length scales they make accessible. The simple question above really amounts to two questions: (1) How close can we get to the current-carrying elements on chips, and what is the limiting mechanism?; and (2) How close can we get to the dielectric chip surface, and what is the limiting mechanism? While the limiting mechanisms of (1) Johnson-noise-induced spin flips and (2) the Casimir-Polder force were not unexpected, our experiments quantify both of these limits in terms of chip and magnetic trap parameters. Our results are generally applicable to all atom chip experiments that rely on miniaturization.

It turns out that a relatively simple interpolation formula is capable to describing the spin-flip rate due to Johnson noise very accurately; it can therefore be used as a good guideline to predict the performance of a given chip wire layout at a particular length scale. In terms of engineering atom chips, a good strategy seems to be to only

approach close to chip conductors when the high gradient made possible by doing so is necessary for the desired experiment; magnetic fields that don't need to have large gradients should be produced by conductors far away from the atom sample and stages of the experiment that don't require close proximity to chip conductors (such as the evaporation to BEC in our experiment) should be performed far away from the chip. Furthermore, the fact that the spin-flip rate does not depend on current in the conductor suggests that there should be no unused conductors on the chip surface.

The sharpness of the Casimir-Polder potential means that only a few microns' separation from the dielectric chip surface is enough to render it completely benign. Placing chip wires on low dielectric ridges above the surface, as done on our modified chip, is thus one way to circumvent the dielectric surface's destructive effects while still taking advantage of its capabilities both for microfabrication and for heat dissipation.

We are encouraged by our ability to use the magnetic microtrap to split a BEC into two clouds with equal atom number. While the splitting method suffers from several technical limitations, it suggests a possible direction of research to achieve trapped atom interferometry on an atom chip.

One of the interesting peripheral aspects of the surface effects experiments is the extreme sensitivity of ultracold atoms to external potentials and their resulting usefulness as probes of their environment. By using loss of trapped atoms as our signal, we are able to relatively easily determine atom positions and resolve surface structures to within less than 100 nm, and technical improvements in this direction should increase this technique's resolution substantially. The potential of ultracold atoms as sensors of surface potentials and of electric and magnetic fields is only beginning to be realized.

Another important problem in atom chip experiments that is addressed by our work is one of signal detection. One of the primary appeals of atom chips is their ability to make possible various kinds of quantum atom optics devices, characterized by interesting and novel statistics and correlations between atoms. In order to be able to implement such devices, there needs to be a way to sensitively measure these statistics and correlations. To this end, we implement an atom detector using a medium-finesse cavity mounted on an atom chip. The detector, integrated in a

straightforward manner into our existing experimental apparatus, is sensitive enough to detect single atoms with high quantum efficiency and a fast maximum count rate, while at the same time it is relatively easy to optimize and use (so that experimenters' energies can be focused on the rest of the experiment), robust, and versatile (in the sense that it should be possible to integrate it into many different kinds of atom chip experiments). The existence of such an on-chip atom detector is an important step toward many microtrap applications.

It should be noted that many of the atom optics applications for which microtraps show promise can also potentially be implemented in other experimental regimes, such as in optical traps and lattices, or with untrapped atoms, and, indeed, in some areas, these other techniques are far ahead of what has been realized in microtraps. However, microtraps offer unparalleled versatility in terms of creating complex potentials for manipulating cold atoms, and this versatility suggests that the area of microtraps on chips should be a fertile source of exciting experimental results for a long time to come.

7.2 Future Work

We have plans for several different experiments that would extend or build upon the work described here, in both the atom chip and in the cavity directions.

Our surface effects experiments suggest that continued miniaturization of some parts of the experiment is both possible and desirable. In particular, if we wish to create a magnetic double well where the central barrier is sharp enough to permit an appreciable amount of tunneling between the two sides (this would be an implementation of a BEC Josephson junction), we need to go to distances on the order of a micron from the conductor that generates that potential barrier. All other magnetic fields and gradients can be generated by conductors much farther away from the atoms, and the magnetic field barrier does not require a large current, so it could be generated by a wire just a few hundred nm wide, which would give slow spin-flip loss rates at micron distances. The modified chip used in the cavity experiment has

this kind of small splitting wire on it, which, unfortunately, is damaged and unusable; future atom chips should be able to implement this kind of splitting. Two such sharp tunnel barriers could be used to create an atomic Fabry-Perot resonator, a single-atom source where all atoms come out with the same energy.

In terms of cavity-aided atom detection, there are several improvements that are currently being made to the apparatus. The mirrors in our cavity have maximum reflectivity at 850 nm, so the cavity's finesse, which is 8600 at 780 nm, could be significantly higher at longer wavelengths. The ^{87}Rb D₁ line ($5^2\text{S}_{1/2} \rightarrow 5^2\text{P}_{1/2}$) wavelength is 795 nm, and we expect mirror transmission at that wavelength to be 5-10 times lower than at the 780 nm D₂ wavelength used in all the experiments presented here, which would correspond to a cavity finesse higher by the same factor, assuming losses are still negligible. A higher finesse achieved in the same cavity implies a cavity linewidth narrower by the same factor, so taking full advantage of it would require a laser at 795 nm with a linewidth of less than 100 kHz. A 795 nm external-cavity diode laser that uses an external Fabry-Perot cavity to narrow its linewidth via optical feedback [72] is currently being integrated into our apparatus. The cavity length stabilization also must be improved to work with a narrower linewidth. If successful, using a narrow laser to probe the D₁ line may increase our atom detection signals by as much as an order of magnitude.

In addition, we are currently trying to implement a scheme for using the cavity to achieve spin squeezing on an atomic clock transition in ^{87}Rb . While the most precise atomic clocks are currently at the standard quantum limit, for which the precision scales as \sqrt{N} , where N is the number of atoms composing the clock, spin squeezing can be used to achieve precision below the standard quantum limit, with the ultimate fundamental precision limit being the Heisenberg limit, for which precision increases as N . The general idea for spin squeezing is to exploit the cavity's enhancement of optical depth and the indistinguishability of forward scattering into the cavity – since all atoms scatter into the same cavity mode, it is fundamentally impossible to tell which atom scattered a given photon that comes out of the cavity. By observing light coming out of the cavity, it is possible to create many-atom correlated states that will have different behavior from an uncorrelated coherent state. In particular, it is

possible to use the atoms' tuning of the cavity resonance to extract the population difference between two atomic clock states without obtaining any information about the states of the specific atoms. The uncertainty in relative populations can thus be reduced below the shot noise limit (while the uncertainty in relative phase will be increased), and this "squeezed" uncertainty can then be mapped onto the final projection noise of a Ramsey-type clock measurement. While atomic spin squeezing can be achieved in free space without a resonator[73], the large optical depth that can be achieved in a cavity promises substantial improvement over free-space measurements. The details of our implementation of this technique are still being worked out, but the initial calculations given our experimental parameters appear very promising.

Appendix A

Derivations of Correlation Relations

A.1 Fluorescence

In order to derive the formula used to analyze the fluorescence data, let us consider a series of experimental shots, in each of which a sample with a atoms scatters photons from an incoming pump beam, n of which are detected by a photon counter, and define $p = n/a$ as the number of photon counts per atom for each experimental shot. Let $\langle a \rangle$, $\langle p \rangle$, and $\langle n \rangle$ be the expectation values for a , p , and n , respectively. Then,

$$\begin{aligned} n &= ap \\ \langle n \rangle &= \langle a \rangle \langle p \rangle \end{aligned} \tag{A.1}$$

If c_i is the probability of having i atoms in the sample, and $m_{j,i}$ is the probability of detecting j photons given i atoms, then

$$\langle n^2 \rangle = \sum_{i,j} c_i m_{j,i} j^2 = \sum_i c_i \langle n_i^2 \rangle, \tag{A.2}$$

where $\langle n_i^2 \rangle$ is the expectation value of n^2 at fixed atom number i . Assuming

Poisson statistics for the photons at fixed atom number,

$$\langle n_i^2 \rangle = \langle n_i \rangle^2 + \langle n_i \rangle = \langle p \rangle^2 i^2 + \langle p \rangle i. \quad (\text{A.3})$$

Substituting this into Eq. A.2 yields

$$\langle n^2 \rangle = \sum_i c_i (\langle p \rangle^2 i^2 + \langle p \rangle i) = \langle p \rangle^2 \sum_i c_i i^2 + \langle p \rangle \sum_i c_i i. \quad (\text{A.4})$$

By definition,

$$\begin{aligned} \sum_i c_i i^2 &= \langle a^2 \rangle \\ \sum_i c_i i &= \langle a \rangle, \end{aligned} \quad (\text{A.5})$$

and therefore,

$$\langle n^2 \rangle = \langle p \rangle^2 \langle a^2 \rangle + \langle p \rangle \langle a \rangle = \langle p \rangle^2 \langle a^2 \rangle + \langle n \rangle. \quad (\text{A.6})$$

Rearranging terms, we get,

$$\begin{aligned} \frac{\langle n^2 \rangle}{\langle n \rangle} - 1 &= \frac{\langle p \rangle^2 \langle a^2 \rangle}{\langle n \rangle} = \frac{\langle p \rangle \langle a^2 \rangle}{\langle a \rangle} = \frac{\langle n \rangle \langle a^2 \rangle}{\langle a \rangle^2} = \frac{\langle n^2 \rangle}{\langle n \rangle} - 1 = \langle n \rangle \left(\frac{\langle a^2 \rangle}{\langle a \rangle^2} - \frac{1}{\langle a \rangle} + \frac{1}{\langle a \rangle} \right) = \\ &= \langle n \rangle \left(\frac{\langle a^2 \rangle - \langle a \rangle}{\langle a \rangle^2} + \frac{1}{\langle a \rangle} \right) = \langle n \rangle g_{aa} + \langle p \rangle, \end{aligned} \quad (\text{A.7})$$

where $g_{aa} \equiv \frac{\langle a^2 \rangle - \langle a \rangle}{\langle a \rangle^2}$ is the atom-atom correlation function. This is the relation in Eq. 6.5 that we wanted to show.

A.2 Absorption

The overall form of this derivation is identical to the one above; only details of the formulas differ. For absorption, let us consider a series of experimental shots, in each

of which an incoming probe beam is attenuated due to absorption from a atoms, with the result that n photons are detected by a photon counter. Let us define n_0 as the number of photons that would have been detected by the photon counter if the atoms were not present, and let $s = (1 - n/n_0)/a$ be the fractional absorption per atom for each experimental shot. Let $\langle a \rangle$, $\langle s \rangle$, $\langle n_0 \rangle$ and $\langle n \rangle$ be the expectation values for a , s , n_0 , and n , respectively. Then,

$$\begin{aligned} n &= n_0(1 - as) \\ \langle n \rangle &= \langle n_0 \rangle(1 - \langle a \rangle \langle s \rangle) \end{aligned} \quad (\text{A.8})$$

If c_i is the probability of having i atoms in the sample, and $m_{j,i}$ is the probability of detecting j photons given i atoms, then

$$\langle n^2 \rangle = \sum_{i,j} c_i m_{j,i} j^2 = \sum_i c_i \langle n_i^2 \rangle, \quad (\text{A.9})$$

where $\langle n_i^2 \rangle$ is the expectation value of n^2 at fixed atom number i . Assuming Poisson statistics for the photons at fixed atom number,

$$\langle n_i^2 \rangle = \langle n_i \rangle^2 + \langle n_i \rangle = \langle n_0 \rangle(1 - i\langle s \rangle)^2 + \langle n_0 \rangle(1 - i\langle s \rangle). \quad (\text{A.10})$$

Substituting this into Eq. A.9 yields

$$\begin{aligned} \langle n^2 \rangle &= \sum_i c_i [\langle n_0 \rangle(1 - i\langle s \rangle)^2 + \langle n_0 \rangle(1 - i\langle s \rangle)] = \\ &= \langle n_0 \rangle^2(1 - 2\langle s \rangle) \sum_i c_i i + \langle s \rangle^2 \sum_i c_i i^2 + \langle n_0 \rangle(1 - \langle s \rangle) \sum_i c_i i. \end{aligned} \quad (\text{A.11})$$

By definition,

$$\sum_i c_i i^2 = \langle a^2 \rangle$$

$$\sum_i c_i i = \langle a \rangle, \quad (\text{A.12})$$

and therefore,

$$\begin{aligned} \langle n^2 \rangle &= \langle n_0 \rangle^2 (1 - 2\langle s \rangle \langle a \rangle + \langle s \rangle^2 \langle a^2 \rangle) + \langle n_0 \rangle (1 - \langle s \rangle \langle a \rangle) = \\ &= \langle n_0 \rangle^2 (1 - 2\langle s \rangle \langle a \rangle + \langle s \rangle^2 \langle a^2 \rangle) + \langle n_0 \rangle^2 \langle s \rangle^2 (\langle a^2 \rangle - \langle a \rangle^2) + \langle n \rangle = \\ &= \langle n \rangle^2 + \langle n \rangle + \langle n_0 \rangle^2 \langle s \rangle^2 \langle a \rangle^2 \left(g_{aa} - 1 + \frac{1}{\langle a \rangle} \right), \end{aligned} \quad (\text{A.13})$$

where $g_{aa} \equiv \frac{\langle a^2 \rangle - \langle a \rangle^2}{\langle a \rangle^2}$ is the atom-atom correlation function. Rearranging terms, we get,

$$\frac{\langle n^2 \rangle - \langle n \rangle^2 - \langle n \rangle}{\langle n_0 \rangle^2 \langle a \rangle \langle s \rangle} + \langle a \rangle \langle s \rangle = \langle a \rangle \langle s \rangle \left(g_{aa} - 1 + \frac{1}{\langle a \rangle} \right) + \langle a \rangle \langle s \rangle = \langle a \rangle \langle s \rangle g_{aa} + \langle s \rangle. \quad (\text{A.14})$$

This is the relation in Eq. 6.7 that we wanted to show.

Bibliography

- [1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269:198–201, 1995.
- [2] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969–73, 1995.
- [3] C. C. Bradley, C. A. Sackett, and R. G. Hulet. Bose-Einstein condensation of lithium: observation of limited condensate number. *Phys. Rev. Lett.*, 78:985–9, 1997.
- [4] D. M. Kurn D. S. Durfee C. G. Townsend M.-O. Mewes, M. R. Andrews and W. Ketterle. Output coupler for Bose-Einstein condensed atoms. *Phys. Rev. Lett.*, 78:582–5, 1997.
- [5] B. P. Anderson and M. A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, 282:1686–9, 1998.
- [6] I. Bloch, T. W. Hänsch, and T. Esslinger. Atom laser with a cw output coupler. *Phys. Rev. Lett.*, 82:3008–11, 1999.
- [7] E. W. Hagley, I. Deng, M. Kozuma, J. Wen, K. Helmerson, S. L. Rolston, and W. D. Phillips. A well-collimated quasi-continuous atom laser. *Science*, 283:1706–9, 1999.

- [8] D. Müller, D. Z. Anderson, R. J. Grow, P. D. D. Schwindt, and E. A. Cornell. Guiding neutral atoms around curves with lithographically patterned current-carrying wires. *Phys. Rev. Lett.*, 83:5194–7, 1999.
- [9] N. H. Dekker, C. S. Lee, V. Lorent, J. H. Thywissen, S. P. Smith, M. Drndić, R. M. Westervelt, and M. Prentiss. Guiding neutral atoms on a chip. *Phys. Rev. Lett.*, 84:1124–7, 2000.
- [10] J. Fortágh, H. Ott, A. Grossman, and C. Zimmermann. Miniaturized magnetic guide for neutral atoms. *App. Phys. B*, 70:701–8, 2000.
- [11] D. Cassettari, B. Hessmo, R. Folman, T. Majer, and J. Schmiedmayer. Beam splitter for guided atoms. *Phys. Rev. Lett.*, 85:5483–7, 2000.
- [12] W. Hänsel, P. Hommehoff, T. W. Hänych, and J. Reichel. Bose-Einstein condensation on a microelectronic chip. *Nature*, 413:498–501, 2001.
- [13] H. Ott, J. Fortágh, G. Schlotterbeck, A. Grossman, and C. Zimmermann. Bose-Einstein condensation in a surface microtrap. *Phys. Rev. Lett.*, 87:230401/1–4, 2001.
- [14] A. E. Leanhardt, Y. Shin, A. P. Chikkatur, , D.Kielinski, W. Ketterle, and D. Pritchard. Bose-Einstein condensation near a microfabricated surface. *Phys. Rev. Lett.*, 90:100404/1–4, 2003.
- [15] Y. Shin, M. Saba, T. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt. Atom interferometry with Bose-Einstein condensates in a double-well potential. *Phys. Rev. Lett.*, 92:050405/1–4, 2004.
- [16] Y.J. Wang, D. Z. Anderson, V. M. Bright, E. A. Cornell, Q. Diot, T. Kishimoto, M. Prentiss, R. A. Saravanan, S. R. Segal, and S. Wu. Atom Michelson interferometer on a chip using a Bose-Einstein condensate. *Phys. Rev. Lett.*, 94:090405/1–4, 2005.

- [17] Y. Shin, C. Sanner, G.-B. Jo, T. A. Pasquini, M. Saba, W. Ketterle, D. E. Pritchard, M. Vengalattore, and M. Prentiss. Interference of Bose-Einstein condensates split with an atom chip. *Phys. Rev. A*, 72:21604/1–4, 2005.
- [18] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger. Matter-wave interferometry in a double well on an atom chip. *Nature Physics*, 1:57–62, 2005.
- [19] M. Albiez, R. Gati, J. Fölling, S. Hunsmann, M. Cristiani, and M. K. Oberthaler. Direct observation of tunneling and nonlinear self-trapping in a single bosonic Josephson junction. *Phys. Rev. Lett.*, 95:010402/1–4, 2005.
- [20] M. Wilkens, E. Goldstein, B. Taylor, and P. Meystre. Fabry-Perot interferometer for atoms. *Phys. Rev. A*, 47:2366–9, 1993.
- [21] I. Carusotto. Nonlinear atomic Fabry-Perot interferometer: From the mean-field theory to the *atom blockade* effect. *Phys. Rev. A*, 63:023610/1–12, 2001.
- [22] M. Olshanii. Atomic scattering in the presence of external confinement and a gas of impenetrable bosons. *Phys. Rev. Lett.*, 81:938–41, 1998.
- [23] V. Dunjko, V. Lorent, and M. Olshanii. Bosons in cigar-shaped traps: Thomas-Fermi regime, Tonks-Girardeau regime, and in between. *Phys. Rev. Lett.*, 86:5413–16, 2001.
- [24] B. L. Tolra, K. M. O’Hara, J. H. Huckans, W. D. Phillips, S. L. Rolston, and J. V. Porto. Observation of reduced three-body recombination in a correlated 1D degenerate Bose gas. *Phys. Rev. Lett.*, 92:190401/1–4, 2004.
- [25] B. Paredes, A. Widera, V. Murg, O. Mandel, S. Fölling, I. Cirac, G. V. Shlyapnikov, T. W. Hänsch, and I. Bloch. Tonks-Girardeau gas of ultracold atoms in an optical lattice. *Nature*, 429:277–81, 2004.
- [26] T. Kinoshita, T. Wenger, and D. S. Weiss. Local pair correlations in one-dimensional Bose gases. *Phys. Rev. Lett.*, 95:190406/1–4, 2005.

- [27] L. Tonks. The complete equation of state for one, two, and three-dimensional gases of hard elastic spheres. *Phys. Rev.*, 50:955–63, 1936.
- [28] M. Girardeau. Relationship between systems of impenetrable bosons and fermions in one dimension. *J. Math. Phys.*, 1:516–23, 1960.
- [29] A. E. Leanhardt, A. P. Chikkatur, D. Kielpinski, Y. Shin, T. L. Gustavson, W. Ketterle, and D. E. Pritchard. Propagation of Bose-Einstein condensates in a magnetic waveguide. *Phys. Rev. Lett.*, 89:040401/1–4, 2002.
- [30] J. Fortágh, H. Ott, S. Kraft, A. Günther, and C. Zimmermann. Surface effects in magnetic microtraps. *Phys. Rev. A*, 66:041604/1–4, 2002.
- [31] S. Kraft, A. Günther, H. Ott, D. Wharam, C. Zimmermann, and J. Fortágh. Anomalous longitudinal magnetic fields near the surface of copper conductors. *J. Phys. B*, 35:L469–74, 2002.
- [32] M. P. A. Jones, C. J. Vale, D. Sahagun, B. V. Hall, and E. A. Hinds. Spin coupling between cold atoms and the thermal fluctuations of a metal surface. *Phys. Rev. Lett.*, 91:080401/1–4, 2003.
- [33] D. M. Harber, J. M. McGuirk, J. M. Obrecht, and E. A. Cornell. Thermally induced losses in ultra-cold atoms magnetically trapped near room-temperature surfaces. *J. Low Temp. Phys.*, 133:229–38, 2003.
- [34] C. Henkel, S. Pötting, and M. Wilkens. Loss and heating of particles in small and noisy traps. *App. Phys. B*, 69:379–87, 1999.
- [35] C. Henkel and S. Pötting. Coherent transport of matter waves. *App. Phys. B*, 72:73–80, 2001.
- [36] Y.-J. Lin, I. Teper, C. Chin, and V. Vuletić. Impact of the Casimir-Polder potential and Johnson noise on Bose-Einstein condensate stability near surfaces. *Phys. Rev. Lett.*, 92:050404/1–4, 2004.

- [37] H. B. G. Casimir and D. Polder. The influence of retardation on the London-van der Waals forces. *Physical Review*, 73:360–72, 1948.
- [38] C. I. Sukenik, M. G. Boshier, D. Cho, V. Sandoghar, and E. A. Hinds. Measurement of the Casimir-Polder force. *Phys. Rev. Lett.*, 70:560–3, 1993.
- [39] P. Horak, B. G. Klappauf, A. Haase, R. Folman, J. Schmiedmayer, P. Domokos, and E. A. Hinds. Possibility of single-atom detection on a chip. *Phys. Rev. A*, 67:043806/1–9, 2003.
- [40] R. Long, T. Steinmetz, P. Hommelhoff, W. Hänsel, T. W. Hänsch, and J. Reichel. Magnetic microchip traps and single-atom detection. *Phil. Trans. R. Soc. Lond. A*, 361:1375–89, 2003.
- [41] I. Teper, Y.-J. Lin, and V. Vuletić. Cavity-aided single atom detection in a microfabricated chip. *Phys. Rev. Lett.*, 97:023002/1–4, 2006.
- [42] E. M. Purcell. Spontaneous emission probabilities at radio frequencies. *Phys. Rev.*, 69:681, 1946.
- [43] C. E. Wieman and L. Hollberg. Using diode lasers for atomic physics. *Rev. Sci. Instrum.*, 20:1–20, 1991.
- [44] H. Loh, Y.-J. Lin, I. Teper, M. Cetina, J. Simon, J. K. Thompson, and V. Vuletić. Influence of grating parameters on the linewidths of external-cavity diode lasers. (to be published in *App. Opt.*).
- [45] C. Wieman and T.W. Hänsch. Doppler-free laser polarization spectroscopy. *Phys. Rev. Lett.*, 36:1170–3, 1976.
- [46] D. A. Steck. Rubidium 87 D line data. <http://steck.us/alkalidata/>, 2003.
- [47] *CRC Handbook of Chemistry and Physics*. CRC Press, Boca Raton, FL, 86th edition, 2005-2006.
- [48] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. Behavior of atoms in a compressed magnetic-optical trap. *J. Opt. Soc. Am. B*, 11:1332–5, 1994.

- [49] W. Ketterle and N.J. van Druten. Evaporative cooling of trapped atoms. *Adv. Atom. Mol. Opt. Phys.*, 37:181–236, 1996.
- [50] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. *Phys. Rev. A*, 53:381–9, 1996.
- [51] J. M. McGuirk, D. M. Harber, J. M. Obbrecht, and E. A. Cornell. Alkali-metal adsorbate polarization on conducting and insulating surfaces probed with Bose-Einstein condensates. *Phys. Rev. A*, 69:1–6, 2004.
- [52] S. Scheel, P. K. Rekdal, P. L. Knight, and E. A. Hinds. Atomic spin decoherence near conducting and superconducting films. *Phys. Rev. A*, 72:042901/1–4, 2005.
- [53] J. Reichel, W. Hänsel, and T. W. Hänsch. Atomic micromanipulation with magnetic surface traps. *Phys. Rev. Lett.*, 83:3398–401, 1999.
- [54] Z.-C. Yan, A. Dalgarno, and J. F. Babb. Long-range interactions of lithium atoms. *Phys. Rev. A*, 55:2882–7, 1997.
- [55] E. L. Surkov, J. T. M. Walraven, and G. V. Shlyapnikov. Collisionless motion and evaporative cooling of atoms in magnetic traps. *Phys. Rev. A*, 53:3403–8, 1996.
- [56] T. A. Pasquini, Y. Shin, C. Sanner, M. Saba, A. Schirotzek, D. E. Pritchard, and W. Ketterle. Quantum reflection from a solid surface at normal incidence. *Phys. Rev. Lett.*, 93:223201/1–4, 2004.
- [57] J.-Y. Courtois, J.-M. Courty, and J. C. Mertz. Internal dynamics of multilevel atoms near a vacuum-dielectric interface. *Phys. Rev. A*, 53:1862–78, 1996.
- [58] D. M. Harber, J. M. Obbrecht, J. M. McGuirk, and E. A. Cornell. Measurement of the Casimir-Polder force through center-of-mass oscillations of a Bose-Einstein condensate. *Phys. Rev. A*, 72:033610/1–6, 2005.
- [59] P. Goy, J. M. Raimond, M. Gross, and S. Haroche. Observation of cavity-enhanced single-atom spontaneous emission. *Phys. Rev. Lett.*, 50:1903–6, 1983.

- [60] D. J. Heinzen, J. J. Childs, J. E. Thomas, and M. S. Feld. Enhanced and inhibited visible spontaneous emission by atoms in a confocal resonator. *Phys. Rev. Lett.*, 58:1320–3, 1987.
- [61] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414:413–8, 2001.
- [62] J. Reichel and J. H. Thywissen. Using magnetic chip traps to study Tonks-Girardeau quantum gases. *J. Phys. IV*, 116:265–74, 2004.
- [63] A. Öttl, S. Ritter, M. Kohl, and T. Esslinger. Correlations and counting statistics of an atom laser. *Phys. Rev. Lett.*, 95:090404/1–4, 2005.
- [64] P. Münstermann, T. Fischer, P. W. H. Pinkse, and G. Rempe. Single slow atoms from an atomic fountain observed in a high-finesse optical cavity. *Opt. Comm.*, 159:63–7, 1999.
- [65] J. McKeever, J. R. Buck, A. D. Boozer, and H. J. Kimble. Determination of the number of atoms trapped in an optical cavity. *Phys. Rev. Lett.*, 93:143601/1–4, 2004.
- [66] T. Calarco, E. A. Hinds, D. Jaksch, J. Schmiedmayer, J. I. Cirac, and P. Zoller. Quantum gates with neutral atoms: controlling collisional interactions in time-dependent traps. *Phys. Rev. A*, 61:022304/1–11, 2000.
- [67] A. E. Siegman. *Lasers*. University Science Books, Sausalito, CA, 1986.
- [68] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *App. Phys. B*, 31:97–105, 1983.
- [69] E. D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *Am. J. Phys.*, 69:79–87, 2001.

- [70] R. Loudon. *The Quantum Theory of Light*. Oxford University Press, 3rd edition, 2000.
- [71] B. R. Mollow. Power spectrum of light scattered by two-level systems. *Phys. Rev.*, 188:1969–75, 1969.
- [72] B. Dahmani, L. Hollberg, and R. Drullinger. Frequency stabilization of semiconductor lasers by resonant optical feedback. *Opt. Lett.*, 12:876–8, 1987.
- [73] J. M. Geremia, J. K. Stockton, and H. Mabuchi. Real-time quantum feedback control of atomic spin squeezing. *Science*, 304:270–3, 2004.