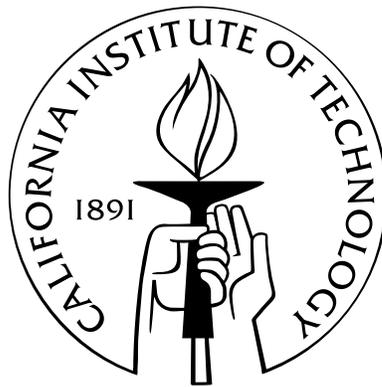


Magnetic Microtraps for Cavity QED, Bose-Einstein Condensates, and Atom Optics

Thesis by
Benjamin L. Lev

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2006

(Defended September 15, 2005)

© 2006

Benjamin L. Lev

All Rights Reserved

To my parents, Daniel and Susan

Acknowledgements

One could not ask for a better advisor: throughout my years in graduate school, Hideo Mabuchi has been a friend as well as a mentor. His passion and dedication to science breathes life into research and inspires all who are fortunate enough to work with him. He has not only taught me how to build an experiment, but also, how to think as a scientist. I hope to honor his efforts by carrying his unceasing inquisitiveness wherever I go.

In addition to Hideo, I credit Jakob Reichel with encouraging my development as a scientist. Jakob and his wife, Johanna, have been amazingly kind to me over the years, and I am grateful to both Jakob and Professor T. W. Hänsch for allowing me to participate in their experimental efforts in Munich. Tilo Steinmetz, Philipp Treutlein, Peter Hommelhoff, and Yves Colombe have been wonderful labmates and hosts, and I thank them for teaching me to be a true Münchener. I feel fortunate to have been welcomed into Oskar Painter's lab and appreciate his interest in my research and scientific development. It has been a joy to work with his students, Paul Barclay and Kartik Srinivasan.

Hideo has assembled a group of students and colleagues second to none, and I feel privileged to have had the opportunity to work with them: Michael Armen, John Au, Andrew Berglund, André Conjusteau, Andrew Doherty, JM Geremia, Salman Habib, Asa Hopkins, Chungsook Lee, Tim McGarvey, Kevin McHale, Anthony Miller, Benjamin Rahn, Jen Sokol, John Stockton, Ramon Van Handel, and Jon Williams. Grad school would not have been the same without the tomfoolery of Mike, Andy, Stockton, McKeever, and Noah Stern—truly the best partners-in-crime a guy could ever want.

Jeff Kimble and his students in the Caltech Quantum Optics group have been especially kind in their assistance throughout the years. David Boozer, Joseph Buck, Kevin Birnbaum, Jason McKeever, and Christoph Nägerl deserve special thanks. The assistance of Axel Scherer, Michael Roukes, and their students and postdocs—notably Darrell Harrington and Eyal Buks—has been crucial.

To all my LA, East Coast, and Georgenstraße friends, teachers, and family, I say thank you for encouraging me throughout the years. My mother and father have been both my inspiration and my anchor. Lastly, I want to thank my wife-to-be, Elizabeth, for her love and support.

Abstract

The system comprised of an atom strongly coupled to photons, known as cavity quantum electrodynamics (QED), provides a rich experimental setting for quantum information processing, both in the implementation of quantum logic gates and in the development of quantum networks. Moreover, studies of cavity QED will help elucidate the dynamics of continuously observed open quantum systems with quantum-limited feedback.

To achieve these goals in cavity QED, a neutral atom must be tightly confined inside a high-finesse cavity with small mode volume for long periods of time. Microfabricated wires on a substrate—known as an atom chip—can create a sufficiently high-curvature magnetic potential to trap atoms in the Lamb-Dicke regime. We have recently integrated an optical fiber Fabry-Perot cavity with such a device. The microwires allow the on-chip collection and laser cooling of neutral atoms, and allow the magnetic waveguiding of these atoms to an Ioffe trap inside the cavity mode. Magnetically trapped intracavity atoms have been detected with this cavity QED system. A similar experiment employing microdisks and photonic bandgap cavities is nearing completion. With these more exotic cavities, a robust and scalable atom-cavity chip system will deeply probe the strong coupling regime of cavity QED with magnetically trapped atoms.

Atom chips have found great success in producing and manipulating Bose-Einstein condensates and in creating novel atom optical elements. An on-chip BEC has been attained in a miniaturized system incorporating an atom chip designed for atom interferometry and for studies of Josephson effects of a BEC in a double-well potential.

Using similar microfabrication techniques, we created and demonstrated a specular magnetic atom mirror formed from a standard computer hard drive. This device, in conjunction with micron-sized charged circular pads, can produce a 1-D ring trap which may prove useful for studying Tonks gases in a ring geometry and for creating devices such as a SQUID-like system for neutral atoms.

This thesis describes the fabrication and employment of these atoms chips in experiments at both Caltech and Munich, the latter in collaboration with Professors Theodore Hänsch and Jakob Reichel at the Max Plank Institute for Quantum Optics.

Contents

Acknowledgements	iv
Abstract	vi
Preface	1
1 Introduction	21
2 Atom Chips	26
2.1 Introduction	26
2.2 Microwire traps	27
2.3 Zoology of microtraps	30
2.3.1 The wireguide, U-trap, and Z-trap	31
2.3.2 Variations: H-trap, dimple trap, P-trap, et cetera	33
2.3.3 The Libbrecht-style Ioffe trap	37
2.3.4 Permanent magnets and RF and electrostatic fields	39
2.4 Atom chip loading	41
2.4.1 Macro U-MOT	43
2.4.2 D-MOT	48
2.4.3 Sub-doppler cooling and optical pumping	49
2.4.4 D-trap and rotation	50
2.4.5 Z-trap transfer	51
2.4.6 Imaging	52
3 Experimental Details	54
3.1 Vacuum system	54
3.2 Atom chip assembly	60

3.3	Laser system	62
3.4	Computer control system	67
4	Fabrication of Micro-Magnetic Traps for Cold Neutral Atoms	72
4.1	Fabrication Challenges and Constraints	72
4.2	The elements of atom chip fabrication	74
4.2.1	The photomask	75
4.2.2	The substrate	76
4.2.3	Substrate cleaning	77
4.2.4	Thermal evaporation	78
4.2.5	Photoresist spinning and baking	79
4.2.6	UV exposure	79
4.2.7	Developing	80
4.2.8	Ozone dry stripping	81
4.2.9	Wire contacts	81
4.2.10	The mirror	81
4.3	Specific fabrication techniques: wet etching, ion milling, lift-off method, and electroplating	85
4.3.1	Wet etching and ion milling	85
4.3.2	The lift-off method	88
4.3.3	Electroplating	89
4.4	Trap fragmentation	92
4.5	Conclusion	93
5	Cold Collisions and Cesium Microtrap Losses	94
5.0.1	Collision Basics	94
5.1	Cold collision processes	96
5.1.1	Excited-state collisions: MOT loss mechanisms	96
5.1.2	MOT diffusion losses	98
5.1.3	Background losses	99
5.1.4	Ground-state loss mechanisms	100
5.1.5	Majorana spin-flip losses	102
5.2	Cesium microtrap losses	103

5.2.1	U-MOT losses	103
5.2.2	U-trap losses	120
5.2.2.1	Background collisions	120
5.2.2.2	Ground-state collisions and Majorana spin-flips	120
5.2.3	Elastic collisions	123
5.2.4	Concluding remarks	123
6	The Atom-Cavity Chip: Combining microwire traps with photonic bandgap cavities and microdisks	125
6.1	Introduction	125
6.2	Cavity QED with microcavities	129
6.2.1	Photonic bandgap cavities	130
6.2.2	Microdisks	133
6.2.3	Comparison to other cavity systems	135
6.3	Experimental proposal	137
6.4	Single atom detectability using PBG cavities	144
6.4.1	Single atom detection signal-to-noise	146
6.4.2	Simulated atom transits and cavity induced force	150
6.5	Single atom detectability using microdisk cavities	156
6.6	Conclusion	157
7	Fiber-Gap Fabry-Perot Cavity: The first atom-cavity chip	160
7.1	Introduction	160
7.2	Experimental details	162
7.2.1	Atom chip loading and waveguiding	163
7.2.2	The fiber-gap cavity	167
7.3	Signal-to-noise and spontaneous emission analysis using the master equation	172
7.3.1	Signal-to-noise and spontaneous emission	174
7.3.2	Simulated atom transits and cavity induced force	179
7.4	Experimental results	181
7.5	Outlook	188

8	Atom Mirror Etched from a Common Hard Drive	190
8.1	Overview	190
8.2	Design and fabrication of the hard drive atom mirror	191
8.3	Experimental details	195
8.4	Results	197
9	A 1-D Magnetoelectrostatic Ring Trap for Neutral Atoms	199
9.1	Overview	199
9.2	Ring trap design	200
9.3	Device electrical leads and trap perturbations	203
9.4	Fabrication	204
9.5	Trap loading and surface effects	205
9.6	Discussion	207
10	Splitting a BEC in a Magnetic Double-Well Potential: Atom interferometry and Josephson effects on an atom chip	210
10.1	The double-well chip	210
10.1.1	Josephson effects	212
10.1.2	The magnetic double-well	214
10.2	The glued-cell chamber	220
A	Magnetic Fields, Gradients, and Trap Minima of U- and Z-Traps	225
A.1	Infinitely thin linear wires	226
A.2	Force on an atom in a U-trap	227
A.3	U-trap minimum and gradients	228
B	Matlab Code	234
B.1	Cavity QED: transmission, signal-to-noise, and spontaneous emission	234
B.2	Master equation for a two-level atom	238
B.3	Magnetic field, gradient, and curvature for U- and Z-traps and waveguides	239
B.4	Computer control code for the atom chip experiment	245
	Bibliography	259

List of Figures

1.1	Quantum feedback in a cavity QED setting	22
1.2	Quantum network scheme based on cavity QED	23
2.1	Macroscopic coil-based magnetic traps	28
2.2	The Lamb-Dicke regime	31
2.3	Magnetic microtrap building-blocks	32
2.4	The H-trap, dimple trap, and P-trap	33
2.5	The Libbrecht-style Ioffe trap	38
2.6	Variation of the 3-D planar Ioffe trap and a permanent magnet Ioffe trap . .	40
2.7	An example of an atom chip loading procedure	42
2.8	The mirror MOT	45
3.1	Vacuum chamber	55
3.2	Atom chip assembly	60
3.3	Laser system schematic	62
4.1	A fabricated Libbrecht-style Ioffe trap	74
4.2	Polyimide-coated waveguide atom chips	84
4.3	Fabrication techniques we use to make atom chips	86
4.4	Gold wire patterned using the wet etch technique	87
4.5	An atom chip-based BEC interferometer	92
5.1	Excited-state collisions	97
5.2	MOT steady-state atom number versus background pressure	101
5.3	U-MOT decay upon compression	105
5.4	U-trap coordinate system	106
5.5	B_{bias} response to a $\tau = 2$ ms linear current ramp	109

5.6	Experimental data of the U-MOT distance above substrate versus time . . .	110
5.7	y_0 versus time	110
5.8	U-MOT distance above substrate versus time	111
5.9	Calculated distance and velocity of an atom at the U-MOT center	111
5.10	Loading and collision terms in Equation 5.23 for a low-gradient trap	113
5.11	Loading and collision terms in Equation 5.23 for a high-gradient trap	114
5.12	The reduced χ^2 without the $R(t)$ loading term in Equation 5.23	114
5.13	Typical fits to U-MOT decay curves	116
5.14	U-MOT loss rate, β , versus trap gradient	116
5.15	V_0 versus trap gradient	117
5.16	Covariance of β and V_0 at a confidence interval of 68.3%	118
5.17	V_0 with β fixed at $12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$	118
5.18	The reduced χ^2 for the V_0 fits in which $\beta = 12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$	119
5.19	RMS of $\chi^2 - 1$ for various V_0 fits in which β is fixed	119
6.1	Model system for neutral atom cavity QED	126
6.2	Photonic bandgap cavities	127
6.3	SEM of microdisks	128
6.4	Schematic of a photonic bandgap cavity	129
6.5	The PBG cavity with magnetic microtrap	130
6.6	SEM images of PBG cavity and waveguide plus fiber taper	132
6.7	Schematic of the fiber taper coupling system	133
6.8	Schematic of the fiber taper-to-microdisk coupling scheme	134
6.9	Fiber taper coupled to AlGaAs microdisk cavity	134
6.10	Schematic of the atom-cavity chip experiment	140
6.11	The atom-cavity chip to be used in the experiment	140
6.12	The fiber taper-to-atom chip alignment set-up	143
6.13	Jaynes-Cummings ladder of atom-cavity eigenstates for $[\Delta, \theta] = 0$	146
6.14	PBG cavity transmission as a function of various detunings	148
6.15	On-resonance PBG cavity transmission, signal, and signal-to-noise	151
6.16	Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [9.2, 0]$ GHz.	151
6.17	Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [9.2, 9.2]$ GHz.	152

6.18	Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [16, 0]$ GHz.	152
6.19	Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [16, 16]$ GHz.	153
6.20	Simulated photon counts due to atom transits through the cavity	154
6.21	Same as Figure 6.20 but for different detunings	154
6.22	Force on an atom traversing the PBG cavity mode	155
6.23	Microdisk cavity transmission as a function of various cavity detunings	157
6.24	On-resonance microdisk cavity transmission, signal, and signal-to-noise	158
6.25	Same as Figure 6.24 except $[\Delta, \theta]/2\pi = [1.5, 0]$ GHz.	158
6.26	Same as Figure 6.24 except $[\Delta, \theta]/2\pi = [1.5, 1.5]$ GHz.	159
7.1	Fabry-Perot atom-cavity chip using commercial supermirrors	161
7.2	Schematic of the fiber-gap cavity atom chip experiment	162
7.3	The vacuum cell assembly	164
7.4	Fiber-gap Fabry-Perot cavity and atom chip assembly	164
7.5	Close-up view of the fiber-gap cavity atom chip	166
7.6	Magnetic field plot of a waveguide plus dimple	168
7.7	Image of the fiber-gap Fabry-Perot cavity	168
7.8	The fiber mirror	170
7.9	Schematic of the fiber-gap cavity set-up	170
7.10	On-resonance fiber-gap cavity transmission, signal, and signal-to-noise	174
7.11	Fiber-gap cavity transmission as a function of various detunings	176
7.12	Fiber-gap cavity off-resonant signal-to-noise	177
7.13	Fiber-gap cavity signal-to-noise for various detunings	177
7.14	Transmission, signal, and signal-to-noise for higher finesse fiber-gap cavity	178
7.15	Force on an atom traversing the fiber-gap cavity mode	179
7.16	Simulated photon counts due to an atom transit through the cavity	180
7.17	Image of waveguided ^{87}Rb atoms traversing the fiber-gap cavity	181
7.18	Data from first atom-cavity chip	183
7.19	Atom transits through the fiber-gap cavity for different drive powers	183
7.20	Single shot measurement of atom transits in fiber-gap cavity	185
7.21	Fiber-gap cavity uptransit signals	186
7.22	Possible optical-bistability signal	187

7.23	Cavity signal post-polarization optimization	188
8.1	Cross-section of the etched hard drive	192
8.2	Relative size of the hard drive sliver used in this experiment	192
8.3	Example of etched hard drive	193
8.4	AFM and MFM images of the etched hard drive	194
8.5	Atom mirror mounted inside vacuum chamber	195
8.6	Diagram of the atom mirror experimental set-up	196
8.7	Atom mirror reflection data	197
9.1	The 1-D ring trap potential	202
9.2	Schematic of the magnetoelectrostatic ring trap	204
10.1	Microwire layout for producing a magnetic double-well potential	211
10.2	Potential on the wires of the double-well atom chip	215
10.3	Currents flowing through the wires of the double-well atom chip	216
10.4	Magnetic field in the \hat{x} - \hat{z} plane of the double-well trap	217
10.5	Magnetic field in the \hat{x} - \hat{y} plane of the double-well trap	217
10.6	Magnetic field in the \hat{y} - \hat{z} plane of the double-well trap	218
10.7	Glued-cell vacuum chamber for atom chip BEC production	220
10.8	Atom chip and base chip assembly	222
A.1	Difference between y_0 and $y_{0\text{ approx}}$	229
A.2	Log plot of z_0 as a function of η	230
A.3	The U-trap gradient in \hat{x} at the trap minimum	231
A.4	The U-trap gradient in \hat{y} at the trap minimum	231
A.5	The absolute value of the U-trap gradient in \hat{z} at the trap minimum	232
A.6	The trap gradient in \hat{y} : y_0 versus the simpler $y_{0\text{ approx}} = 2\pi\eta$	232
A.7	The U-trap gradient in \hat{x} plotted for various values of L	233
A.8	Anisotropy of the U-trap in \hat{y} versus \hat{x} as a function of trap height	233

List of Tables

2.1	Schedule of magnetic and optical fields for atom chip loading and trapping	43
5.1	Majorana spin-flip and ground state losses for U- and Z-traps	121
6.1	Comparison of cavity designs	138
6.2	Comparison of S/\mathcal{N} for various cavity detunings	150
9.1	1-D ring trap parameters for Cs ¹³³ and Rb ⁸⁷	202

Preface

Rearing the Caltech Atom Chip: a chronicle spanning six years and two countries

I consider myself fortunate to have had the privilege of working with many talented physicists throughout my graduate school career at Caltech. In addition to my thesis advisor, Professor Hideo Mabuchi, two scientists in particular have profoundly influenced my research and approach to science: Professor Jakob Reichel jointly of the Max Planck Institute for Quantum Optics (MPQ) in Garching, Germany and the Ludwig-Maximilians-Universität (LMU) in Munich, who is now a professor at the Laboratoire Kastler Brossel de l'École Normale Supérieure (ENS) in Paris; and Professor Oskar Painter at Caltech. This preface relates the history of my graduate research, beginning with the inception of Hideo's lab and the construction of my atom trapping experiments in the first half of my graduate career, and continuing through the collaborations with Jakob and Oskar's groups in the latter half.

In the beginning, there was Hideo...

My six years in what we fondly call *MabuchiLab* began in the spring of 1999 as I was finishing college and visiting prospective grad schools. By a stroke of good fortune, I happened upon Hideo while visiting the group of Professor Jeff Kimble at Caltech. Hideo and I soon found ourselves at the Athenaeum chatting over a pitcher of beer, and he began to tell me about the excitement brewing in the fields of quantum computation, cavity QED, quantum optics, and quantum feedback control. These were subjects of which I had largely been unaware and had not considered for my doctorate, but his enthusiasm left me eagerly wanting to learn more.

Roughly a year before, Hideo had finished his doctorate in Jeff's group and had been hired as an assistant professor at Caltech. He was spending time—during what happened

to be my senior year—at Princeton as a visiting fellow in their chemistry department before starting his research program at Caltech. We met a couple of times over the weeks subsequent to my Caltech visit, and through the course of our discussions his exceptional teaching ability and fresh ideas became evident. I had wanted to join a condensed matter physics group, but his proposed research program was unlike any that I had encountered and was incredibly unique and exciting. It combined the applied with the fundamental, experiment with theory, all motivated by table-top experiments. Hideo was seeking to hire his first students and I wanted something more adventurous than the standard graduate student experience. The excitement of working with a newly-minted professor, the building of a fresh lab, and his brand-new research direction was just what I wanted. I signed-on.

MabuchiLab: the first year

In mid-summer 1999, I moved to Pasadena and began work in Hideo’s new lab. Actually, his lab space wasn’t fully ready to be occupied until the end of the summer, but in the early fall the laser tables were installed and Michael Armen—who had also joined Hideo’s lab as a graduate student in the spring—and I began to fill the space with tables, computers, nuts and bolts, etc. Our first experiment would be to magnetically manipulate cold atoms for the purpose of learning how to controllably trap them inside a high finesse, low mode volume cavity such as the exotic photonic band-gap (PBG) cavities [1]. The long-term goal of the research was—and still is—to develop a cavity QED device to function as the hardware of quantum communication network. In addition, such a device would allow the experimental exploration of the dynamics of a quantum system under continuous measurement by incorporating quantum-limited feedback control.

Hideo recognized that the newly invented atom chip could solve the problem of how to robustly trap atoms inside a microcavity. It was our task to build such a device and use it to manipulate atoms, and the first step was learning how to cool and trap atoms near the room temperature surface of the atom chip. Fortunately, in the late-summer of 1999 the group of Jakob and Prof. Theodore Hänsch at the MPQ/LMU demonstrated how to do just such a thing [2]. By using the atom chip as a mirror to form two of the six MOT beams via reflection, they could move a cloud of cold rubidium atoms anywhere from millimeters to microns above a surface. Hideo, Mike, and I set-out to build, as quickly as possible, a MOT that utilized a plain mirror in such a fashion.

For the mirror MOT, we used our new Coherent MBR 110 Ti:Sapphire for the trapping

lasers and an old diode laser from Prof. Libbrecht's group as the repumper. Mike concentrated on the vacuum chamber and laser electronics, while I worked on the fabrication of our first atom chip, magnetic coils, and the methods for loading atoms into magnetic microtraps. By the last week in January 2000 we had our first mirror MOT, albeit a small one. Our mirror MOT was one of the first and—to the best of my knowledge—continues to be the only one to use cesium.

The mirror MOT turned-out to be a fickle beast, much unlike its free-space, 6-beam cousin. The late winter and spring of 2000 was spent trying in vain to improve the trapping. Meanwhile, beginning in November 1999 I was spending most of my time in Prof. Michael Roukes' clean room learning photolithography and trying to co-opt these techniques for atom trapping. The idea of using microfabricated wires to produce trapping potentials for neutral atoms was first proposed by Weinstein and Libbrecht—at Caltech, incidentally—in 1995 [3]. Prentiss and Westervelt's groups had succeeded in fabricating the micron-sized wire patterns of Libbrecht's design, and by 1999 the first experiments utilizing microfabricated wires began to produce results [2, 4, 5, 6]. (The term “atom chip” only came into wide use after a publication by Joerg Schmiedmayer's group in 2000 [7].)

While we needed to play catch-up to the other groups in terms of atom chip technology development, most of the groups were striving to create a BEC on a chip, and were not pursuing cavity QED. Of course, this has changed and there are now several groups interested in magnetically coupling atoms to a cavity. At the time, all of the atom chip groups to some extent out-sourced their atom chip fabrication, either by buying them from a company, or by having a collaboration with a dedicated fabrication group at the same or even a remote university. Hideo felt that, in contrast, the best way for us to come up to speed would be to have the same student (me) do both the fabrication and atom trapping simultaneously. This would allow a synergy to emerge whereby complications and advantages in both experimental disciplines could be dealt with or harnessed efficiently. To a large extent, I believe the strategy has worked-out as intended.

By April 2000, I had made our first batch of atom chips. Although the techniques of photolithography are well known, the constraints imposed by atom trapping require several modifications to the standard procedures that, when taken together, produce a formidable task. It was during this five-month period that I worked-out the basics of atom chip fabrication. Though it was grueling, it provided me with an opportunity to learn various

tricks that I later compiled for an atom chip fabrication tutorial [8] and have employed over the years for both our chips and those of Jakob’s group at the MPQ/LMU. Of course, I wasn’t alone in the Roukes’ clean room: Darrell Harrington and Eyal Buks—grad student and post-doc in Roukes’ group, respectively—were generous with their advice. These first chips were made with the lift-off technique. It’s not the easiest to begin with—and I didn’t know any better at the time—but can produce very thin features. Libbrecht-style trap patterns were made from gold thermally evaporated on large sapphire substrates, with some devices having wire cross-sections of one-by-one micron.

Mike and I visited Salman Habib and Tanmoy Bhattacharya at the Los Alamos National Laboratories for three weeks in June 2000. Hideo and Salman have been collaborating since Hideo was in grad school, and we spent our time learning about atom trapping and cooling techniques. In the late summer, we realized that our mirror MOT could be greatly improved with a different arrangement of the quadrupole coils, and I began to re-design the trapping apparatus. In the spring, MabuchiLab had acquired a new grad student, John Au who began work on an adaptive phase measurement experiment [9], and in late summer Mike joined that project. Since I already had atom chip fabrication experience, I took-over the atom chip trapping experiment and ran it alone until the summer of 2005 when I began to hand it over to Oskar Painter’s grad student Paul Barclay.

Building and rebuilding

You need to know how everything works to be a good experimentalist, but to really learn, you must either break it or build it from scratch. From fall 2000 to November 2001, I was figuring-out for myself the do’s and don’ts of atom trapping. In sum, all aspects of the experimental apparatus needed redesigning and construction. Over these months, I rebuilt the Cs oven and integrated a Pfeiffer turbo-pump permanently into the set-up; redesigned and rebuilt the mirror MOT coils and the bias magnetic field coils; rearranged the layout of the vacuum chamber; integrated a new diode laser (of the latest plexiglas Kimble group design that Mike had largely assembled) into the system; began to develop a computer system for control of experimental timing, CCD camera triggering, and data acquisition; and learned how to computer control and fast-switch power-supplies for the magnetic coils and microwires.

The new mirror MOT was running by January 2001, using the atom chip with the Libbrecht-style Ioffe microtraps. The goal was to learn how to load atoms into these high-

curvature traps for eventual loading into photonic bandgap cavities [3]. The central difficulty was in transferring atoms from the mirror MOT millimeters above the surface to a small-volume, purely magnetic trap just 10 microns above the micron-sized Ioffe trap. No group had attempted to load a Libbrecht-style Ioffe trap, and to the best of my knowledge, no one has yet. I first tried loading the atoms from the mirror MOT to a quadrupole trap formed from one of the rings on the substrate, but this didn't work well since the atoms couldn't follow the changing orientation of the quadrupole field. Moreover, they were too hot because the cooling lasers need to be extinguished during the field rotation. I found my laser, imaging, and computer control system to be inadequate, and realized that I was trying to reinvent the wheel on too many fronts.

In August 2001, I tore-down the system and began to rebuild once again, incorporating the lessons learned in my previous attempt. The key simplification was to not reinvent magnetic trapping with microwires. Jakob's group had already found a great method for doing this by combining the field from a U-shaped wire and a bias magnetic field to form a cigar-shaped quadruple trap. The great advantage over my ring-based quadrupole trap was that it automatically produced the correct field orientation for forming a mirror MOT: an externally generated quadrupole field could be used to collect a large number of atoms, then without changing any field orientation, this quadrupole could be smoothly replaced by that of the U-trap while never needing to extinguish the cooling lasers. Once the atoms are comfortably in the U-MOT, the lasers could be turned-off leaving the atoms in a magnetic microtrap. I decided to switch to the U-trap since it was sure to work and would help me narrow-down which other aspects of the experiment needed fine-tuning.

Putting the U-MOT to use

I made our first U-MOT by late November 2001. The atom chip had a 300 μm wide, 1 μm tall U-wire that I made using a different fabrication technique. Since the device feature-sizes weren't too small, I used a transparency mask and the wet etching technique which allowed very quick design-to-finish turn-around times. A positive-process-with-negative-photoresist photolithography technique was employed and was a great simplification over lift-off. Loading the U-MOT from the mirror MOT worked as planned, and I spent the next six months improving all aspects of the experiment by using the U-MOT loading efficiency, stability, lifetime, atom population, and image quality as a benchmark.

The computer system for controlling the lasers, magnetic field coils, and CCD cameras

was not adequate and was based largely on programming pulse delay generators. I began using LabView combined with some Matlab code to make a hybrid system for image acquisition and processing, for controlling DAQ boards to provide analog outputs to the coil and microwire power supplies, and for controlling digital outputs for various triggers to the remaining pulse delay generators. It still wasn't perfect, but more refinements wouldn't come until late 2003 with the use of Matlab as the primary experimental sequence scheduler (see Chapter 3).

The laser system was the most significant thing to be taken-apart and revamped. By December 2001, the Ti:Sapphire laser was becoming more heavily used by others in our lab and I wanted to replace it with my own diode laser. Moreover, the plexiglas diode laser was practically unusable due to its mechanical instability. Over the past year I had spent a good amount of time reworking it, but it seemed to have an inherent affinity to mode-hop nearly every other day which required a re-optimization of the grating each time. I didn't realize that I shouldn't have to put-up with this, but by the end of 2001 I decided that a new laser was required. Fortunately, a few months before, Christoph Nägerl in Kimble's lab finished making a brand-new diode laser design that was the ultimate in stability. I built one by mid-January and it worked like a charm as my master trapping laser. The repumper was built with JM Geremia who wanted to learn how to build diode lasers for an experiment he and John Stockton were designing.

I noticed that the mirror MOT was extremely sensitive to intensity and pointing fluctuations and to imperfections in the laser beam profiles. To ameliorate these problems, I found it best to put the laser through a fiber to filter it into a nice Gaussian mode. It also became apparent that one needed polarization-maintaining fiber to avoid polarization noise being mapped onto the intensity. This was all incorporated, but due to the power reduction from the various AOM's and the fiber, I didn't have enough MOT beam power for the experiment. The old plexiglas laser came in handy, and I turned it into a slave laser (with advice from Dave Boozer) by injection-locking it with the master diode laser. The slave diode laser provided enough power to go through the AOM's and fiber while leaving ample additional power from the master laser for an absorption imaging beam and an optical pumping beam. This laser set-up was largely completed by March 2002. It remains as a nice turn-key element of my current experimental apparatus, and a similar system has been adopted in John and JM's experiment.

In June 2002, I put all the computer, laser, and imaging control systems together to polarization gradient (PG) cool the atoms above a surface, which is more difficult than in the free-space case. In free-space MOTs, PG cooling is relatively straightforward because neither the laser beams nor the imaging access are obstructed. One can zero the magnetic fields by watching how the atom cloud expands as the magnetic field is suddenly extinguished. If the cloud expands asymmetrically, and the opposing laser beams are well matched in intensity, then one knows that the fields are non-zeroed in the direction of the errant expansion. However, in a mirror MOT configuration the situation is a bit more complicated. The mirror is not perfect, which poses five problems: 1) If the microwires imprint defects on the mirror, then these defects will create shadows in the 45 degree MOT beams and in effect dice-up the regions in space that can trap a MOT; 2) The mirror might have defects—either from dust or from the microwires—that obscure the image of the atoms by scattering unwanted light into the CCD camera. This makes it difficult to take careful images of the atom cloud; 3) The horizontal laser beam is partially blocked and scattered by the mirror creating a similar problem as in (2); 4) In some experiments, the mirror might not be 100% reflective, causing the reflected 45 degree MOT beams to be mismatched in intensity from their counterpropagating beams; 5) The cloud quickly crashes into the mirror surface if the mirror is not positioned upside-down. Problem (1) is most severe, in that it hampers one’s ability to zero the magnetic fields by causing the MOT to fragment and not spatially translate in a linear fashion. This prevents one from easily observing in which direction the magnetic field is not zeroed. By careful construction of the mirror and the use of multiple viewing angles, one can overcome these difficulties even with a right-side-up mirror that is not perfectly reflecting. In these early experiments I was able to cool the atoms to between 3 and 10 μK .

With the addition of PG cooling to the toolbox, I was immediately able to load a magnetostatic U-trap from the U-MOT. A problem arose, however, in the trap loading efficiency. My mirror MOT started-out at around 10^6 Cs atoms, but the loading efficiency was so poor that only 1% or so made it into the U-trap. At that time, the detection sensitivity wasn’t that great, and these 10^4 atoms could barely be resolved. Hideo and I weren’t sure whether it was a flawed loading procedure or whether there was a fundamental problem with trapping Cs in such a manner. We were aware that many groups had had trouble trapping Cs for BEC production, and we were concerned that similar problems

would plague us. Since our group didn't have a great understanding of cold collisions, I investigated this experimentally with my U-MOT system and wrote some review-like notes that are Chapter 5 in this thesis. I made measurements of the atom loss for the U-MOT in August of 2002 and spent the next two months primarily working on the analysis. The upshot is that Cs is ~ 100 times more sensitive to collisions than other commonly trapped atoms, like Rb, and that the loading inefficiency was primarily due to loss from excited-state collisions. This means that while my loading procedure would have worked fine for Rb, it was not optimal for Cs, and that I'd need to minimize the excited-state collisions by reducing the time the atoms spent in a compressed U-MOT. This led to a redesign of my experiment in favor of using a macro U-MOT to trap atoms from vapor directly. By the way, the reason we use Cs as opposed to more benign atoms such as Rb is largely historical. While Cs is heavier (good for trapping in the Lamb-Dicke regime) and has a smaller scattering rate (good for cavity QED) than the other trappable alkali's, I believe I've been told that the reason Kimble's group began using Cs is that mirror coatings in the late 80's were much better at Cs's 852 nm wavelength than at Rb's 780 nm, which was crucial for obtaining high finesse cavities for cavity QED. (I believe it is also the case that the non-linear crystals used for creating squeezed light operate better at 852 nm than at shorter wavelengths.) Since that group is interested in single atoms, not BECs, collisional loss isn't such a big problem and Cs remains in use to this day, and we continued in this tradition.

Bouncing atoms

While current-carrying wires generate heat and require an electrical input/output connection, permanent magnets are completely passive, miniaturized, and can generate similarly large magnetic field curvatures and gradients. When I arrived in Hideo's group, a Caltech undergraduate by the name of Clifford Hicks was working on a summer research project to design permanent magnet geometries for Ioffe traps in the Lamb-Dicke regime. He came-up with some promising designs, but we didn't know how to actually make permanent magnets of the required size and shape. It became my side-project to investigate fabrication strategies.

In April 2001, Hideo and I decided that the first thing to do would be to start using currently available miniature permanent magnets to build something useful for atom manipulation. This would give us an opportunity to learn more about their fabrication and

operation that could be applied to future experiments. We settled on making a permanent magnetic atom mirror for our initial atom optics project. The types of atom mirrors that had recently been developed by the group of Ed Hinds used magnetic material from floppy disks and video tape [10]. This group wrote sinusoidal patterns of magnetization by programming the drive head, but were limited to around 10 to 12 μm periodicity. The smaller the periodicity, the sharper the turn-on of the repulsive magnetic field barrier. We recognized that a hard drive contains much smaller magnetic granules with higher coercivities, allowing much smaller magnetic modulation. Moreover, the remanent magnetic field would be much higher, enabling the reflection of higher energy particles. Following Hinds' method of writing-in the magnetic pattern wouldn't teach us about permanent magnetic fabrication, so we decided to investigate patterning thin magnetic films.

Around this time, Yves Lassailly—a visiting researcher from France—had come to work in Axel Scherer's group, but found himself interested in learning about atom trapping. We co-opted him to help on this project since his background was in fabrication. Our first idea was to deposit cobalt alloy in thin, micron-sized strips defined by photoresist. The stripes would then be magnetized perpendicular to their axis and parallel to the substrate's plane. In a discussion with Mladen Barbic, a postdoc with Axel, we became convinced that the magnetic domains would not line-up correctly and abandoned this idea. (Subsequently, we found-out that the group of Peter Hannaford was doing this exact thing and successfully made an atom mirror with it [11].) Yves and I went back to the hard drive as the source of a magnetic thin film since it automatically had nice magnetic granules that could be oriented in in-plane stripes (these are how bits are formed). Through much effort on the part of Yves, by November 2001 we had made nice micron-sized photolithography patterns on the hard drive substrate. In July 2002, we got the ion etching of the hard drive to work and MFM scans confirmed that we had made nice magnetic strips. There were still a few more bugs to work-out, but Yves had to return to France in the fall. I made a few of the hard drive devices for optimization myself, but the last few ion etching runs were made by our new postdoc, Chungsook Lee, whom we wanted to involve in the project. By November, I had a nice device, but was too busy with other projects to put it into the vacuum chamber until early February 2002. The first attempt to bounce atoms off the hard drive failed because during the chamber bake-out the temperature had risen too close to the Curie temperature of the magnetic granules, erasing the pattern. The atoms would not

bounce, and sure enough, an MFM scan after the experiment revealed the missing pattern. Chungsook and I re-magnetized the sample, and I returned it to my vacuum chamber, this time without baking during pumped-down. The etched hard drive surface was a lousy mirror (see Chapter 8 for details), but I managed to form a mirror MOT above the surface and PG cool the atoms to $\sim 10 \mu\text{K}$. It was a simple matter to drop the atoms and capture images of them falling and bouncing a few times: The hard drive atom mirror was born on March 6, 2003 [12].

A racetrack for atoms

The atom mirror project has grown to be much more than a testbed for permanent magnet fabrication. We believe that the hard drive atom mirror has many potential applications beyond the simple reflection of thermal atoms or matter waves. After the completion of this project, Hideo and I began thinking of what to do next with this technology. This atom mirror could be useful for creating atom optical devices such as corner cubes and tightly-confining waveguides for neutral atoms or neutrons. Another route would be to explore the use of microfabricated wires or charged-pads on the surface to provide time-dependent fields that would work in conjunction with the high gradient, high field atom mirror potential to produce new devices. Sometime in mid-2002, I met Tomasso Calarco, a theorist at Innsbruck and Trento, who told me he had been thinking of ways to use a magnetic atom mirror's potential in combination with electric fields to perform quantum logic gates [13]. In the spring of 2003, Hideo and I decided that learning how to trap atoms with the combination of electric fields from charged pads and the magnetic field from our atom mirror would be a useful first step toward such an end.

Asa Hopkins, a new graduate student in our group, joined me on this project in the summer of 2002. We first looked at using the motion of the atom mirror to stop an impinging cloud of atoms. However, it seemed unlikely that we'd find a way to move the hard drive substrate fast enough. We then investigated the possibility of using the atom mirror as a device to create a 2-D gas as Hinds and colleagues had proposed [14]. Unfortunately, the corrugation of the magnetic field emanating from the patterned hard drive surface was too large, preventing the trapping of atoms. In February of 2004, I met a Seagate engineer at a conference who told me he could provide us with hard drive platters with ten times the remanent magnetic field [15]. I believe that atom mirrors made of these hard drives would not exhibit the corrugation problem and the 2-D gas experiment could work with these

materials.

Finally, we went back to looking at what kind of trap we could make with charged pads on the mirror surface. It quickly became clear that the magnetic repulsion of the etched hard drive could be balanced by the attractive force on the atoms due to the electric field. Specifically, a charged disk on the atom mirror could create a trap from the circular ring of potential minima suspended above its edges, and Asa calculated that the curvature of the trap could be as large as 100 kHz. This got me quite excited, as I had just been reading about 1-D Tonks gases and recognized that this magnetoelectrostatic trap could force a BEC of alkali atoms to be in the 1-D Tonks gas regime in a ring geometry. Definitely a novel device! Moreover, by perturbing the ring potential with underlying charged-wires, we could create Josephson-like junctions for the confined matter wave. This, I believe, would be akin to a SQUID for neutral atoms, and I became even more excited. We spent some time working out various schemes for loading the ring trap with atoms and testing its robustness to perturbations. We wrote-up the result as a Phys. Rev. A in January and February 2004. Asa will take-over the hard drive atom mirror project when I leave, and is currently attempting to fabricate the magnetoelectrostatic trap and build a suitable experimental set-up (see Chapter 9)

To Munich, young man

One of the most fortuitous encounters in my life occurred in late December 2001. On an overcast Sunday morning, Philip Grangier—a professor at CNRS in France—and Jakob Reichel stopped-by our lab. They had been to a conference in San Francisco (I believe), and I was expected to give a lab tour to only Prof. Grangier. Jakob was a pleasant surprise. I had read all of his papers on microwire trapping, but had never met him and was eager to learn more details from him. Likewise, I believe Jakob was surprised. I don't believe he knew there was an atom chip project at Caltech, let alone one that was fabricating chips in-house and using his U-trap method. What was intended as a quick lab tour stretched-on for hours as we exchanged experimental tips and thoughts about future directions.

We realized that an extremely useful collaboration could be formed between our groups. They were experts at producing BECs on a chip with all its difficult technical details, but were forced to buy their atom chips from fabrication companies—a situation that hampered design creativity since these companies were not willing to invest as much effort into pushing the technology as would a graduate student. I had spent a lot of time in creating an in-

house atom chip fabrication capability, and was greatly interested in learning how to make a BEC on a chip. The deal was hatched: I'd make their next generation BEC on a chip device, and they'd have me visit his and Prof. Hänsch's labs in Munich for an extended stay. At the time I just thought this would be a one-time exchange. To my amazement and great pleasure this turned into a full-fledged collaboration involving six trips to Munich and lasting three years and counting... They have kindly treated me as if I were a full member of their group.

The groups in Innsbruck and Trento invited me to give talks on my research in April 2002, and Jakob and I decided that this would provide us with a good chance to discuss prospective projects in Munich. I spent about a week in Munich meeting with his students, touring his labs and those at the MPQ campus in Garching, and discussing various ideas for a joint experiment. Over the following summer, we exchanged a few more ideas, and in the fall, Jakob and his students, Tilo Steinmetz and Peter Hommelhoff made the final proposal. Peter had been working on separating a BEC into two clouds with microwires and letting them recombine and interfere. A device such as this would be useful for atom interferometry. The more closely spaced and smaller the wire pattern, the smaller the separation between the BEC clouds and the more pronounced the interference [16]. Unfortunately, their current chip's wire dimensions were too large, and it would be my goal to fabricate a new device with several parallel wires no larger than 2 to 3 microns with 2 micron spacings and several microns tall to allow a significant current to flow. Making such a double-well device was a considerable challenge. They sent me the mask design in October of 2002, and by the end of November I had worked-out the fabrication method. I used the technique of electroplating with a positive photoresist to make the narrow wires 4 microns tall and patterned on the AlN substrates they provided. A lift-off plus electroplating technique worked well also, but was more complicated. I made ten chips of various designs by early December, and in mid-January 2003 flew to Munich for three and a half weeks.

Upon arrival, Tilo and I immediately began assembling the vacuum chamber. We attached the atom chip to a base chip, and this assembly glued-on as one face of the chamber's glass cell. Within two weeks we had pumped-down the chamber to 2×10^{-10} Torr and had finished making the electrical feedthrough cables. By the end of the trip, Tilo, Peter, and I had aligned the MOT optics and made our first mirror MOT in this compact chamber. It should be noted that the entire set-up (minus the laser system) is extremely compact and

is one of the smallest BEC machines in the world. The glued-cell technique originated two years or so before in a collaboration between Jakob and Prof. Dana Anderson of University of Colorado/JILA in Boulder, and Dana's group has also put a lot of effort into developing these compact systems [17].

This experiment was built on the same table as their BEC-on-a-chip experiment and shares the same laser and computer control systems (the lasers are fiber-coupled and the power supply connections can be exchanged). Philipp Treutlein, another grad student of Jakob's, and Peter occupied the laser and computer system while they completed a measurement of the coherence lifetime of a superposition of atomic internal states as a function of the atom height above the chip surface [18].

I returned to Munich again for three weeks in late May 2003, a time in which we were hopeful that the coherence experiment would be finished. In the end, the double-well experiment was put on hold until October 2003, when the coherence experiment was fully completed. I decided to spend my time during this spring trip working on calculations exploring the possibility of observing Josephson oscillations with a BEC in the double-well potential formed by our atom chip. I wrote a Matlab code—based on the equations in a paper by J. E. Williams [19]—that finds the ground-state of an azimuthally symmetric double-well potential by solving the Gross-Pitaevskii (GP) equation. With a Gaussian mode approximation, a separate code then backs-out the Josephson oscillation frequency and the atom number current amplitude.

After a comment from Markus Greiner, we realized that the current in our planar double-well trap might not flow in the ideal pattern that was originally intended and described in Reference [16]. I wrote a Laplace solver in Matlab for our exact wire pattern (with some investigation of the 3-D nature of the wires), and sure enough, the current flowed through the wires in a non-ideal manner. I then wrote a 3-D Biot-Savart solver to investigate the extent to which this current flow deviation affected the formation of a double-well potential suitable for observing Josephson effects. (I've since realized that all of this could much more easily be done in the Femlab software package.) The result was that the current deformation skewed the double-well axis away from the axis of the wire pattern, limiting the minimum achievable well spacing and junction surface area. Later, in February 2004, Philipp wrote a more versatile code based on Femlab and C for finding the 3-D field, and these calculations reproduced the skewed trap geometry. This skewing was unfortunate, and we found that

the ground-state of our double-well would only exhibit a Josephson effect slightly smaller than what we could experimentally detect. However, by this time the original device was no longer functional, and we had already begun the process of designing—using these codes—a new and improved double-well chip based on a double layer device.

Following the completion of their coherence experiment in the late fall, the laser and computer systems were moved back to the double-well BEC experiment. I arrived at the end of October for another three week stay to work with Philipp and Tilo on the experiments. In spite of the aforementioned non-idealities of the wire layout for detecting Josephson effects, we still wanted to push forward with the experiment. There was still plenty to learn with regards to the atom trapping and BEC formation with the new vacuum chamber system as well as the possibility of performing atom interferometry with this chip. During the trip, we refurbished the vacuum chamber, reinstalled the laser system, and got the absorption imaging working in two viewing axes. We completed the following trapping steps: transferred the atoms from a macro U-MOT to a P-trap, rotated the atoms 90° with the P-trap, and finally transferred the atoms to a Z-trap which formed the entrance to a waveguide meant to move the atoms to the splitting region of the chip. In the last day or so before leaving, we began to RF evaporatively cool. However, a big problem developed: the lifetime of the atoms in the Z-trap was around only a second and was too short. For proper evaporative cooling to a BEC using the chip’s microwires, we would have to have a lifetime more on the order of 8 to 10 seconds. Over the subsequent six months, Philipp worked hard to increase the atom number while decreasing the pressure in the chamber. By early spring 2004 he obtained a BEC on our chip. The main difficulty seems to have been getting the dispenser to not emit “dirt” along with the rubidium, and to operate it at a current that allowed the effective use of a UV lamp for pulsed Rb desorption from the cell walls.

In May and June 2004, Philipp began to attempt to run current through the splitting microwires with a BEC in position. Unfortunately, a freak accident led to the overheating and breaking of the smallest wires and the chip was rendered dysfunctional. In late May I flew-out for another three week visit, this time concentrating on a cavity QED experiment with Tilo (to be discussed below). While there, Philipp and I brainstormed variations of our scheme for splitting a BEC. He eventually settled on a double-layer technique involving small splitting wires fabricated on top of a polyimide layer coating a larger guiding wire. It was, in essence, the same design as our last chip, but divided the wires into two lay-

ers with an insulator in between. This would solve the skewed double-well problem and perhaps prevent a repeat of the wire burning incident (for details on this experiment see Chapter 10). A splitting scheme involving an atom mirror, charged pads, and microwires might be superior, but we choose not to pursue this since it would involve a complicated fabrication process. Since the spring of 2004, Philipp has been involved in developing their own in-house fabrication capability and in resolving the complications of fabricating such a double-layer device.

The macro U-MOT and waveguide

Early on—in mid 2001—Hideo and I recognized the need for a waveguide to move the atoms from the initial laser cooling and trapping zone of the atom chip to a region of the chip that was more suitable for coupling atoms to a cavity. Before we could begin to magnetically waveguide atoms, we had to develop a more robust method for loading atoms into a U-trap. During my first trip to Jakob’s group in January 2003, I took a side-trip for a two day stay at the University of Heidelberg, visiting Joerg Schmiedmayer’s group and his postdoc Ron Folman (now at the Ben-Gurion University in Israel). While there, I learned of a technique Joerg and his students were developing that would allow the trapping—from vapor—of atoms in a MOT formed by a macroscopic U-shaped wire (or rather more of a block of metal than a wire). This macro U-MOT seemed like an elegant solution to the problems I was having using a MOT made with external coils to load the microwire U-MOT. The external coils blocked optical access, rendering useless much of the 6” optical window area of my vacuum chamber. These coils were also very difficult to position so that the MOT formed optimally above the U-wires—day to day adjustments were necessary. Overheating of the coils was a problem, causing the MOT atom number to fluctuate from the heated air currents affecting the MOT lasers (cooling water would have solved this, but would have required a coil redesign). More importantly, the low inductance of the U-wire and the precise controllability of the macro U-MOT’s position would allow the fast, spatially mode-matched loading of subsequent microwire traps.

With traveling to Munich and working on those projects, it wasn’t until mid-summer 2003 that I got the chance to build my own version of the macro U-MOT. I milled-out a copper U-shaped block of roughly half-centimeter proportions and supporting 30 amps of current (see Chapter 3 for details). I mounted it on another copper block that supported three pairs of wire coils: two for creating nulling fields and one for creating the macro

U-MOT bias field. The entire assembly is the size of a softball and fits snugly in the vacuum chamber. On top of the U-block is attached a copper cladded teflon circuit board, and glued onto this is the sapphire atom chip. The circuit board is milled to allow the connection of wires ultrasonically soldered onto the atom chip's gold pads to macroscopic brass pins that are attached to wires in the vacuum chamber (see Chapter 2). Though ugly-looking and originally only intended as a temporary trial device before precision machining a permanent one, the atom chip carrier assembly has worked quite well and I still use it for my experiments. In addition to allowing larger area atom chips (for on-chip cavity experiments) and great optical access, a main benefit is its reusability: in contrast to the glued-cell technique, one can exchange atom chips or cavities on the atom chip without rebuilding the whole system. In fact, it only takes a day or two to open the vacuum chamber, take-out the chip assembly, attach a new device, replace it inside the chamber, and begin pumping down again.

I got the macro U-MOT to work in early September 2003. For this trial, the atom chip was only a mirror, but the teflon circuit board had a set of mesoscopic U- and Z-traps milled into it. The copper wires forming these mesotraps were $\sim 70 \mu\text{m}$ tall and $\sim 500 \mu\text{m}$ wide. These traps are registered to the macro U-MOT, and I was able to form macro U-MOTs of 1-2 million cesium atoms and transfer them to the mesoscopic U-MOT with little loss. The mesoscopic U-MOT configuration was also able to trap atoms from vapor, albeit with a factor of ten fewer population.

From November 2003 to May 2004, I went through several iterations of waveguiding chip experiments, each time improving the chip design and fabrication technique, chip carrier assembly, vacuum chamber, and computer control and optical detection system. The waveguiding chip employed a single wireguide with many crossing wires. These crossing wires act to gate the otherwise free expansion of the trapped atoms down the guide by forming either H-traps or Z-traps. The waveguide axis was oriented 90° from the original U-MOT axis to enable the movement of the atoms out of the MOT laser beams. The atoms were loaded into this waveguide with a P-trap transferring to a Z-trap and performed in a similar manner to what we did in the BEC experiments in Munich. I first got this working in our Caltech lab in March 2004, and took some nice images of the transfer procedure. The cross wires for this first device were a little too far separated (7 mm), and my absorption imaging wasn't good enough to detect the atoms diffusing between

gates. The next versions of the chip had gates spaced by 2 mm, and increasingly improved surface mirrors that eventually allowed good vertical fluorescence imaging. By May 2004 the waveguiding system was quite reliable and reproducible. I finished writing a versatile Matlab code that serves as the control interface for the experiment. In a fairly turn-key manner, it controls the National Instruments analog and digital output boards, sets the triggering of the CCD cameras, makes plots of the wire current control sequences, and has a nice modular integration of experimental sequences.

Around this time I improved the electroplating technique, made the chip carrier assembly more modular, took-apart the vacuum chamber for a thorough cleaning, and continued to improve the trapping technique. In early January 2005, I switched from a Z-trap to a more controllable, H-trap method for conveying atoms down the waveguide. It became apparent in March 2005, that the 90° rotation technique was unnecessary for delivering atoms to the PBG and microdisk cavities since they would not block MOT beams due to their naturally low aspect ratio. I built a new, simpler atom chip without the P-trap, and this chip should be able to guide the atoms to the cavities with higher transfer efficiency.

Cavities galore

The original goal of this research has always been to magnetically trap and guide atoms to the mode of a photonic bandgap cavity. However, it wasn't until June 2003 when we began collaborating with Oskar Painter's group that we believed that we could in the near term obtain a PBG cavity suitable for our experiments. In the meantime, Hideo and I explored the possibility of using other types of cavities in our atom chip experiments. Hideo and Kimble's group had a lot of experience with making and testing microspheres [20, 21, 22], and we thought it would be worthwhile to magnetically guide atoms to the evanescent field at the edge of the microsphere. In September 2001, I designed a waveguide scheme that would shuttle atoms from the laser trapping and cooling zone to the edge of a cut-out in the chip in which the microsphere would reside (or alternatively, the microsphere would be mounted on $\sim 100 \mu\text{m}$ tall pedestals on the atom chip). Several months later, we learned that Jakob and Prof. Hänsch's were also building such an experiment [23], and it seemed like this was a very difficult cavity QED scheme to implement: we decided to pursue other types of cavities.

We had in our labs several supermirrors from REO, and in February 2002 I worked-out a scheme for moving the atoms between two, 7.5 mm diameter Fabry-Perot mirrors counter-

sunk into the atom chip (see Chapter 7). Though the cavity mirrors would be 1 mm apart (a space necessary for inserting the microwire substrate between the mirrors), the high finesse of the cavity mirrors would enable the achievement of strong coupling with magnetically trapped atoms. Although we didn't implement this idea and instead moved-on to the PBG cavity experiment, we entered into a collaboration with Jakob's group which began building a Fabry-Perot cavity on a chip using tiny mirrors glued to the ends of fiber optics. In late May 2004, I flew-out to Munich to work on this experiment for three weeks. By the time I arrived, Tilo had already attached the cavity—of finesse 600—to the chip. The chip is of similar design as my waveguiding chip and was installed in a glued-cell vacuum chamber of design similar to the apparatus used for the BEC double-well experiment. I spent part of my time helping to get the coils, power supply, and laser system system set-up and by the end of the trip we had loaded atoms into a U-MOT. The main focus, however, was in thinking of how to optimize and understand the prospective cavity QED signal and how to lock the cavity in the presence of detunings caused by microwire wire heating. At the end of my trip, we got the locking system working, and by October 2004 Tilo and Yves Colombe—a new postdoc with Jakob—detected atom transits of magnetically trapped atoms being guided through the cavity (see Chapter 7). My main contribution since that visit has been to run master equation calculations for the purpose of determining which drive powers and cavity and laser detunings are optimal for maximizing signal-to-noise and minimizing spontaneous emission as an atom transverses the cavity. So far, the calculations qualitatively concur with the data from the experiment. To the best of my knowledge, this experiment demonstrates the first atom-cavity chip device. In the late winter of 2005, Yves took some data that hint at the observance of optical bistability in the system. I returned to Munich in June 2005 to work on this experiment some more. We had wanted to trap a BEC inside the resonator, but the UV lamp used for rubidium desorption broke and we decided to table this effort until it could be replaced. We instead concentrated on improving the signal-to-noise of the atom transit detection, and by the end had obtained much improved signals (see Chapter 7).

At Caltech over the last two years, I have been trying to integrate a PBG and/or microdisk cavity with the waveguide atom chip. The original plan was to use the PBG cavities designed by Axel Scherer's group [24]. However, the problem of efficiently coupling of light into and out of this devices hadn't been solved. In June 2003, Hideo and I met with Oskar Painter and his grad students Paul Barclay and Kartik Srinivasan, and we

learned that they had solved exactly this problem. They demonstrated extremely efficient input/output coupling to their PBG and microdisk cavities with the use of fiber tapers. We formed a collaboration and they spent the last year learning to fabricate their PBG and microdisk cavities in AlGaAs—which was difficult—and in SiN. By late 2004, they demonstrated small mode volume but so far relatively low Q devices in AlGaAs, but in early 2005 high Q's were achieved in SiN.

Ideally, we'd like to use PBG cavities for our experiments—they have much smaller mode volume than microdisks—but since microdisks are easier to fabricate and diagnose, we're exploring both options. In January 2004, I wrote a paper with Oskar, Paul, and Kartik detailing our proposed scheme for integrating a PBG cavity with the magnetic microwire traps to form an atom-cavity chip. The paper's focus was on demonstrating the feasibility of detecting single atoms with this system by solving the master equation in a two-level approximation (see Chapter 6 for details and updated calculations). A few months later, a similar analysis was performed for microdisks. It seems that we can achieve a system that is further into the strong-coupling regime than ever demonstrated before, and achieve this with compact, experimentally robust devices. For the first experimental demonstration, we plan on magnetically guiding atoms into the microcavity's mode, albeit loading in a non-deterministic fashion. After we build expertise with this technique, we will add more complicated microwire traps and eventually deterministically load atoms into the cavity while they are trapped in the Lamb-Dicke regime. Currently, Paul and I are working-out the experimental details of mounting fiber tapers to the atom chip while maintaining fiber-to-resonator coupling. We are in the process of installing this device inside the vacuum chamber, and we hope to test the first devices in the presence of trapped atoms in the fall of 2005.

Publications based on graduate work

- B. Lev, K. Srinivasan, P. Barclay, O. Painter, and H. Mabuchi, “Feasibility of Detecting Single Atoms using Photonic Bandgap Cavities,” *Nanotechnology* **15**, S556 (2004). Preprint: quant-ph/0402093
- A. Hopkins, B. Lev, and H. Mabuchi, “Proposed Magneto-Electrostatic Ring Trap for Neutral Atoms,” *Physical Review A* **70**, 053616 (2004). Preprint: quant-ph/0402037
- B. Lev, “Fabrication of Micro-Magnetic Traps for Cold Neutral Atoms,” *Quantum*

Information and Computation, Vol. **3**, No. 5, 450-464, (2003). Preprint: quant-ph/0305067

- B. Lev, Y. Lassailly, C. Lee, A. Scherer, and H. Mabuchi, “Atom Mirror Etched from a Hard Drive,” Applied Physics Letters, **83**, 395-397, (2003). Preprint: quant-ph/0304003
- H. Mabuchi, M. Armen, B. Lev, M. Loncar, J. Vuckovic, H. J. Kimble, J. Preskill, M. L. Roukes, and A. Scherer, “Quantum Networks Based on Cavity QED,” Quantum Information and Computation **1**, Special Issue, 7-12, (2001).

Chapter 1

Introduction

Central to all modern technology—from airplanes and computers, to everyday kitchen appliances—is the application of feedback control to extract desired device functionality. As we push towards designing ever smaller devices, two questions arise: how to make devices rooted in the classical domain cope with the emergence of quantum effects; and how to make quintessentially quantum components—atoms, ions, photons, nuclear and electron spins—work in concert with one another in the face of environmental perturbations. As in the classical domain, feedback control may allow the engineering of quantum systems to exhibit useful dynamics which would otherwise be unattainable. Quantum feedback, an emerging subfield of physics and information science, pertains to the control of systems in which quantum dynamics and behavior—measurement backaction, entanglement—are non-negligible.

Cavity quantum electrodynamics (QED) describes the system composed of one or more atoms, ions, or quantum dots coupled to the mode of a high finesse resonator. Feedback experiments in the strong coupling regime of cavity QED—in which the coherent dynamics dominate the dissipative—will provide an excellent setting to explore the real-time actuation and measurement of an open quantum system (see Figure 1.1). Moreover, it is a proving ground for the efficacy of current theoretical tools, such as quantum trajectory theory [25, 26], in analyzing continuously-observed quantum systems for the study of the quantum-classical transition or for the purpose of designing feedback controllers in the presence of measurement backaction. Recent proposals and experiments have highlighted the capability of quantum feedback in the neutral atom cavity QED setting [27] to perform motional cooling of intracavity atoms [28, 29] and for joint atom-cavity quantum state manipulation via the combination of detection and actuation of the intensity and phase of the cavity probe

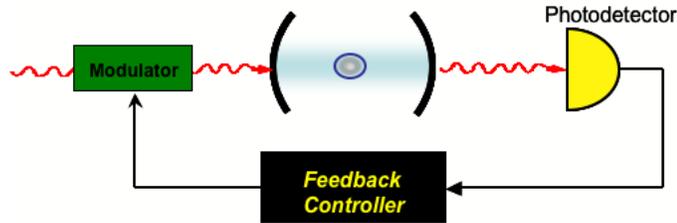


Figure 1.1: Quantum feedback in a cavity QED setting.

laser [30, 31]. A quantum to (semi)classical transition occurs in a strongly coupled cavity QED system as the number of intracavity atoms are increased from $N = 1$ to $N \gg 1$ while holding constant the total atom-cavity coupling strength, $g_{\text{eff}} = g\sqrt{N}$. Quantum feedback may aid in the study of this transition by localizing and observing the single-atom cavity QED dynamics in regions of phase-space that in the semiclassical regime exhibit nonlinear behavior such as subcritical pitchfork and supercritical Hopf bifurcations [32]. The *Science* paper by H. Mabuchi and A. C. Doherty provides an excellent review of recent experimental and theoretical research and motivation in cavity QED [33].

In addition to investigating open quantum systems under continuous measurement, cavity QED in the strong coupling regime holds great promise for the field of quantum information processing (QIP). Whereas atoms are useful for qubit storage and performing quantum logic gates, photons are optimal for transporting—via fiber optics—this quantum information over long distances. Cavity QED facilitates the reversible transfer of quantum information from atoms to photons, enabling the use of both media for building quantum networks for quantum communication and entanglement distribution [34, 35, 36]. Figure 1.2 sketches a quantum network scheme based on cavity QED.

Until recently, state-of-the-art cavity QED experiments in the optical domain have largely been performed by dropping [37], tossing [38], or optically guiding [39] cold atoms from a magneto-optical trap (MOT) between the high finesse mirrors of a Fabry-Perot cavity. A major challenge in accomplishing the aforementioned goals is to develop an experimentally robust cavity QED system in which the atom is tightly confined inside the mode of a high-finesse, low mode volume cavity for long periods of time. This long lifetime and the elimination of the stochastic variation of atom-cavity coupling would enable the implementation of quantum feedback or QIP schemes. A major breakthrough toward these

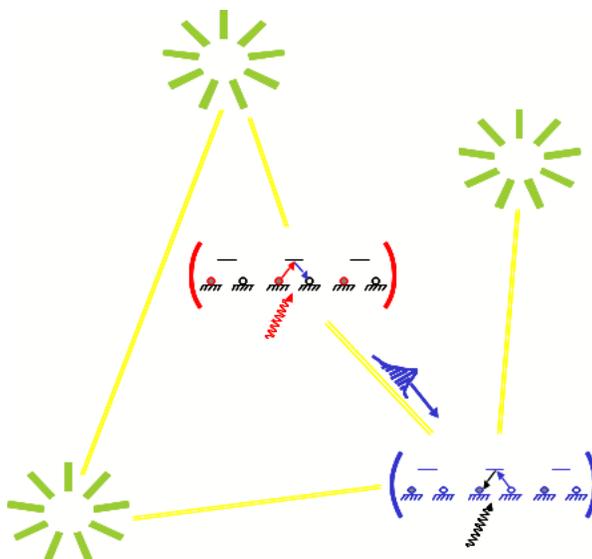


Figure 1.2: Quantum network scheme based on cavity QED. At each node resides an atom-cavity system. The state of a qubit is mapped from the atom to the cavity field. A photon leaking out of the cavity transports the qubit to a remote node where the state mapping is reversed. In this manner a quantum repeaters and entanglement distribution could be implemented [34, 35]. Figure courtesy of H. Mabuchi.

ends was achieved in Prof. Kimble's group in 2002 by the successful use of an intracavity far-off resonance trap (FORT) formed by coupling a second laser into the cavity [40]. Atom trapping lifetimes on the order of seconds were achieved, and demonstrations of a single atom laser [41], the deterministic generation of single photons from one atom [42], and the measurement of the vacuum-Rabi spectrum for one trapped atom [43] soon followed.

In 1999, our group chose an alternative route for single atom trapping inside an optical resonator. This involves the magnetostatic confinement of atoms, and has the potential advantage of enabling experimentally robust, scalable, and fully integrated cavity QED systems to be built on a chip. This thesis documents our burgeoning efforts to combine the tools of nanotechnology with atom trapping and cooling to produce such a chip-based cavity QED system. Moreover, it describes how we have begun to harness the capabilities of micro- and nanofabrication to enable the development of better ways to manipulate atoms for the purpose of investigating quantum phenomena associated with Bose-Einstein condensates (BEC) and to create new atom optical devices such as atom mirrors and 1-D ring traps.

Microfabricated wires on a substrate—known as an atom chip—can create a sufficiently high-curvature magnetic potential to trap atoms in the Lamb-Dicke regime. We have de-

veloped a state-of-the-art atom chip fabrication capability [8] for making these devices “in-house.” Chapters 3, 4, and 5 detail our atom chip fabrication and trapping techniques as well as the apparatus we have built for chip-based atom trapping and cooling. Over the past six years we have developed a streamlined experimental system for producing custom-made atom chips and for using them to trap and convey micro-Kelvin atoms ten to one hundred microns above the room-temperature chip surface into an optical cavity.

Chapter 6 discusses the experiments being performed in collaboration with the Caltech Applied Physics group of Professor Oskar Painter to couple single cesium atoms into the mode of an on-chip photonic bandgap (PBG) or microdisk cavity via the aforementioned atom chip system. The natural proximity of the atoms to the chip surface introduces the possibility of replacing the Fabry-Perot cavity with these more compact and rigid cavities of much smaller mode volume. This enables stronger atom-cavity coupling, and the planarity of these cavities allows the straightforward integration with fiber tapers and photonic waveguides. Scalability and miniaturization is inherent in these micro- and nanofabricated systems, and an architecture for a multi-qubit quantum network device on a single atom-cavity chip may be achievable [1].

A great challenge recently tackled by the Painter group was the fabrication of these devices in AlGaAs and SiN_x—materials suitable for cavities resonant with cesium—and the development of robust techniques for coupling light into and out of the cavity. At the time of this writing, we are in the process of integrating these new PBG and microdisk cavities with our atom chip trapping apparatus. Numerical master equation calculations indicate single atom detectability with this system, as discussed in Chapter 6 and in Reference [44].

In 2002, we began a collaboration with the group of Professors Theodore Hänsch and Jakob Reichel at the Max Planck Institute for Quantum Optics in Munich. Originally intended as an exchange of our atom chips for an opportunity to learn on-chip Bose-Einstein condensate (BEC) production in their labs, this collaboration has blossomed into a full partnership—including chip design, theoretical modeling, experimental apparatus assembly, and data analysis—on two experiments during several visits of the author to Munich. In a cavity QED experiment similar to the one detailed above, we have successfully integrated a miniature optical fiber-gap Fabry-Perot cavity with an atom chip in the Munich labs. Magnetically trapped intracavity atoms have been detected with this cavity QED system. Chapter 7 details the results of master equation calculations that suggest this cavity—

though in the weak-coupling limit—should be able to detect single atoms. We are currently improving detection sensitivity and are attempting to magnetically load an on-chip BEC into the cavity.

The atom chip and associated fabrication techniques have proven valuable in research outside the field of cavity QED. In our labs at Caltech, we fabricated and demonstrated a specular magnetic atom mirror formed from an ordinary hard drive, which is discussed in Chapter 8 and Reference [12]. Ion milling of the hard drive surface patterned large area, high resolution permanent-magnetic structures on these flat, rigid, and inexpensive substrates. The periodically modulated magnetic domains produce an exponentially repulsive, high remanent magnetic field for reflecting cold neutral atoms, making the device ideal for creating waveguides, corner cubes, or other atom optical devices. Moreover, electric fields from integrated charged pads on the mirrors surface can perform quantum logic gates, and the device’s large coercivity allows the placement of microwires on its surface for additional time-dependent magnetic fields.

We recently proposed a device incorporating micron-sized charged circular pads fabricated on the hard drive atom mirror and is discussed in both Chapter 9 and Reference [45]. Through the balance of the repulsive magnetic potential and the attractive electrostatic, this device can produce a 1-D ring trap that may prove useful for both studying Tonks gases in a ring geometry and for creating a SQUID-like system for cold neutral atoms.

Our second joint experiment in Munich involves an atom chip with wire features capable of splitting a BEC in a magnetic double-well potential. The first generation of this chip was fabricated by ourselves at Caltech, and we hope to be able to perform atom interferometry and investigate Josephson effects with future versions of the device. Solutions to the Gross-Pitaevskii equation for a BEC in this double-well trap indicate that Josephson oscillations are observable with such an atom chip device, and this study is presented in Chapter 10. We constructed a compact atom chip vacuum chamber system and obtained an on-chip BEC. This type of chamber system will be useful for future experiments in our lab at Caltech, and its design is also described in Chapter 10.

Chapter 2

Atom Chips

2.1 Introduction

Cold samples of neutral atoms and Bose-Einstein condensates have become readily available using the techniques of laser cooling and trapping [46], and it has been widely recognized that cold atoms are a rich resource for experiments in quantum information science and quantum feedback. For many proposals, however, quantum control of the atomic motional degrees of freedom is essential. For example, many proposals for quantum computation in a cavity QED setting or through controlled cold collisions require the ability to trap and control single atoms in the Lamb-Dicke regime [34, 47, 48]. In 1995, Weinstein and Libbrecht noted that micron-sized wires, fabricated on a substrate, are capable of producing the large magnetic field gradients and curvatures required for trapping atoms in this regime [3]. Westervelt *et al.*, in 1998, succeeded in fabricating the wire patterns used in the trap designs of Weinstein and Libbrecht [49]. These microwire devices, now commonly known as atom chips [7], have been used to great success in atom optics and in the production of Bose-Einstein condensates (BEC), and are promising tools not just for quantum computation, but for atom interferometry, cavity QED, and the study of cold collisions as well [2, 7, 4, 6]. In this chapter we describe the design of atom chips and the methods we use for cooling and trapping atoms with these magnetic micropotentials.

Atom optical elements, such as mirrors, waveguides, splitters, traps, and conveyor belts have been demonstrated using atom chips [50, 51, 52, 10, 53, 11, 12]. Cesium cold collisions in the presence of light have been studied using a magnetic microtrap (see Chapter 5). The use of fiber-gap cavities [54, 23] has recently been realized for on-chip atom detection (see Chapter 7), and on-chip microsphere [23, 55], microdisk [56, 57], and photonic bandgap cav-

ities [44] are being explored (see Chapter 6). Ion trap experiments are now using substrates with microfabricated electric pads for the purpose of controlling ion position [58, 59].

On-chip production of a BEC has been one of the most successful uses of the atom chip thus far [60, 61, 62, 63]. Ioffe traps formed from microwires can produce extremely large trap compressions that enhance the efficiency of evaporative cooling. Consequently, condensate production time can be reduced from one minute to less than ten seconds [60]. This allows magneto-optical trap (MOT) loading from a thermal vapor in a glass cell with a vacuum of only a few 10^{-10} Torr (see Chapter 10). All of the required magnetic fields can be produced on-chip [64, 5], removing the necessity of large, high power external coils. The atom chip greatly miniaturizes BEC production and will enable the integration of matter waves with chip-based atom optics and photonics.

Another exciting avenue of research involves the use of an atom chip to trap, in the Lamb-Dicke regime, one or more atoms in the mode of a high finesse cavity. The combination of magnetic microtraps and photonic bandgap (PBG) cavities would be an excellent cavity QED system for the implementation of scalable quantum computation, or for the study of continuous measurement and quantum-limited feedback. One technical proposal involves the integration of a PBG cavity with an Ioffe trap formed from microwires patterned on the same surface [1]. The combination of small mode volume and modest optical quality factor that should be obtainable with PBG structures would enable strong atom-cavity coupling. This would be an interesting alternative to present experiments that utilize a Far Off Resonance Trap (FORT) to confine atoms inside optical Fabry-Perot cavities [40]. Several PBG cavities, each with an independent microwire trap, could be fabricated on the same substrate and coupled together with a network of line-defect optical waveguides.

2.2 Microwire traps

Magnetic traps exploit the interaction potential, $V = -\vec{\mu} \cdot \vec{B}$, between an atom's magnetic moment, $\vec{\mu}$, and a magnetic field, \vec{B} , to trap or guide weak-field seeking states of a neutral atom. Typical magnetic traps are formed from coils of wire arranged outside the vacuum chamber. This results in macroscopic coil radii on the order of 10 cm which require the use of tens to hundreds of amps of current, I . These coils produce either homogenous biasing fields or trapping fields in a quadrupole or Ioffe configuration. The on-axis field from a

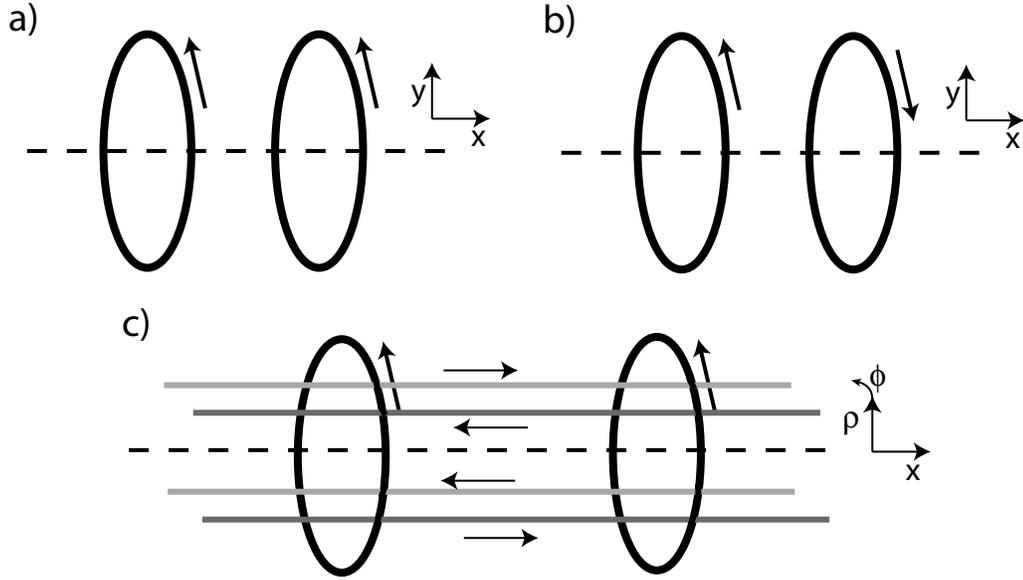


Figure 2.1: a) Helmholtz coils for homogeneous magnetic field production. The arrows denote the current direction. b) Anti-Helmholtz coils for creating a quadrupole trap. c) Ioffe trap.

single coil of radius, R , and positioned a distance A from the origin $x = 0$ is,

$$\mathbf{B}_{coil}(x) = \frac{\mu_0 n I R^2 \hat{x}}{2[(x - A)^2 + R^2]^{3/2}}, \quad (2.1)$$

where $\mu_0 = 4\pi \times 10^{-7}$ N/A². In this equation, n is the number of turns of wire carrying current I . A pair of coils positioned on axis with one another, separated by $2A$, and having the same orientation of current flow produce a homogenous magnetic field, $2\mathbf{B}_{coil}$, at the midpoint between the coils (see Figure 2.1 [a]). Maximum field homogeneity is achieved when $A = R/2$. This cancels the quadratic field variation at $x = 0$, and this coil arrangement is called the Helmholtz configuration:

$$2\mathbf{B}_{coil}(0) = \left(\frac{4}{5}\right)^{3/2} \frac{\mu_0 n I}{R}. \quad (2.2)$$

This method for producing a homogenous magnetic field proves quite useful for nulling ambient fields, defining a quantization axis, and for producing bias fields that can be combined with non-homogenous fields from magnetic traps and waveguides. Expressions for the off-axis field of this coil configuration and those discussed below can be found in Reference [65].

The most basic magnetic trap is formed by simply reversing the current direction in one of the Helmholtz coils (see Figure 2.1 [b]). This anti-Helmholtz coil arrangement forms a quadrupole field which is zero at the trap center and increases linearly from the center. Near the trap center—located at $x = 0$, $y = 0$, and $z = 0$ —the field is:

$$\mathbf{B}_{quad} = \frac{3\mu_0 n I A R^2}{(R^2 + A^2)^{5/2}} \left[x\hat{x} - \frac{1}{2}y\hat{y} - \frac{1}{2}y\hat{z} \right]. \quad (2.3)$$

For $A/R \approx 0.63$, the trap depth in \hat{x} is equal to that in the $\hat{y} - \hat{z}$ plane. The quadrupole trap is easy to create, and forms a MOT when combined with red-detuned laser beams impinging the trap center from six directions along the quadrupole trap's axes [66]. The MOT has become the workhorse for modern atomic physics. The quadrupole trap has a disadvantage in that the field vanishes at the trap center which can cause atoms to be lost due to Majorana spin-flips (see Chapter 5). To avoid this loss mechanism, an Ioffe trap¹ is used which plugs-up the trap minimum with a finite field. This transforms the linearly varying quadrupole potential into a harmonic potential. As pictured in Figure 2.1 (c), the Ioffe trap is formed from a coil pair more widely separated than the Helmholtz configuration to provide a harmonic trap in \hat{x} and four straight current-carrying wires arranged symmetrically about the coil axis. These wires provide confinement in the $\hat{y} - \hat{z}$ plane perpendicular to the coil axis. At the trap minimum, the field is to quadratic order:

$$\begin{aligned} B_z &= C_1(z^2 - \frac{1}{2}\rho^2) + B_{bias} + \dots, \\ B_\rho &= -C_1 z \rho + C_2 \rho \cos(2\phi) + \dots, \\ B_\phi &= -C_2 \rho \sin(2\phi) + \dots, \\ B &\approx B_{bias} + C_1 z^2 + (C_2^2/B_{bias} - C_1)\rho^2/2. \end{aligned} \quad (2.4)$$

The constants, C_1 and C_2 are determined for a specific wire layout [65, 3]. The bias field, B_{bias} , is the residual field at the trap center which prevents Majorana spin flips, and the Ioffe trap is stable for $C_2^2/C_1 > B_{bias} > 0$.

A major goal of atom trapping is the confinement of atoms in the Lamb-Dicke regime. This is important, for instance, in cavity QED where one wants to localize the atom to a nearly constant region of an optical field. In ion trapping experiments, confinement in

¹This is sometimes referred to as an Ioffe-Pritchard trap.

the Lamb-Dicke regime is crucial for Raman sideband cooling which requires the spectroscopic resolution of the trap’s vibrational levels. The Lamb-Dicke regime is defined as $\eta = (E_{recoil}/E_{vib})^{1/2} < 1$, where η is the Lamb-Dicke parameter [67] (see Figure 2.2). An η less than one implies that the scattering of a photon of wavelength λ is of insufficient energy to excite the atom to higher vibrational levels. In other words, excitation out of the ground state—or out of any other state—of the trap is suppressed by a factor of η . This condition is equivalent to viewing the trapped atom’s wavepacket extent, x_0 , as smaller than the wavelength, λ , of resonant light: $\lambda \gg x_0 = \sqrt{\hbar/2m\omega_{vib}}$.

In general, the field’s magnitude, gradient, and curvature scale as I/r , I/r^2 , and I/r^3 , respectively, where I is the wire’s current and r is its characteristic dimension such as radius. To create traps with field curvature high enough to trap atoms in the Lamb-Dicke regime, one needs to either increase the coil current I or shrink the magnetic field coils. Currents in excess of a few hundred amps become difficult to produce and work with in the lab. Moreover, macroscopic coils are space-inefficient, and it can be difficult to align their trap minima with microscopic devices—such as microfabricated cavities—inside the vacuum chamber (see Chapters 6 and 7). Alternatively, microscopic wire patterns maximize field gradients and curvatures while keeping power dissipation to a minimum. Reducing r to the 1 to 10 μm scale allows extremely high trap gradients and curvatures to be produced for modest < 1 A currents and is sufficient to trap atoms in the Lamb-Dicke regime. Of course, microscopic wires are difficult to arrange inside a vacuum chamber. We overcome this by co-opting the techniques of computer chip microfabrication—namely, photolithography—to pattern these wires on ridged, thermally conductive substrates. These devices are commonly known as atom chips and the next section discusses how to create analogues of the quadrupole trap, waveguide, and Ioffe trap using planar pattern of wires and easily generated homogeneous bias fields.

2.3 Zoology of microtraps

Since the proposal of magnetic microtraps by Libbrecht’s group in 1995 [3], there has been an explosion in the field of atom chips with approximately ten groups employing them worldwide. In this section, we present the basic microtrap building blocks that are used to construct the various atom chip devices used today. The groups of J. Reichel,

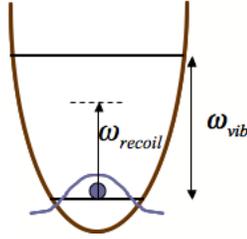


Figure 2.2: An atom is trapped in the Lamb-Dicke regime when the atomic recoil energy is much less than the vibrational levels of the trap.

J. Schmiedmayer, E. Hinds, and C. Zimmermann have written excellent reviews on this subject [50, 68, 51, 10, 69].

2.3.1 The wireguide, U-trap, and Z-trap

The most commonly used magnetic microtraps are those based on the wireguide which is composed of a single straight wire and a bias field, B_{bias} , perpendicular to the wire [70]. Operated by itself, this device functions as a 2-D magnetic waveguide for transporting neutral atoms. Figure 2.3 (a) depicts how a quadrupole field is formed with this device. Current in a straight wire produces a curling magnetic field that decreases as r^{-1} from the wire. A bias field, B_{bias} , perpendicular to the wire cancels the curling magnetic wire along a line suspended above the wire at a distance

$$y = \frac{\mu_0 I}{2\pi B_{bias}}. \quad (2.5)$$

The resulting field, shown in the vector plot of Panel (a), is a 2-D quadrupole waveguide that extends the length of the wire. Atoms in weak-field seeking states will be drawn to this field minimum and can freely propagate along the length of the wire, hence the name wireguide for the wire and field arrangement pictured in Panel (b). The gradient of this 2-D quadrupole field is

$$\nabla B = \frac{2\pi B_{bias}^2}{\mu_0 I}. \quad (2.6)$$

The trap height and gradient can be controlled by choosing the proper combination of I and B_{bias} . This assumes that the wire is infinitely thin, cylindrical, and can withstand arbitrarily high currents. The former assumption breaks-down when y is less than a wire width from the surface. For more information on this and the maximum power that the

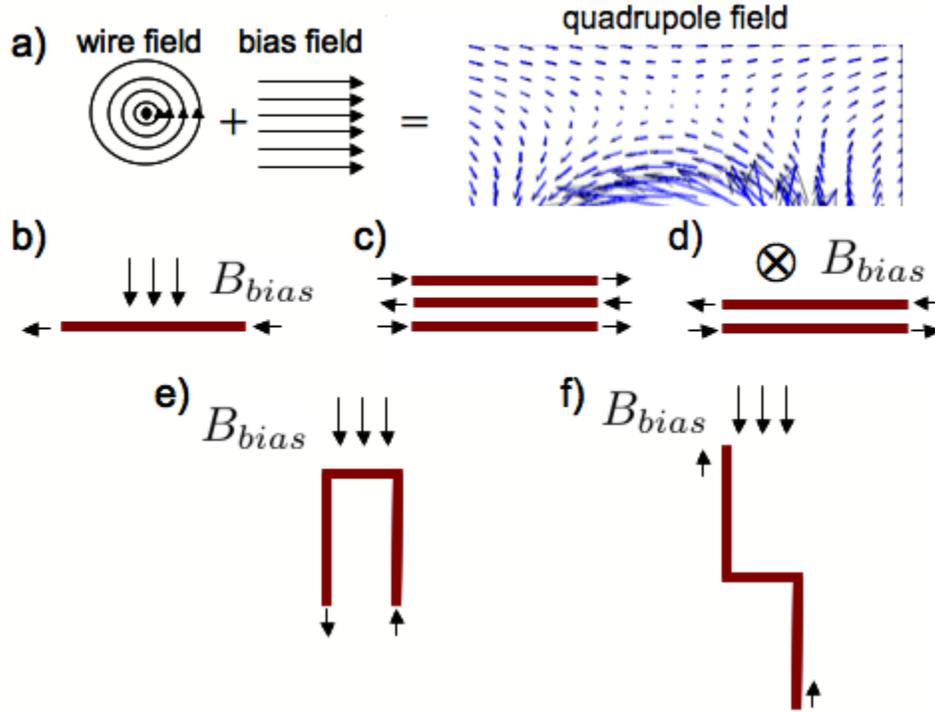


Figure 2.3: a) A 2-D quadrupole field—for use as a waveguide—may be formed from a bias field perpendicular to a current carrying wire. b) Simple single-wire waveguide [70]. c) Same as (b) except the externally generated bias field is replaced by two wires with oppositely flowing current [71, 72]. d) A waveguide formed from two wires and a bias field oriented perpendicular to the plane of the wires [4, 6]. This device may be used as an atom interferometer: the trap minimum may be split into two depending on the current and bias field values [73]. e) The 3-D quadrupole trap referred to as a U-trap [2]. f) The Ioffe trap referred to as a Z-trap [2].

wire can support, see Chapter 4. Reference [68] contains a good discussion regarding the constraint placed on wire size and trap gradient due the microwire power dissipation.

Figures 2.3 (c) through (f) show the various types of waveguides and traps formed in a similar manner. Panels (c) and (d) show alternative methods for waveguiding atoms. In Panel (c) the bias field is replaced by microwires with currents running in the direction opposite to the central wire. Panel (d) shows a double wire guide employing a bias field oriented perpendicular to the wire plane. This offers the advantage of enabling the waveguide to bend around curves and even into a spiral [74].

Perhaps the most useful atom chip devices are the U- and Z-traps pictured in Panels (e) and (f), respectively [2]. The U-trap forms a 3-D quadrupole trap by “pinching-off” the waveguide field at the two sharp wire bends. The “side” wires have currents in opposite

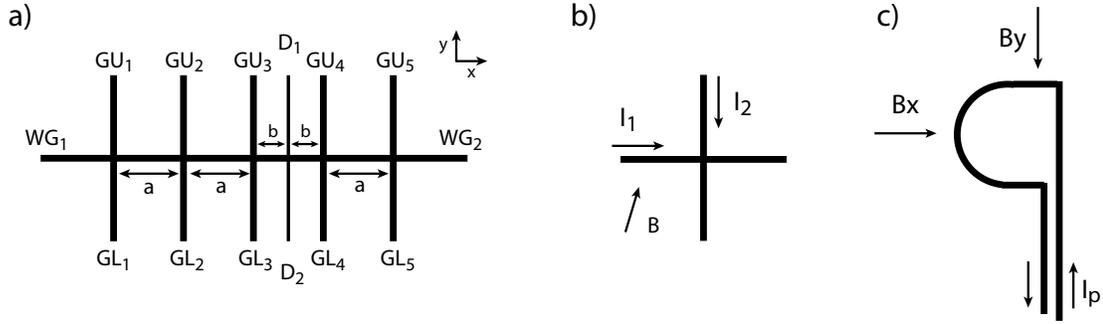


Figure 2.4: a) Microwire layout for creating an H-trap, gated waveguide, and dimple trap. b) Dimple trap: two crossed wires with $I_1 > I_2$ and bias field direction as shown. c) The P-trap: the quadrupole trap can be rotated 90° by swapping B_x with B_y while maintaining constant current, I_P .

directions that produce fields that cancel at the trap center, thereby completely forming a 3-D quadrupole field. In this manner a cigar-shaped trap is formed whose trap minimum, y , and gradient (in the two tightly trapping dimensions) is approximately given by Equations 2.5 and 2.6. The Z-trap is similar to the U-trap, but one of the side wires is extended in the opposite direction. The fields from the parallel currents in the side wires add to form a harmonic potential minimum at the trap center. This is an easy way to make a microwire Ioffe trap and is employed extensively in creating on-chip BECs (see Chapter 10). Appendix A contains the exact field expressions for the U- and Z-traps.

2.3.2 Variations: H-trap, dimple trap, P-trap, et cetera

The basic microtrap elements presented in the previous section can be combined to form more complicated devices. While there have been many such variations [68, 51, 75], we present here only those that have played a role in the experiments discussed in this thesis.

While the single-wire waveguide performs as intended, a guided atomic cloud quickly expands throughout the entire volume, reducing its phase-space density and making it hard to control or image. A more useful device is one that has cross-wires that act as gates for the atoms. Manipulating these gates like locks in a river allows one to transport the atoms over centimeter distances while maintaining reasonably high phase-space densities. Figure 2.4 (a) illustrates a typical wire layout for transporting atoms, and actual atom chips that employ this technique are pictured in Figures 2.7 and 4.2. The wireguide (WG) is

defined by endpoints $[WG_1, WG_2]$. Gate wires—defined by their upper (GU) and lower (GL) endpoints—intersect the wireguide at regular intervals, a . These gate wires actually make electrical contact with the wireguide. It is important to note that since these intersections only define the potential at one point along the gate wires, we are free to self-consistently adjust the potentials at the various terminals to allow arbitrary current flow through the wires. As an example, let us assume that the wires defined by $[WG_1, WG_2]$ and by $[GU_1, GL_1]$ are each connected to a floating, current-controlled power supply. All other wire terminals (i.e. GU_i and GL_i for $i = 2$ to 5 and D_1 and D_2) are disconnected (floating). We essentially have two wires crossed at a single point, P , and we (or in practice, the current supplies) are free to define the potential at this point. If, for example, we want to have 2 A flow through wire $[+WG_1, -WG_2]$ and 1 A through $[-GU_1, +GL_1]$, where the \pm denotes the terminal with a more positive voltage and hence defines the direction of current flow. Assuming, for instance, that the wire resistances in each segment are $[WG_1, P] = [P, WG_2] = 3 \Omega$ and $[GU_1, P] = [P, GL_1] = 1 \Omega$, a possible voltage scheme that would achieve this current flow is: $WG_1 = 10 \text{ V}$, $WG_2 = -2 \text{ V}$, $GU_1 = 3 \text{ V}$, and $GL_1 = 5 \text{ V}$. The voltage at the point P would have to be 4 V. Since the power supplies are floating, the same currents would flow even if an arbitrary constant voltage were to be added to each wire terminal and intersection point, P . To prevent the power supplies from railing, one can ground either the WG_1 or the WG_2 terminal. However, at most only one terminal in the circuit can be grounded—the rest must float for this scheme to work. The strength of this technique lies in the fact that the current in one wire, say $[GU_1, GL_1]$, can be changed without affecting the current in the other wire(s). For the case above, if we instead want a current of 3 A to flow in wire $[+GU_1, -GL_1]$, then the voltages can be changed to $GU_1 = 7 \text{ V}$, and $GL_1 = 1 \text{ V}$ without changing the voltage at point $P = 4 \text{ V}$, which maintains the current in $[+WG_1, -WG_2]$ constant at 3 A. This technique can be extended to simultaneously include all the wires shown in Figure 2.4 (a). The only condition is that each wire have its own floating power supply and only one wire can be multiply connected (and only this wire can be grounded at one terminal). In practice, the multiply connected wire, $[+WG_1, -WG_2]$, sets the potential at each of the gate wire intersection points, and the gate wire floating power supplies adjust their \pm terminal potentials to accommodate the potentials at the intersection points while simultaneously driving the required current.

The wire pattern shown in Figure 2.4 (a) is quite versatile: U-, Z-, and H-traps can

be formed as well as a gated waveguide and dimple trap. U- and Z-traps can be formed in several ways: for instance, by running current in wires such as $[\text{GL}_1, \text{GL}_2]$ and $[\text{GL}_1, \text{GU}_2]$, respectively. The H-trap is a more flexible version of the U- and Z-traps, and it can be formed from a three-wire pattern such as $[\text{WG}_1, \text{WG}_2]$, $[\text{GU}_1, \text{GL}_1]$, and $[\text{GU}_2, \text{GL}_2]$. With a bias field in $+\hat{y}$, an Ioffe trap is formed with a $+\hat{x}$ current in $[\text{WG}_1, \text{WG}_2]$ and parallel currents in the gate wires $[\text{GU}_1, \text{GL}_1]$ and $[\text{GU}_2, \text{GL}_2]$. Adding a bias field in $\pm\hat{x}$ either adds or subtracts from the field created by the gate wires, and the former increases the potential barrier of the gate wire and is useful for both increasing the Ioffe trap's oscillation frequency in \hat{x} and for increasing the magnitude of the field at the trap center. Making the currents run anti-parallel in the gate wires produces a quadrupole trap. The H-trap is desirable since the strength of the fields in the two gate wires can be independently controlled. For instance, the H-trap mentioned above can form the starting trap for loading the gated waveguide: as the current in $[\text{GU}_2, \text{GL}_2]$ is ramped down, the atom cloud is released into the waveguide formed from $[\text{WG}_1, \text{WG}_2]$. Moreover, the current in $[-\text{GU}_1, +\text{GL}_1]$ can be ramped-up in the presence of a bias field in $+\hat{x}$ to give the atoms an extra push.

While a waveguide can be formed solely from a bias field in $+\hat{y}$, $[\text{WG}_1, \text{WG}_2]$, and $[\text{GU}_1, \text{GL}_1]$, it is sometimes more convenient to use an extended Z-trap formed by $[\text{GL}_1, \text{GU}_5]$, for instance. A short Z-trap, such as $[-\text{GU}_1, +\text{GL}_1]$, can be transformed into a longer one, $[-\text{GU}_5, +\text{GL}_1]$, by using two independently-controlled, floating power supplies connected to the same wire for one terminal but to different wires for the other terminal. Gradually turning-off the current from the power supply connected to terminal GU_1 while turning-on the current from the other power supply connected to GU_5 produces this transfer without heating the atoms. Switching the currents on and off with an “error function,”

$$\text{erf}(t) \equiv \frac{2}{\sqrt{\pi}} \int_0^t e^{-z^2} dz, \quad (2.7)$$

provides a smooth transition while maintaining constant current in the shared segments of the wire.

Regardless of the method for releasing the atoms, the gated waveguide transports atoms in the following manner. At $t = 0$, the atoms are released into the waveguide and the gate $[-\text{GU}_4, +\text{GL}_4]$ is turned-on along with a bias field in $+\hat{x}$. The atoms expand down the waveguide (in the $+\hat{x}$ direction) until they are past the gate $[\text{GU}_2, \text{GL}_2]$. Current is turned-

on in this gate, $[-\text{GU}_2, +\text{GL}_2]$, which prevents the atoms from expanding backwards to the starting point. We now have the atoms trapped between gates $[-\text{GU}_2, +\text{GL}_2]$ and $[-\text{GU}_4, +\text{GL}_4]$. This procedure can be repeated with gates $[-\text{GU}_2, +\text{GL}_2]$ and $[-\text{GU}_4, +\text{GL}_4]$, and in this manner the atoms can be transported down a waveguide arbitrarily long. We demonstrated a gated waveguide of length 8 mm both in our Caltech lab and in those in Munich. Atom transfer usually occurs in ~ 100 ms. A similar style waveguide transported atoms over a ~ 7 cm distance in an experiment in Munich [76, 23]. We found that gate spacings of $a \approx 2$ mm maintain a reasonably high level of trap phase-space density while minimizing the number of gates. At Caltech, we originally tried to use a gated waveguide with $a = 7$ mm, but found that the atom cloud became too diffuse for diagnosing and optimizing transport.

The thinner wire, $[D_1, D_2]$, depicted in Figure 2.4 (a) may be used to form a dimple trap. Figure 2.4 (b) sketches the basic element of a dimple trap: two crossed wires and a bias field, B . For $I_1 > I_2$ and bias field, $\mathbf{B} = B_x \hat{x} + B_y \hat{y}$ where $B_y > B_x$, the confinement in the \hat{x} direction is stronger than in \hat{y} . We call this a dimple trap since the I_2 current forms a small potential well in the otherwise perfectly 2-D quadrupole trap formed from I_1 and B_y . This dimple trap is centered above the intersection of the two wires and is an Ioffe potential. The wire $[+D_1, -D_2]$ in Figure 2.4 (a) forms a dimple trap when the bias field is in $+\hat{x}$. The dimple trap is commonly used in conjunction with an H- or Z-trap, augmenting trap compression and collecting the cold atom cloud directly above the dimple center. The width of the wire used for making the dimple trap is typically smaller since the atoms in this potential well will be closer to the substrate and less current is required. It should be noted that the dimple trap may be viewed as the most fundamental magnetic microtrap building block since the U- and Z-traps are each formed from the conjunction of two dimple traps [50, 68]. Appendix B contains a Matlab code that calculates the field, gradient, and curvature from an arbitrary combination of dimple traps.

The P-shaped microwire pattern shown in Panel (c) of Figure 2.4 is a device for rotating the axes of a quadrupole microtrap by 90° [76]. This is useful for moving magnetically trapped-atoms outside the zone of laser beam convergence used in making a mirror MOT (see Chapters 6 and 7). The P-trap is essentially a two-sided U-trap. Typically the atoms are loaded into the quadrupole field at the tip of the ‘‘P’’ opposite to the B_x bias field. Even though the wire is curved in a half-circle there, a U-like trap can be formed. Once the atoms

are loaded in this section, they can be transported to the top, straight section by turning the \hat{x} field off while turning on the bias field in \hat{y} . This is best accomplished using a sinusoidal ramp: $B_x(t) = \cos(\pi t/2\tau)B_x$, $B_y(t) = \sin(\pi t/2\tau)B_y$, where τ is the ramp time which should be ~ 100 ms to ensure that the movement does not heat the atom cloud. This transfer time can be shortened if the magnetic trap lifetime is too small due to a high vacuum chamber pressure. We note that the atoms can be rotated in the opposite direction (to the Z-trap like zone), by ramping B_y in the opposite orientation: $B_y(t) = -\sin(\pi t/2\tau)B_y$. The U-trap side is more useful since the P-trap can be embedded in a larger H-trap wire network and the atoms transferred between the two at this wire segment (see Figure 2.7 and 4.2). We have found that extending the ‘‘P’’ shape to a ‘‘D’’ by adding a straight section of wire in the middle of the half-circle increases the trapping volume. Rotation in this D-trap works just like the P-trap except that one has to first add a bias or gradient field to shift the atoms from the center of the straight-wire region of the D-trap to the beginning of the arced region before sinusoidally ramping the bias fields. This ‘‘shift-then-rotate’’ procedure minimizes atom loss. A good test for measuring the rotation transfer efficiency is to rotate the atoms back to the trap starting point to measure the remaining atom number with the same detection viewpoint.

While the gated wireguide is simple and sufficient for transporting a thermal cloud of atoms, it does not maintain high enough phase-space density to transport a BEC over long distances. A microwire pattern developed by the Munich group can act as a conveyor belt to move a BEC, and this device was incorporated into some of the atom chips we made for this group (see Figure 4.5). See References [77, 60] for more details regarding this conveyor belt.

2.3.3 The Libbrecht-style Ioffe trap

The Z-trap can only confine atoms two-dimensionally in the Lamb-Dicke regime, and the conductive wire is always directly between the trapped atoms and the substrate. Other trapping schemes, such as the double-wire guide or the dimple trap, improve upon one or the other of these constraints but not both. In contrast, the Ioffe traps proposed in Reference [3] can confine atoms three-dimensionally in the Lamb-Dicke regime with equal trap frequencies in all directions. Moreover, the atoms are trapped above bare substrate, opening-up the possibility of integrating microcavities directly below atoms confined in the Lamb-Dicke

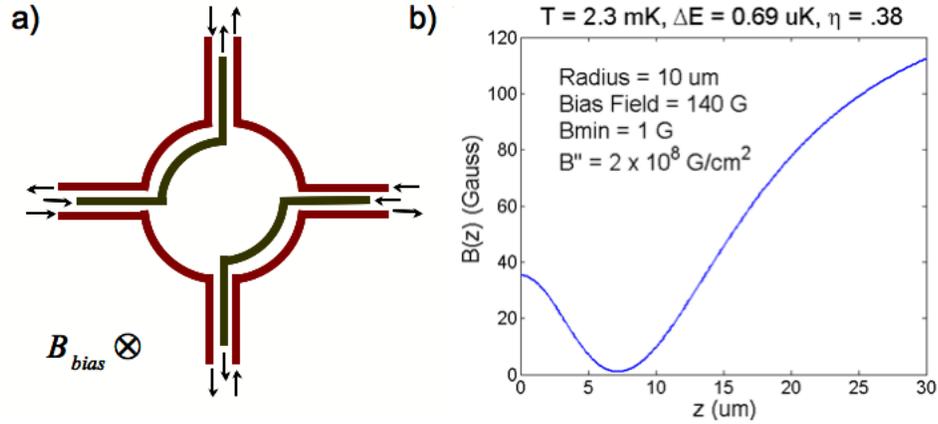


Figure 2.5: a) The Libbrecht-style Ioffe trap. A 3-D Ioffe trap is formed a distance approximately equal to the wire radius above the wire substrate. b) Plot of the magnetic field along the axis of the trap perpendicular to the wire substrate. The current in the wires is $I = 1$ A and the outer ring radius is $R_o = 15 \mu\text{m}$ and the inner is $R_i = 10 \mu\text{m}$. The trap curvature in the plane parallel to the wire substrate is $2 \times 10^{10} \text{ G/cm}^2$ with $\eta = 0.11$.

regime (see Chapter 6 for more details on this proposal). Figure 2.5 (a) shows the wire layout for this kind of trap. In contrast to the cylindrically symmetric quadrupole field produced by a U-trap, an approximately spherically symmetric quadrupole field is formed from a ring-shaped wire on an atom chip and a homogeneous bias field perpendicularly opposing the field produced by this conductor. (The external bias field may be replaced by the field from a second, concentric wire of different radius carrying an oppositely flowing current.) The Libbrecht-style Ioffe trap is formed by using two wires in quarter-circles to perturb the quadrupole field formed from an external bias B_{bias} and an approximately ring-shaped wire on the substrate. The field from the inner ring segments “plug-up” the zero at the center of the quadrupole field, forming 3-D harmonic potential described by Equations 2.4. There are several different variations of the planar 3-D Ioffe trap proposed by Weinstein and Libbrecht in Reference [3], and we chose to display version “(c)” due to its optimal balance of large trap depth and high trap frequencies. As noted in that paper, the optimal trap is formed when the currents in each wire segment are equal and the radius of the outer ring is 1.5 times the radius of the inner arcs. Panel (b) of Figure 2.5 shows the trapping field as a function of the distance from the wire substrate taken through the central axis of the wire pattern. This is plotted for $B_{bias} = 140$ G and an outer radius of $R_o = 15 \mu\text{m}$ and an inner of $R_i = 10 \mu\text{m}$. For cesium to be confined in the Lamb-Dicke

regime, the trap curvature must exceed 2×10^6 G/cm². For a wire current of 1 A, the trap curvature in \hat{z} is 2×10^8 G/cm², resulting in a Lamb-Dicke parameter of $\eta = 0.38$. In the plane parallel to the wire substrate, the trap curvature is 2×10^{10} G/cm² with $\eta = 0.11$. The trap depth is 2.3 mK— 10^3 times larger than the typical temperature of sub-Doppler cooled atoms—and the vibrational level splitting is ~ 0.7 μ K. Although a bias field of 140 G is difficult to generate, this is not impossible. A more severe difficulty lies in the fact that the trap minimum is only $z_{min} = 7$ μ m above the surface. This height scales with the outer radius wire ($z_{min} = 0.72R_o$), and for a cesium atom and a trap current of 1 A, the Lamb-Dicke regime is attained for a trap radius (and trap height) ≤ 40 μ m. Atoms have been trapped as close as ~ 2 microns from a surface before surface effects quench trap lifetime severely (see Chapter 9 and the references within). The small radius of the trap wires poses another difficulty in that the wire widths can only be a few microns wide if the arcs are to be sufficiently approximated. These small wires must be able to support current pulses of nearly 1 A. The chapter on atom chip fabrication discusses these issues in more detail.

A simple modification to the original Libbrecht-style trap yields higher trap frequencies. This trap was designed by an undergraduate in our lab, Arjun Menon, in 2000 and is pictured in Figure 2.6 (a). For similar currents and wire radii as the trap in Figure 2.5 (a), this design produces a trap with curvature 10 times higher in \hat{z} and a factor of 2 to 5 times higher in the plane parallel to the wire substrate. This produces a factor of ~ 2 decrease in the z -axis η . Unfortunately, this trap requires a factor of 3 increase in B_{bias} , and the trap minimum, z_{min} , is closer to the substrate.

2.3.4 Permanent magnets and RF and electrostatic fields

DC currents are not the only possible tool used to manipulate atoms in micropotentials. Permanent magnets can provide high gradients and curvatures without the large heat dissipation associated with electromagnets. Moreover, their naturally compact size is an advantage to system miniaturization and integration. Unfortunately, they are static, but traps formed from permanent magnets generally require additional homogenous bias fields to operate. These fields, when generated by Helmholtz coils can be time-dependently controlled to turn on and off the trap. The review paper by E. Hind's group discusses some of these types of traps [10], as do the References [78, 79]. Clifford Hicks, an undergraduate in our

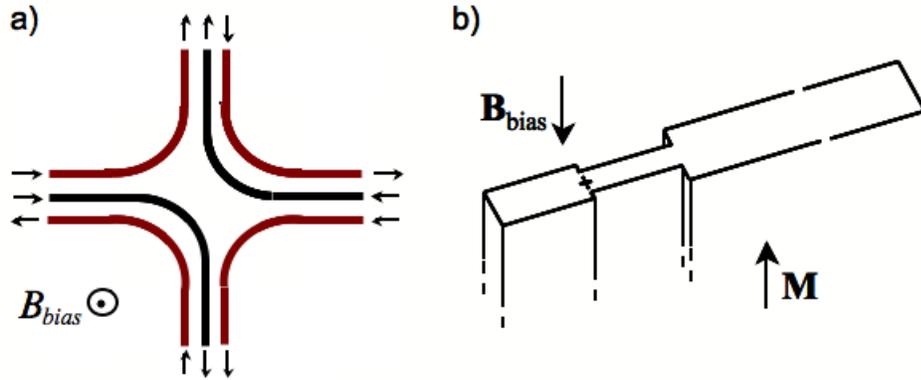


Figure 2.6: a) Modification to the Libbrecht-style Ioffe trap by Arjun Menon. This design improves the trap frequencies. b) Permanent magnet-based Ioffe trap designed by Clifford Hicks.

lab in 1999, designed a permanent magnet-based quadrupole and Ioffe trap. Pictured in Panel (b) of Figure 2.6 is a sketch of his Ioffe trap design. A bias field of a few hundred G opposes a magnetization, M , of roughly ten kG in a long and thin permanent magnet, forming a quadrupole field. The width of the magnet is modulated to add stray magnetic fields that convert the quadrupole trap into a harmonic, Ioffe potential. The achievable gradients and curvatures— 10^7 G/cm and 10^{10} G/cm²—are comparable to those achievable in the tightest microwire traps, but for less power dissipation. Chapter 8 discusses an atom mirror experiment employing a thin film of permanent magnets.

Atom chips have begun to employ electrostatic and RF fields to create novel trapping potentials for atoms. The group of J. Schmiedmayer has made great progress using both: charged conductive pads were used to split a cold cloud of atoms into a string of “sausage-link” like potentials [80]; and more recently they have used RF fields to create a double-well trap [81]. Electrostatic fields form attractive potentials for alkali atoms that can create novel traps when combined with high gradient repulsive magnetic potentials. In Chapter 9 of this thesis, we discuss the use of electrostatic fields and a magnetic atom mirror to form a 1-D ring trap. The use of RF and microwave fields on atom chips presents the possibility of incorporating atomic internal state-dependent potentials [18]. This could be of use for implementing on-chip quantum computation.

2.4 Atom chip loading

The previous sections have shown how one can use micron-scale current carrying wires to create high-gradient and curvature traps for neutral atoms, and Chapter 4 describes how to fabricate such devices. But how do we actually load atoms into these cubic micron volume traps suspended as little as a few microns above a room temperature surface? Our central challenge is to cool a gas of >300 K atoms and collect them in a manner amenable for loading into magnetic microtraps. After loading the atoms into the magnetic traps, the atom manipulation becomes more straight-forward since the atoms follow the minima of the magnetic potentials. These magnetostatic traps are easy to calculate, and the Matlab code in Appendix B.3 may be used for their simulation. The main difficulties we experience once the atoms are magnetically confined is to ensure that our tasks are performed before atoms are lost due to the trap's finite lifetime (see Chapter 5) and to have a good enough imaging set-up to be able to diagnose the trapping system and make accurate atom density measurements. We begin this section by presenting a figure and table that summarize a typical atom trap loading procedure. Each trapping step is discussed in more detail further in the section.

Figure 2.7 shows the general loading procedure: a) The atoms are loaded into the macro U-MOT for ~ 5 s. From $t = 0$ to 20 ms they are transferred to a D-MOT using the D-wire on the atom chip. This is shown in Panel (b). b) From $t = 20$ to 40 ms the D-MOT compresses the atoms and brings them closer to the surface. At $t = 40$ ms, all magnetic trapping fields are shut-off and the atoms are sub-Doppler cooled and optically pumped for 3.6 ms. All lasers are then extinguished. c) At $t = 43.6$ ms, the D-trap is turned-on and the atoms are compressed into this quadrupole trap for 5.4 ms. d) The D-trap rotation procedure begins at $t = 49$ ms. A weak bias field applied for 8 ms in \hat{y} pushes the atoms to the edge of the circular arc of the D-trap. A sinusoidal swapping of the strong bias field in \hat{x} for a strong field in \hat{y} rotates the atom cloud around the arc and into the straight wire portion that is indicated by the (d') label in Panel (g). This transfer requires at least 20 ms. In step (d''), the atoms are transferred from the D-trap to the Z-trap by ramping down the D-trap current while simultaneously ramping up the Z-trap current. This is performed under constant bias field and takes 8 ms. Finally, the atoms are released into the waveguide from the Z-trap by extending the short Z-trap to one that stretches the length of the wireguide. The gate

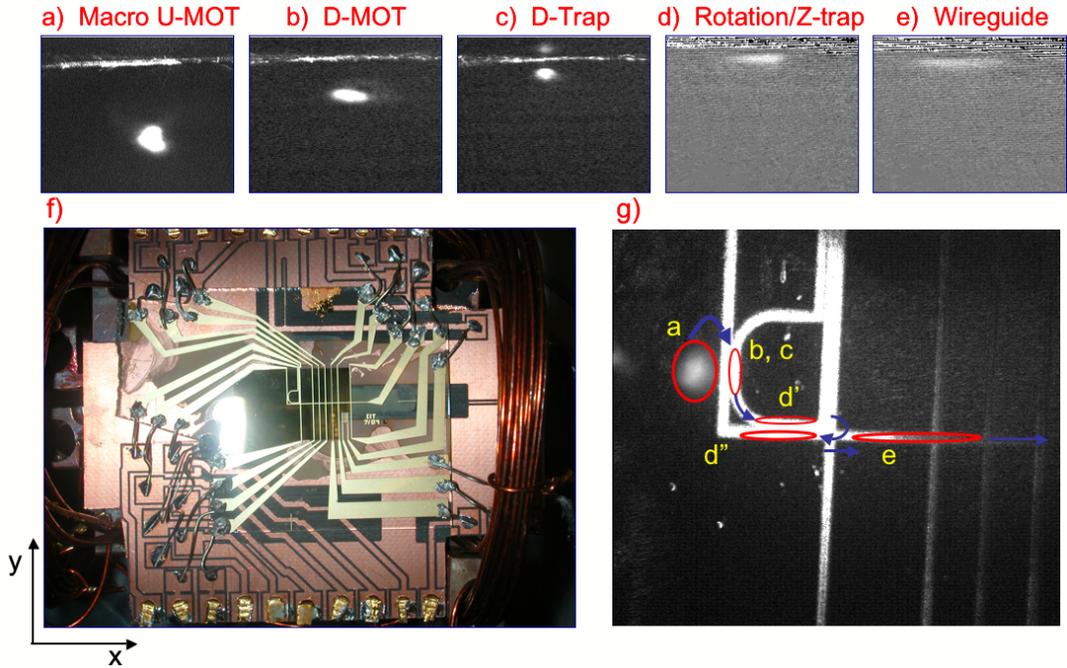


Figure 2.7: Atom chip loading procedure. Images (a) through (c) are taken with absorption imaging in the \hat{y} direction as defined by the axes in panel (f). Images (d) and (e) are taken with absorption imaging, also in the \hat{y} direction. Image (g) is a fluorescence image of the atoms from a top view of the substrate taken at the end of the macro U-MOT loading period. The red circles and blue arrows overlaid onto image (g) show the positions and transfer progression of the atom cloud. The yellow letter labels in (g) refer to the steps shown in panels (a) through (e). Steps (d') and (d'') are discussed in the text.

wires can be operated at will to govern the progress of the atom cloud or to create a dimple trap.

Table 2.1 shows the schedule of bias coil currents, microwire currents, and optical beams required to load the atom chip. Note: bias fields in \hat{x} , \hat{y} , and \hat{z} are quoted not in Gauss, but in the amps required to produce the field. As long as the same coils are used for future experiments, quoting in amps is more useful since we can measure the computer controlled currents more accurately. Use these conversions for a rough calibration of the current-to-magnetic field: $B_x = 2I$, $B_y = 2.4I$, and $B_z = 4.2I$, where I is in amps and the magnetic fields are in Gauss. The field in \hat{x} is produced by a pair of coils inside the vacuum chamber with 20 turns each. The coil radius is 3.5 cm and their separation is 7 cm. Coils external to the vacuum chamber—wrapped around the 6" diameter windows—produce the field in \hat{y} . The \hat{z} field is produced by a single coil of 5 turns and radius 4 cm placed in the vacuum

chamber and wrapped around the atom chip assembly.

atom chip trapping steps	x	y	z	U	D	Z	MOT		RP	OP	Abs
	(A)	(A)	(A)	(A)	(A)	(A)	$-\Delta$	I/I_s			
$t=[-5,5]$ s; U-MOT loading (a)	0.9	0.5	0.9	20	0	0	3.4	24	on	off	off
$t=[0,20]$ ms; U- to D-MOT	0.5	0.5	0	erf \rightarrow 0	erf \rightarrow 2.2	0	3.4	24	on	off	off
$t=[20,20]$ ms; D-MOT hold (b)	1.2	0.4	0	0	2.2	0	3.4	24	on	off	off
$t=[40,2.7]$ ms; sub-D cooling	0	0	0	0	0	0	14	1	on	off	off
$t=[42.7,0.3]$ ms; field decay	0	0	0	0	0	0	-	0	off	off	off
$t=[43,0.6]$ ms; optical pumping	2	0	0	0	0	0	-	0	off	on	off
$t=[43.6,5.4]$ ms; D-trap hold (c)	erf \rightarrow 3	0.9	0	0	erf \rightarrow 3.5	0	-	0	off	off	off
$t=[49,8]$ ms; D-trap shift	3	-1	0	0	3.5	0	-	0	off	off	off
$t=[57,20]$ ms; D-trap rotate (d')	cos \rightarrow 1	sin \rightarrow -3.5	3.7	0	3.5	0	-	0	off	off	off
$t=[77,8]$ ms; D- to Z-trap (d'')	1	-3.5	3.7	0	erf \rightarrow 0	erf \rightarrow 3.5	-	0	off	off	off
$t=[85,100]$ ms; Z-trap (e)	1	-3.5	3.7	0	0	3.5	-	0	off	off	off
$t=[185,10]$ ms; free fall	0	0	0	0	0	0	-	0	off	off	off
$t=[195,40]$ ms; fluorescence	0	0	0	0	0	0	3.4	24	on	off	off
$t=[195,40]$ ms; absorption	0	4.5	0	0	0	0	-	0	on	off	on
$t=[235,1]$ ms; reset	0.9	0.5	0.9	20	0	0	3.4	24	on	off	off

Table 2.1: Schedule of magnetic and optical fields for atom chip loading and trapping. The time for each step is quoted as “ $t = [\alpha, \beta]$,” where α is the start time and β is the step duration. The letters in parenthesis label the steps depicted in the pictures in Figure 2.7. The columns are (from left to right): the current in the coils providing the bias field in \hat{x} , \hat{y} , and \hat{z} ; the current in the macro U-wire; the current in the chip-based D-wire; the current in the chip-based Z-wire; the detuning, Δ , of the MOT beams in units of the transition linewidth, $\gamma = 2\pi \cdot 5.2$ MHz; the intensity of MOT beams with respect to the saturation intensity for this transition, $I_s = 1.1$ mW/cm²; the status of the repumper (RP), optical pumping (OP), and absorption imaging (Abs) beams. The orientation of the bias field axes are shown in Figure 2.7. The $+\hat{z}$ axis is pointing away from the plane of the substrate. For the bias fields, see text for the coil current-to-field conversions. The entries “erf \rightarrow ,” “sin \rightarrow ,” and “cos \rightarrow ” indicate the function used to change the current in the time β to the quoted final value.

2.4.1 Macro U-MOT

Magnetic microtraps are typically no deeper than 1 mK in energy, and our first task is always to cool the diffuse gas of cesium atoms in the vacuum chamber from the temperature that they emitted from the source oven (> 300 K) to a temperature below the trap depth. (Atom trapping and cooling is most commonly performed with alkali atoms, and cesium is our specific atom of choice².) We also want to collect these cooled atoms in a compact volume for efficient funneling into ever tighter magnetic traps. The MOT can accomplish both tasks: the atoms are cooled to 100 to 200 μ K and collected in a ball a few millimeters in diameter. The MOT works by using red-detuned laser beams and a quadrupole magnetic field to exert a damped, harmonic force on each atom traversing the trapping zone [66]. The lasers impinge upon the trapping zone from all six orthogonal directions and exert a damping, Doppler force on the atoms [46]. For cesium, these lasers are red-detuned by a few

²D. Steck has compiled an excellent set of notes containing information regarding the properties of cesium and other trappable alkali atoms [82].

linewidths ($\gamma = 2\pi \cdot 5.2$ MHz) from the $F=4$ to $F'=5$ hyperfine transition and are circularly polarized. A repumping laser beam is necessary to prevent the shelving of atoms into the lower, dark hyperfine state. This laser is tuned to the $F=3$ to $F'=4$ hyperfine transition and is introduced to the trap region most easily by overlapping its beam with one or more of the MOT beams.

For the MOT beams, the two along the axis of the anti-Helmholtz field are of opposite circular polarization to the four in the plane perpendicular to the axis (i.e. the plane parallel to the face of the wire coils). In practice, three of the six independent laser beams can be replaced by retroreflecting the remaining three beams. This automatically produces the correct polarization for each pair of laser beams if there is a quarter-wave plate in front of the retroreflection mirror (this plate's orientation is unimportant). A good strategy for getting the correct combination for the orientation of the laser beam polarization and quadrupole field is to first set the circular polarization orientation of the three retroreflected beams correctly with respect to the orientation of the anti-Helmholtz coils as mention above. There are two correct relative polarization possibilities. Pick one orientation and the sign (orientation) of the quadrupole field will at most need to be flipped once to attain a MOT. Of course, one needs to be careful that ambient magnetic fields do not push the center of the quadrupole field out of the center of the intersection of the lasers. It might be necessary to manually adjust the field compensation coils to ensure that the quadrupole field is centered in the optical field.

For atom chip experiments, we have an extra difficulty in that the chip itself blocks laser beam access to a would-be MOT positioned a few millimeters above the chip's surface. This problem is solved using the technique of the mirror MOT [2], which uses the chip surface as an optical mirror to reflect laser beams in two of the three cardinal axes. Figure 2.8 illustrates this scheme. The mirror eliminates the need for beams to traverse the lower half-space: four of the six beams are formed by reflecting two beams off of the mirror at 45° , and the other two are formed by grazing horizontal beams above the atom chip and perpendicular to the plane formed by the other four beams. Conveniently, the polarization of the reflected beams off the chip's mirror is maintained in the correct orientation. Note that the axes of the quadrupole field must be rotated 45° with respect to the surface of the atom chip. This field may be created by anti-Helmholtz coils outside the vacuum chamber, but a more compact and integrated solution is to use the field created by a macro U-trap

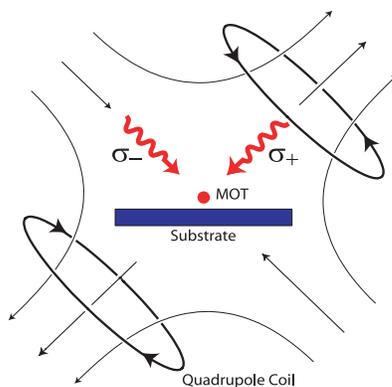


Figure 2.8: Diagram of the mirror MOT experimental set-up. A quadrupole field, two 45° laser beams, and one retroreflected grazing beam (perpendicular to the page and not shown) form a MOT 1.5 to 4 mm above the substrate. The polarization orientation of the horizontal grazing beam is the same as the 45° beam parallel to the plane of the coils (i.e. the beam into the page is σ_- polarized).

located inside the vacuum chamber [64]. As shown in the vector field plot in Figure 2.3 (a), the axes of the quadrupole field from a U-trap are also rotated 45° from the plane of the wire substrate. This is fortunate since the mirror MOT can be formed by simply using the U-trap’s quadrupole field which eliminates the need for cumbersome coils of wire draped around the chamber at 45° . We will denote the mirror MOT formed with a U-trap as a “U-MOT,” and if a macroscopic sized-wire is used for the U, we will designate it a “macro U-MOT” to distinguish it from microwire-based U-MOTs.

The macro U-MOT uses a block of copper milled into a U-shape and the dimensions of the base of the U are $10 \times 5 \times 5$ mm and the sides are 2 mm wide and 5 mm thick. We typically run 20 to 30 A of current through the copper U-shaped block, and when combined with a few Gauss bias field, a U-trap of a gradient of a few G/cm is formed in the tightly confining directions. This is sufficient to create the macro U-MOT from a vapor of cesium atoms. A macro U-MOT is easier to form than a mirror MOT generated by external quadrupole fields for three reasons: 1) The tip of the U-wire—residing inside the vacuum chamber—is pre-positioned less than a centimeter from the intended trap minimum, and this makes it easy to ensure that the quadrupole field center is within the laser intersection volume. In contrast, a pair of 10 cm scale coils positioned outside the vacuum chamber must be accurately aligned in three dimensions to position its trap center within the one cubic centimeter laser trapping zone. 2) The anisotropy of the macro U-MOT’s quadrupole field makes zeroing the field in

the weakly-confining direction simple. Even small bias fields can displace the MOT along the base of the U and the field is zeroed in this direction when the MOT atoms reside in the center of the U. Of course, this forms an easy control of the MOT's position, and we did not notice a drastic decrease in MOT population when the atoms were displaced less than approximately 3 mm from the center of the 10 mm wide base of the U-shaped wire. We used this ability to shift the MOT over a ~ 6 mm distance to load the cooled atoms into two distinct microwire traps fabricated on the same chip (i.e. cold atoms could be fed into two separate microtrap networks). 3) The U-wire base forms a nice guide-to-the-eye for aligning the horizontal MOT beam parallel to this straight edge. We have found that the macro U-MOT's population is highly sensitive to the correct alignment of this laser beam to the weak axis of the U-trap's quadrupole field as well as to the alignment of the two horizontal beams with each other. The horizontal beam should not be retroreflected since it is deformed after grazing the chip surface (the chip edge should clip the beam at approximately one third of its diameter to ensure that the field maximum is only a few millimeters from the substrate surface), but rather two independent, counterpropagating beams should be overlapped. The MOT population is much more sensitive to this alignment than to that of the two 45° beams which we found can—and should be—misaligned. We have found that the macro U-MOT is spatially unstable and contains fewer atoms—by an order of magnitude—when these two 45° beams are perfectly overlapping. Misaligning them by a few degrees corrects this problem, and this holds true for normal mirror MOTs as well. A good alignment procedure is to overlap the reflected 45° beams and then slowly walk one off the other until the trap is stable and the atom number is maximized. We found that misaligning the horizontal beam pair is only detrimental to the performance of the MOT.

The atom chip substrate is attached to the top of a 0.8 mm thick copper-clad teflon circuit board, which is in turn attached to the top of the U-shaped copper block. The sapphire atom chip substrate is 1 mm thick, but since the macro U-MOT can be formed up to ~ 4 mm above the top of the U-shaped copper block, there is plenty of room for these non-magnetic materials to slip in between. The teflon circuit board—which is relatively vacuum safe and obtained from a microwave circuit company—is primarily used as a support structure for the electrical feedthrough pads connecting the atom chip's microwires to macroscopic pins. Additionally, this teflon circuit board has a $70 \mu\text{m}$ thick cladding that can be circuit-milled into mesoscopic U-traps and Z-traps. We formed a U-MOT and

U-traps with such a device. The maximum supported current is nearly 7 A for 1 s pulses and the minimum wire width—limited by the drill bit accuracy—is approximately 250 μm .

Figure 2.7 shows the microtrap loading sequence used in our gated waveguide experiments for cavity QED on a chip. The goal is to perform all the atom trapping and cooling in one area of the chip, transport the atoms to a remote region of the chip where a microcavity will not obscure any of the trapping or imaging laser beams, and then magnetically deliver the atoms to the mode of the cavity. The sequence begins with a macro U-MOT. The atom chip hangs upside-down in the vacuum chamber to allow time-of-flight temperature measurements to be taken. This can be performed by either fluorescence or absorption imaging and typically a drop of delay 5 to 20 ms is required. If the chip were right-side-up, the atoms would crash into the chip surface in this time. Having the atoms drop away from the substrate also improves detection image quality since there is less scattered light from the chip surface. Panel (a) of Figure 2.7 shows a fluorescence image of the macro U-MOT beneath the substrate. The horizontal line at the top is scattered MOT light—used for fluorescence imaging—off of defects on the substrate surface. There are typically 10^6 cesium atoms in the macro U-MOT, and it is located 2 mm below the substrate. The cloud width is roughly 1 mm in the tightly confined directions and 2 mm in the weakly confined direction. The image in Panel (a) is taken in the \hat{y} direction as defined by the axes next to the image of the atom chip displayed in Panel (f). This is the chip used for all the images in this figure. The macro U-MOT is loaded from vapor for ~ 5 s. The bias fields in \hat{x} , \hat{y} , and \hat{z} are used to maneuver the macro U-MOT to the tip of the D-shaped microwire on the substrate. This is shown in Panel (g) which is the top view of the atom chip in Panel (f). This image is also taken in fluorescence and the 45° MOT beams illuminate the rough electroplated gold on top of the microwires. The white cloud in the red circle labeled “(a)” is the macro U-MOT photographed at the end of the 5 s loading period which we will define as $t = 0$. The macro U-MOT positioned at the tip of the D-shaped wire is 2 mm above the surface.

We have two sets of biasing coils for each cardinal direction. One set nulls the ambient magnetic field and is driven with an unchanging current. The current in the other set is computer controlled and is used throughout the experiment to either supply shimming fields for adjusting trap positions or to supply the larger bias fields necessary for forming microwire traps. Typical bias fields in \hat{y} and \hat{z} used to align the macro U-MOT above the D-trap are listed in the $t=[-5,5]$ s row of Table 2.1.

2.4.2 D-MOT

At $t = 0$, the macro U-MOT is brought closer to the substrate surface so that the atoms may be trapped with the atom chip’s microwires. To perform the transfer, the current in the macro U-wire is reduced to zero in 20 ms with the smooth “erf” function (see Equation 2.7). The bias field in \hat{x} is held constant (or only reduced slightly) to maintain a mirror MOT as the decreasing wire current brings the trap center closer to the surface. To ensure the current is completely off, an electrical relay disconnects the wire circuit at $t = 20$ ms. This prevents day-to-day offsets in the U-wire supply from disturbing the microtraps. The use of a relay is generally a good idea if the atoms will be trapped near an unused wire that is connected to a power supply. In making BECs, current fluctuations in unused wires can heat the atoms out of the trap during the long evaporative cooling sequence. In practice, the power supply for the U-shaped copper block (PowerTen 60 V/55 A) takes ~ 10 ms to respond to control signals and the current is not completely zero by the time the relay disconnects the (low-inductance) circuit.

During the macro U-MOT trap lowering, the atoms are transferred from the macro U-MOT to a mirror MOT generated by the D-shaped wire on the chip surface. We call this a D-MOT, and this is used instead of a P-shaped wire to increase the trap volume. The current in the D-wire is increased to 2.2 A using an erf function, and in 20 ms the MOT of atoms are gradually shifted from the starting position of the macro U-MOT to the position of the D-MOT defined by the D-wire current and the bias field in \hat{x} . The MOT laser fields force the atoms to make this spatial shift in a short period of time and is a nice technique for transferring the atoms without having to ensure that the macro U-MOT is perfectly aligned with the D-MOT, or in practice, that the centimeter-sized U-shaped copper block is accurately positioned with respect to the sub-millimeter sized D-wire. This allows us to align the two by eye during the gluing process without much worry. We have demonstrated this transfer with 50% efficiency: the macro U-MOT typically contains 1.5×10^6 atoms and the D-MOT contains 7.5×10^5 atoms. This efficiency could perhaps be improved by using a wider D-wire.

During the step from $t = 20$ ms to 40 ms, the atoms are allowed to settle in the D-MOT, and we bring the atoms closer to the surface by increasing the \hat{x} bias current to 1.2 A. The D-MOT trap minimum is 1 to 2 mm from the surface with a trap gradient of 13 G/cm in

the tightly confining directions \hat{x} and \hat{z} .

2.4.3 Sub-doppler cooling and optical pumping

At $t = 40$ ms, the atoms are prepared for loading into the magnetostatic D-trap. The atoms need to be further cooled and optically pumped to a strong magnetically trappable state. In the MOT, the cesium atoms are at the Doppler-limited temperature of $124 \mu\text{K}$. To achieve high capture efficiencies in the magnetic microtrap, the atoms are sub-Doppler cooled to $\sim 10 \mu\text{K}$ in 2.6 ms (this exact time is not crucial, but should be longer than about 2 ms and shorter than 5 ms). This is accomplished by red-detuning the MOT lasers and turning-off and zeroing the magnetic fields to the mG level or better. This latter requirement is quite tricky to accomplish and we now describe our technique for doing this in the presence of an atom chip. In a normal free-space MOT, one can zero the fields by seeing how the atoms expand in an optical molasses after the fields are shut-off. If there is a remanent field, the atoms will preferentially expand in that field's direction. One can then adjust the nulling fields in \hat{x} , \hat{y} , and \hat{z} until the atom cloud expands slowly and symmetrically. With the presence of the atom chip only a millimeter below the cloud, this expansion becomes more difficult to detect. Moreover, defects on the mirror surface imprint anti-trapping zones onto the the mirror MOT region. The atom cloud expands in strange ways through these zones which makes the field zeroing difficult. It is important to have as smooth a mirror as possible and to ensure that the MOT beams are of equal intensity before attempting field zeroing (unequal beam intensity can mimic a residual magnetic field due to radiation pressure). By carefully observing the expanding cloud in three dimensions simultaneously (good optical access is required), the fields can be zeroed even in the presence of the atom chip. Other groups have used the Hanle effect to zero the magnetic field. While more sensitive to non-zeroed fields, this is more difficult to perform. The magnetic field nulling can be further optimized by measuring the cloud temperature from free-fall expansions: for a given fall time, a better field zeroing will increase the cloud density which signifies a lower atom cloud temperature.

After sub-Doppler cooling, the system is allowed a short time—0.3 ms—for all fields to decay before the optical pumping process begins. Since the D-trap—turned-on post-optical pumping—will require a strong bias field in the \hat{x} direction, we perform the optical pumping with the quantization axis in \hat{x} to ensure the atoms are not jostled by too many shifting

fields. The optical pumping beam is directed at the cold atom cloud in \hat{x} along a 4 G bias field for approximately 0.6 ms. There are two Zeeman ground states that the atoms can be pumped into for (optimal) confinement in magnetic traps: the $F=4, m_f=4$ state and the lower energy $F=3, m_f=-3$ state. For the upper hyperfine state, a σ_+ beam tuned to the $F=4$ to $F'=5$ transition combined with the repumper will shelve the atoms in the $F=4, m_f=4$ state, while for the lower hyperfine state, the repumper combined with a σ_- beam tuned to the $F=3$ to $F'=2$ transition will pump them to the $F=3, m_f=-3$ state. With respect to eliminating hyperfine changing collisions, it is best to pump the atoms to the $F=3$ ground state (see Chapter 5).

2.4.4 D-trap and rotation

The sub-Doppler cooled and optically pumped atoms are recaptured in a magnetostatic trap using the D-shaped wire on the atom chip. The intervening, non-trapping steps only last ~ 3.6 ms which is too short a time for the cloud to have expanded or dropped significantly. Since the atoms have already been confined in a trap that uses the D-wire (i.e. the D-MOT), the atoms are already spatially mode matched to the D-trap, which is one of the advantages of using the D-MOT. The Munich group uses a slightly different microtrap loading procedure that skips the D-MOT (or in their case, a P-MOT). This alternative procedure requires accurately shifting the macro U-MOT (or externally generated mirror MOT) into the position of the P-trap. The atoms are then sub-Doppler cooled, optically pumped, and finally transferred directly to the magnetostatic P-trap without using a chip-based MOT. An additional step is added before the sub-Doppler cooling that compresses the atom cloud: the MOT lasers are detuned and these beams and the repumping beam are decreased in intensity while the gradient of the quadrupole field of the MOT is increased. This step could be added to our procedure to see whether it helps to improve the D-trap capture efficiency. In our current procedure, the D-trap collects 40% (roughly 6×10^5) of the original macro U-MOT's atoms.

The atoms are allowed to acclimate in the D-trap for a few milliseconds before the rotation begins. The D-trap is formed with a 6 G field and 3.5 A wire current. This is the maximum current we can safely pulse through uncooled wires that are $100 \mu\text{m}$ wide and $5 \mu\text{m}$ tall. More recent atom chips have wires that are $12 \mu\text{m}$ tall and will safely support larger currents. The trap minimum is 0.5 to 1 mm above the surface and the D-trap has a

gradient of ~ 50 G/cm—sufficient to support the atoms against gravity which exerts a force equivalent to 23.3 G/cm on the cesium atom. The atoms are pushed to the start of the curve of the D-trap with a 2.4 G field in \hat{y} for 8 ms before the rotation begins. The rotation can be accomplished in as little as 20 ms, but slower field ramps—up to 200 ms—minimize cloud heating and increase the number of atoms remaining in the trap. However, in a chamber with higher than optimal vacuum pressure decreases trap lifetime and a quicker rotation leaves more atoms to conduct the experiment. For high vacuum pressures there is an optimal rotation time to minimize cloud heating while minimizing background collisional loss. For short rotation times, we measured that the atoms are heated to 50-70 μK after the Z-trap transfer. Alkali atoms other than cesium can be evaporatively re-cooled down to the 10 μK level or lower³. The field in \hat{x} is not wholly extinguished at the end of the rotation since this field helps the shift the atoms into the correct position for transferring into the Z-trap. The trap gradient is increased to > 100 G/cm before the next transferring steps. For atom waveguiding, we found that with the wire currents restricted to 3.5 A, the atoms need to be less than 700 μm from the chip surface to maintain high enough trap gradients.

2.4.5 Z-trap transfer

The final step is to transfer the atoms from the rotated D-trap to the Z-trap. The microwires for each traps are only located a few tens of microns from one another. The traps formed from the D-wire and the Z-wire can therefore be made to coincide if the the trap gradients are small enough. The atoms are transferred from one to the other by slowly swapping the microwire fields. This is most simply accomplished by maintaining the same bias field in \hat{y} while exchanging the current in the two wires. We can transfer the atoms with nearly perfect efficiency in 8 ms. This procedure exchanges the atoms between two independent wire networks and can be used more generally. The efficiency of the entire transfer from the macro U-MOT to the Z-trap is 33% ($\sim 5 \times 10^5$ atoms remaining). Once the atoms are in the Z-trap, we are free to transport the atoms down the waveguide at will. Alternatively, we could transfer the atoms to an H-trap instead of a Z-trap. In this case, we have found that the side wires should have a current of 2 A while the guidewire should have a current of 3.5 A. In either case, an \hat{x} bias field of 2 G increases the oscillation frequency of the Ioffe trap in \hat{x} .

³See Chapter 5 for a discussion on the collisional peculiarities of cesium.

2.4.6 Imaging

After manipulating the atom cloud with the microwires, one generally wants to image the atoms to diagnose the cloud’s position, atom number, and density distribution. Imaging an atomic cloud near the chip surface is difficult, and one usually needs to drop the cloud if the atoms are initially trapped closer than $100\ \mu\text{m}$ from the surface. Imaging can occur within 0.5 to 1 ms of the cloud drop and this time is simply limited by the need for the magnetic trapping fields to decay. The longer the delay time, the larger the cloud expansion. The rate of the cloud expansion is an indicator of the temperature of the cloud, and measuring this is crucial for optimizing the magnetic field zeroing during the sub-Doppler cooling stage and for minimizing trap heating during cloud transport.

There are two methods for imaging atoms and both of which are listed in Table 2.1. Fluorescence imaging uses the light scattered off the atoms into 4π to capture a picture of the atom cloud. A CCD camera placed behind a high numerical aperture lens system collects the scattered light. The light from the MOT is typically used to excite the atoms. Shifting the MOT light into resonance will enhance the image contrast since more photons are scattered during the same CCD integration time. We typically integrate the image for 1 to 2 ms—anything longer produces a streaked image as the cloud begins to fall and expand. We fluorescence image in all three directions which is very important for performing quick experimental diagnostics.

Absorption imaging is an alternative method for imaging atom clouds. While fluorescence imaging is easy to perform, it is very difficult to obtain accurate atom number measurements. This is due to the uncertainty in calibrating the scattered photon rate and collection efficiency: the total laser power and detuning, the lens system collection efficiency, and the CCD camera’s quantum efficiency and photon count-to-voltage calibrations must all be accurately measured. In contrast, all of these calibrations drop-out for absorption imaging. In this case, the CCD camera images the shadow that the atom cloud imprints onto a single, collimated, on-resonant laser beam. The trade-off is that a second, computer controlled beam is required and clear optical access must be reserved for the laser to pass unobstructed into the vacuum chamber, across the chip surface, and out of the chamber to magnifying lenses and into the CCD camera. Furthermore, a bias field should be maintained parallel to the circularly polarized light to provide a quantization axis. This maximizes the

number of absorbed photons. We can absorption image only in \hat{y} , and an 11 G field in \hat{y} is used during the 0.5 ms integration time. It is important to use an absorption beam on resonance and with an intensity much less than I_s to eliminate the detuning and intensity calibration from the atom number estimation. The optical density (OD) per pixel is:

$$\text{OD/pix} = \frac{\sigma_0 N_p}{P_x P_z}, \quad (2.8)$$

where σ_0 is the absorption cross section, $\sigma_0 = (3/2)\lambda^2/\pi$, N_p is the number of atoms in that pixel, and P_x and P_z are the pixel dimensions in \hat{x} and \hat{z} for an image in the \hat{x} - \hat{z} plane. The only calibration required is to measure the spatial area of a rectangular pixel. This can easily be done by taking an image of a finely ruled measuring stick placed at the same focal point as the atoms, and counting the number of pixels per ruler tick provides this calibration. The camera has to be temporarily moved to place the ruler at the same distance from the CCD camera as the atoms in the vacuum chamber. When using analog CCD cameras, we allot 40 ms after the first image to acquire a second image without the atoms present (this 40 ms incorporates the time it takes for the atoms to move out of the image and the ~ 30 Hz repetition rate of the analog camera). We divide the two images to obtain the OD. For more accuracy, a background image should be subtracted from these images before dividing. Note: one should make sure that the images are not saturated—CCD camera gain is usually unnecessary—and the bias field should be optimized to ensure a maximal absorption rate.

Chapter 3

Experimental Details

This chapter discusses experimental details of our atom chip trapping system. We present how and why things are assembled, where to buy certain specialized equipment, and specific do's and don'ts of the experiment.

3.1 Vacuum system

Our vacuum system is UHV compatible—with ConFlat-style gasket fittings—and our experiment usually runs in the high 10^{-9} Torr range. This pressure is higher than what the chamber and pumps could normally achieve and is limited by material forming the atom chip that is not fully vacuum safe. The main part of the chamber is a custom-made 6" diameter "Spherical Square" from Kimball Physics. It is similar to part number MCF600-SS200408-A but with two of the four pairs of mini-ConFlat ports (1.33" diameter) replaced with extended 2.75" diameter ConFlat ports. The chamber is shown in Figure 3.1: Panel (a) shows the bare chamber without the two 6" windows and the atom chip; and Panel (b) shows the sealed chamber with the optics arranged for atom trapping and cooling. The two 6" window ports are used for optical detection, the horizontal MOT beams, and manual access enabling installation of the atom chip assembly. The two lower 2.75" windows—angled at 45° —are used for the 45° mirror MOT beams, and the upper two 45° 2.75" ports are connected to electrical feedthroughs. One of these feedthroughs has a two-pin, high-current feedthrough from Huntington Labs which we use for supplying the 20 to 30 A macro U-MOT current. The other upper 45° 2.75" port has a 4" ConFlat expander that connects to a 25-pin, low-current electrical feedthrough from MDC. Current pulses < 100 ms and < 4 A are possible with this 25-pin feedthrough. Two horizontally situated mini-ConFlat

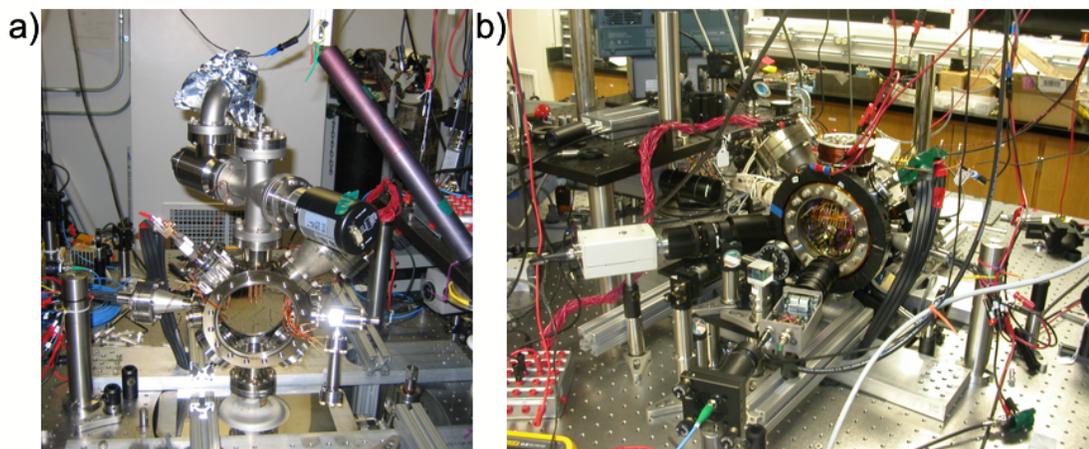


Figure 3.1: a) Open vacuum chamber without atom chip installed. b) Vacuum chamber and trapping optics.

ports are connected to viewports that are used for imaging and optical pumping, and the other two are used for the output of the cesium oven and an 8-pin electrical feedthrough. This latter electrical feedthrough is from MDC (ISI) and is very fragile. It is attached to a 2.75" ConFlat expander to ensure intra-chamber pin contacts do not short. We found that small leaks in the electrical feedthrough pins can be plugged by applying a drop of Epotek 353 epoxy to the base of the offending pin and curing it at 80° using heater tape wrapped around the flange.

A note about vacuum part supply companies. While MDC (ISI) has many products and can ship quickly, we have found that their electrical feedthroughs and mini-ConFlat windows are prone to leak. Varian and Kurt J. Lesker's seem better, but we should explore other companies for electrical feedthrough products. The Kimble group has reported that MDC's UHV valves can seize. We now buy right-angle valves from VAT even though they are more expensive. Kevin Birnbaum's thesis [83] contains more vacuum chamber lore and a good description of vacuum part cleaning procedures¹.

The intra-chamber electrical wires are Kapton coated and are purchased from MDC. We use screw-on connecting pins from MDC to make most of the contacts, though the contacts to the atom chip are made with crimp-on pins which reduce the amount of solder in the chamber. We do not use any low-temperature flux-based solder in the chamber. A

¹The cleaning and assembly procedures used in our lab are not quite the same, but the differences do not warrant description here other than to say that some of the procedures in Birnbaum's thesis are more thorough than what is required for our non-BEC atom chip experiments.

few solder joints on the atom chip use high-temperature silver solder and flux-less ultra-sounding solder—both of which are more vacuum compatible.

We use an oven to create the background vapor of cesium atoms in the chamber. An alkali dispenser could be used for future, lower vacuum pressure experiments, but for now the oven has worked well (see Chapter 10 for more information regarding dispenser use). The cesium oven consists of no more than a cracked cesium ampule in a valved-off section of metal bellows wrapped with heater-tape. One gram cesium ampules can be purchased from Aldrich for less than \$60. The glass ampule comes filled with nitrogen and with a score in the glass for easy cracking. Before placing in the oven, the ampule should be acetone-cleaned in an ultrasound bath and rinsed with methanol like most other vacuum parts. Cesium is flammable when exposed to water and should never be exposed to air. On the rare occasion that the cesium oven must be opened to air, one should be careful not to knock or bang the metal: in air, cesium will form a thin layer of oxide that protects the interior from bursting into flames, and this layer should remain intact until the chamber is resealed or the cesium is disposed of in a controlled manner.

A simple oven is formed from a mini-UHV right-angle valve from Varian that connects the vacuum chamber to a small-radius, 6" long cylindrical bellows. This bellows may be purchased from Varian and should be capped on one end with a mini-ConFlat blank and not a window—the cesium and the repeated heating will destroy the window over time. The bellows from Varian is just slightly wider than the glass ampule, and one may crack-open the ampule under vacuum by simply giving the bellows a quick snap with the hand. This should be performed only under slight vacuum since the cracked ampule will immediately release its nitrogen. Running a turbopump in a low-speed, safely mode will ensure that the chamber pressure is low enough to not introduce water to the newly-exposed cesium while preventing the nitrogen from damaging a faster rotating turbopump. Cesium melts just above room temperature, and to introduce cesium vapor into the chamber, the oven should be heated to 50° to 80° C with heater-tape. Bands of heater-tape can be purchased from Cole-Palmer and are plugged directly into variacs.

The metal of the oven should be heated evenly to prevent the cesium from collecting on one patch of surface instead of entering the chamber. (Some groups purposely chill one area of the cesium oven to have more control over its vapor pressure or to collect the cesium back in to the oven after the day's experimental run is finished.) We found that

a new oven—or one that has not been used in a while—needs to be heated for a day or more before a significant vapor pressure of cesium can be detected in the main chamber. We presume this is due to the need for the cesium to migrate through the bellows and to the valve aperture. Inevitably, the main chamber will have to be opened to air during the lifetime of the cesium ampule, which can be several years at normal usage. (In our atom chip experiments, the chamber is vented every few months to swap-in a new device.) The oven valve should be shut during this chamber venting. The pressure in the oven will still rise, but will not fully come-up to air pressure and the cesium will mostly be isolated from water. We noticed, however, that after pumping-down the chamber, the oven needs to be baked again for a day or more. We assume this is due to the need to bake-off a thin film of oxide that forms on the cesium due to the small amount of water that leaks through the valve while the main chamber is exposed to air. During atom trapping experiments, the oven valve may either be left constantly opened or be periodically opened and closed to release pulses of cesium vapor. The time constant for cesium vapor pressure decay in the chamber is approximately one hour. One can crudely control the cesium vapor pressure by adjusting the oven temperature. Sometimes an imperceptibly small bead of cesium migrates to the main-chamber side of the oven valve. In this case, one will notice that it is unnecessary to open the valve to obtain a usable cesium vapor pressure in the chamber. Rather, one only has to heat the closed oven. Cesium vapor pressure in the chamber can be detected by scanning the MOT lasers over a hyperfine resonance at a few Hz. The chamber will flicker in the NIR at the location of the laser beams when there is sufficient cesium.

The top 2.75" ConFlat flange of the chamber is attached to a four-way cross that connects a Varian Eyesys Mini-IMG vacuum gauge, a VAT valve for connecting extra pumps, and a fiber optics feedthrough. This feedthrough is of the design made by E. Cornell's group at JILA [84], and we use a double-hole Teflon ferrule to provide for the fiber input and output. The bottom 2.75" ConFlat flange is connected to a Varian VacIon Plus StarCell 40 ion pump which is itself connected via a VAT valve to a Pfeiffer TMU 071 P turbopump backed by a Pfeiffer diaphragm pump. For space efficiency, the ion pump and turbopump are located underneath the chamber in a 2' hole custom-bored into the laser table. There is a blanked Kwik-Flange connected to the turbopump that is used to flush the system with nitrogen during chamber venting to minimize the amount of water coating the inner chamber surface. The chamber is pumped-down in the following manner. With the VAT valve

to the ion pump open, the diaphragm pump is turned-on² for 10 min before ramping-up the turbopump. The Pfeiffer turbopump should ramp-up to full speed, 1500 Hz, in a few minutes³. The turbopump must be water-cooled. Make sure the Neslab chillers have been refilled with water—this should be checked every 6 months—and a small fan should be used on the turbopump if the chamber is being baked. As the turbopump turns-on and starts pumping down, the Eyesys Mini-IMG vacuum gauge voltage will begin to fall from around 8 V (10^{-2} Torr) to around 6 V (10^{-4} Torr) in 12 hours⁴. This gauge reads non-linearly from atmosphere to 10^{-2} Torr, but nonetheless provides a decent pressure guide since its output voltage monotonically increases from 4 V to 8 V as the pressure approaches 10^{-2} Torr.

A bake-out is required when the atom chip assembly is in the chamber, and the ion pump will not bring the pressure down into the low 10^{-8} Torr otherwise. One should attempt to turn-on the ion pump before the bake-out begins just to see if there is a vacuum leak. Spraying methanol around leak-prone areas of the chamber will cause the ion pump current—which is proportional to chamber pressure—to spike upwards. To turn-on an ion pump, it is good practice to briefly flip the power switch on a few times. This is presumed to knock-off dirt from the ion pump’s magnets. The ion pump has a voltage monitor that is proportional to the ion pump current (1 V = 1 mA). When first turning-on the ion pump, the monitor voltage is around 12 V. When the chamber pressure is low enough for the ion pump to start working, this voltage hovers around 12 V for a few minutes before quickly shooting downwards. Under normal operating conditions, the voltage will reach 2 to 3 V in around 10 to 20 minutes. Without a bake-out the pressure might not fall too far below this, but achieving this monitor voltage means that the chamber does not have a large leak. The chamber should be tested again with methanol to see if there are any smaller leaks. Of particular concern is the fiber optic feedthrough. The Teflon ferrule will most likely need to be tightened periodically throughout the bake-out process and the methanol test aids in gauging when and by how much to tighten (over-tightening can stress the fiber). During pump-down and bake-out, the cesium oven should be opened periodically to allow its pressure to equilibrate with the main chamber. However, one should make sure that

²When using a fiber taper in the chamber (see Chapter 6), the VAT valve should be opened slowly to prevent the diaphragm pump from creating too large of a pressure differential.

³If unused for more than a month, the first attempt at ramp-up will fail. The turbopump will resume normal operation after stopping and restarting the ramp.

⁴To convert to Torr: pressure in Torr = 10^{V-10} , where V is the Eyesys Mini-IMG output voltage.

the valve is not left open during bake-out which would allow too much cesium to escape the oven. The oven valve should be opened slowly to prevent the pressure differential from breaking a fiber taper.

The bake-out can begin once the chamber has been leak-checked. The chamber is wrapped in oil-less aluminum foil, encircled as evenly as possible with heater-tape, and wrapped once again with aluminum foil. The ion pump has its own built-in heater. These heaters are controlled and powered with variacs and the temperature should be slowly raised over the course of a day. The final temperature—measured by thermistors external to the chamber, on the atom chip, or on both—should be as high as the chamber material will stand to efficiently bake-off water and other chamber contaminants. The honeycomb of the laser table can withstand up to 80° C, but we have been able to heat the chamber and ion pump to nearly 200° C without noticing any deformation. With the atom chip, the chamber’s outside temperature should be no higher than 150° C. The ion pump should be shut-off during the first few hours (or overnight) of the bake-out since the chamber pressure will rise higher than the ion pump’s safe operating limit. Subsequently, the pump should be turned-on during the bake to assist the turbopump in evacuating the chamber. Within the first hour the ion pump’s monitoring voltage should drop from 12 V to 1 V and stabilize at a few hundred mV during the bake. The chamber should be baked for at least a day before slowly turning-off the heat. As the bake comes to an end, the ion pump’s voltage should quickly drop to the sub-100 mV level and continue to drop until it reaches a base pressure (quoted in volts) less than 25 mV. It is important to open and close the cesium oven valve during this process as well. Once the VAT valve and the turbopump are closed, the pressure should drop slightly. Over the next few days, the pressure will rise and fall as the chamber “burps” (i.e. as pockets of gas are expelled). Eventually the chamber should reach a good operating pressure in the sub-10 mV range. The Varian ion pump manual contains a voltage-to-pressure conversion. The voltage readings at the pressure decades are: 1 V \rightarrow 10⁻⁶ Torr; 60 mV \rightarrow 10⁻⁷ Torr; 4 mV \rightarrow 10⁻⁸ Torr; 0.2 mV \rightarrow 10⁻⁹ Torr. We are able to form a MOT at pressures as high as 30 mV, but operation around 4 to 15 mV is better. We do not have an ion pressure gauge installed in the experiment: we have used a nude UHV-24 Varian gauge in the past but found it to be unreliable. The Varian ion gauges in glass bulbs, though ungainly, seem to be more robust.

To vent the chamber, use the Kwik-Flange port to nitrogen backfill the chamber piping

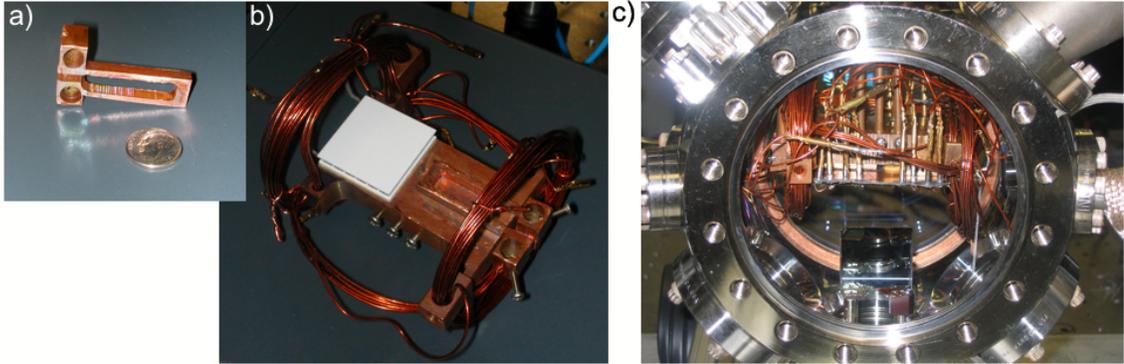


Figure 3.2: Atom chip assembly. a) The U-shaped copper block. b) The main copper support structure with U-block and TEC attached. c) Atom chip assembly in the vacuum chamber with electrical wires attached to pins.

up to the VAT valve. A nearly sealed plastic garbage bag can be attached to the end of the Kwik-Flange to provide an extra volume of nitrogen. Turn-off the ion pump and any vacuum gauges, and make sure that the cesium oven valve is tightly closed (this is very important!) Slowly crack-open the VAT valve until the chamber has reached atmosphere. Ideally, the nitrogen purge should be operated continuously, but this is impractical since the chamber is usually open for several days before pumping down again.

3.2 Atom chip assembly

The atom chip assembly—from top to bottom—consists of the sapphire atom chip, the copper-cladded Teflon circuit board, the macro U-shaped copper block and TEC, and finally the main copper support structure (see Figure 3.2). Onto the latter piece the macro U-shaped block is glued with the thermally conductive Epotek H77 epoxy. A thin Teflon insulator is inserted between the two. The support structure also houses the two pairs of intra-chamber \hat{x} bias field coils and the single \hat{z} coil. Set-screws attach six copper dowels—which are mounted to the top of the vacuum chamber—to the copper support structure. This suspends the atom chip assembly in the center of the chamber. The two high-current electrical feedthrough pins slide into slots in the U-shaped copper block and are fixed in place with set-screws.

The atom chip and Teflon circuit board are glued together using the vacuum-safe Epotek 353 epoxy. Unlike the H77, this glue cures quite quickly at modest temperatures ($> 80^\circ$).

Although it would be nice to use the thermally conductive H77 epoxy, it cures too slowly, which allows the atom chip and the Teflon circuit board to slip and misalign from one another. These two boards are aligned by visually matching marks on the sapphire atom chip substrate to those on the Teflon board that have been pre-printed to ensure that the atom chip is orientated properly in the vacuum chamber and with respect to the macro U-wire underneath. The Teflon board with chip attached is secured to the copper support structure with a screw which allows easy installation. The chip assembly is designed to maximize optical access to the trapping regions.

An ultrasonic soldering iron allows us to connect thin copper wires to the gold patterned on the sapphire substrates. While fluxless and not much more difficult than normal soldering, the ultrasonically-fashioned joints are not perfectly solid and need to be redone on occasion. Care should be taken when handling the atom chip assembly—especially during vacuum chamber installation—that stress is not applied to the joints. The copper wires attached to the sapphire chip are routed to the Teflon board and soldered to the copper cladding with silver solder. This solder contains flux, but this is easily removed with ultrasonic cleaning, and the silver alloy is much more vacuum-safe than normal solder. Attached to the Teflon board are brass pins purchased from ITT Cannon (Newark). These pins connect the Teflon circuit board wires to macroscopic Kapton-coated wires in the chamber. We use the Protel, IsoPro, and QuickCAM software packages and a Quick Circuit 5000 circuit miller to create the wire pads in the Teflon board. The sapphire substrate and Teflon board assembly must proceed in a specific order to ensure that all parts can be ultrasonically cleaned and heat-cured without breaking the microwires or solder joints: 1) make the atom chip and mill the Teflon board; 2) drill holes in the Teflon board and clean the TEC and the mounting nut; 3) glue on the nut with Epotek 353 at the same time as gluing on the TEC with Epotek H77 (heat at 150° C for 1 hour); 4) using silver solder, attach the brass pins and copper wires to the Teflon board; 5) bend pin edges to ensure good optical access above the solder joints; 6) clean in an ultrasound bath of acetone; 7) glue the sapphire substrate to the Teflon board with Epotek 353 for 30 min at 80°; 8) ultrasonically solder wires to the atom chip; 9) attach mirror, cavity, and fiber taper to atom chip using small drops of photoresist and UV curable glue. A thermistor is glued to the sapphire substrate with Epotek H77 to provide *in situ* temperature measurements, and the chip temperature rises by $\sim 8^\circ$ C during typical atom trap operation.

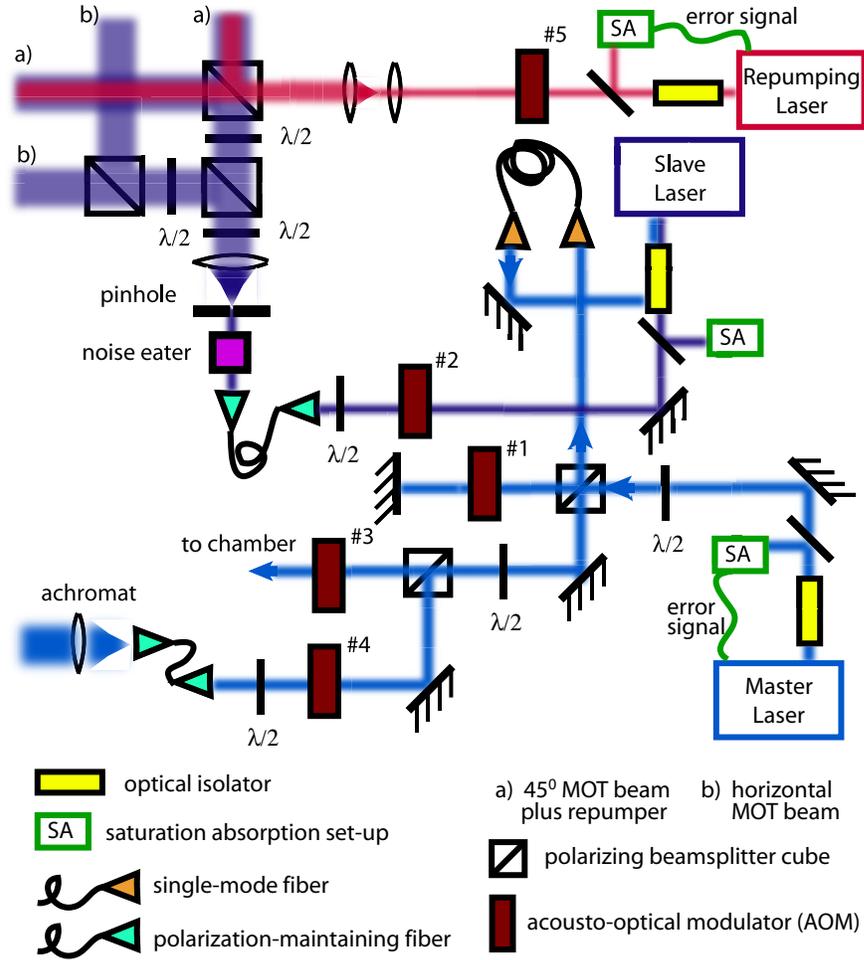


Figure 3.3: Laser system schematic. Note that many mirrors, lenses, waveplates, neutral density filters, and other optics are not shown.

3.3 Laser system

The laser system supplies the MOT trapping and repumping beams, the optical pumping beam, and the absorption imaging beam. In this section we will discuss the diode laser system used to produce and control these beams as well as the detection system we propose for use in the microcavity experiments discussed in Chapter 6.

Figure 3.3 shows the laser system. We use three diode lasers: a master laser for the optical pumping, absorption imaging, and injection lock beams, a slave laser for the MOT beams, and a laser for the repumping beam. The master and repumping lasers are built using the design from J. Buck and C. Nägerl (a former graduate student and a postdoc in Kimble's group, respectively). These are built-upon the original designs discussed in

Reference [85, 86]. The diode is housed in a New Focus mirror-mount and the diffraction grating—mounted in the Littrow configuration—is attached to the face of the mirror mount and can be adjusted with the mount’s micrometers. This assembly rests atop two TEC stages and is housed inside a sealed aluminum box for added temperature and acoustical stability. The laser frequency is most coarsely controlled by the temperature set-points of the TEC stages. A piezo stack mounted between the mirror-mount face and the horizontal micrometer provides sub-1 kHz bandwidth frequency tuning by shifting the grating angle, and a FET board adds a MHz bandwidth modulation signal to the diode current for fine-tuning and adding a modulation for deriving the laser lock error signal. There is a clean $F=4$ cesium resonance at the ~ 4.8 setting of the diode current adjustment dial. The laser outputs 31 mW of power at this setting. The repumper current dial should be set at ~ 6.1 for a clean $F=3$ signal, and this gives 42 mW of output power. The lasers can be continuously locked over a period of a day or more, and the diode current and/or piezo needs only to be adjusted if the room humidity changes. While the air conditioning system does a good job of keeping the temperature at $\sim 22^\circ$ C, the humidity is allowed to fluctuate from 20% to 65% depending on the season and weather. This causes mode-hops, and every six months or so the grating needs to be adjusted with the mirror-mount micrometers.

Both the master and the repumper lasers are locked to transitions in cesium using the saturation-absorption method [87]. The low-frequency but large amplitude perturbations are negated by the piezo branch and high-frequency perturbations are nulled by the diode current feedback branch. The diode current is modulated at 3.5 MHz to derive the error signal. To produce the MOT trapping beams, the frequency-locked laser beam must be further intensity and polarization stabilized, and its cross section must be filtered into a smooth Gaussian with high pointing stability. Moreover, the intensity and frequency need to be adjusted for the sub-Doppler cooling and magnetic trapping stages of the experiment. To accomplish this, the laser must pass through many power-depleting passive and electro-optical elements. The diode current that is optimal for bringing the laser into resonance with the $F=4$ ground-state manifold only supplies 31 mW of laser power. This is not large enough to satisfy all the power needs of the various optical elements and still provide MOT laser intensities, I , much larger than the saturation intensity $I_s = 1.1$ mW/cm². We solve this problem by employing a slave laser to act as a laser power amplifier for the master. It is difficult to simultaneously satisfy the requirement of frequency stability and high power

for a single diode laser, and this master/slave technique distributes the frequency control and high-power generation to two separate lasers.

The master laser is locked to the $F'=4$ to $F'=5$ cross-over transition in the $F=4$ hyperfine ground-state manifold. This cross-over transition is 125.5 MHz red-detuned from the $F=4$ to $F'=5$ cycling transition. The beam from the master laser (shown in blue in Figure 3.3) is split into two at a polarizing beamsplitter cube (PBC) with one 5 mW beam sent to supply the beams for optical pumping and absorption imaging and the other 25 mW beam sent to a double-passed acousto-optical modulator (AOM) to produce the injection lock for the slave laser. The double-passed AOM (#1 in Figure 3.3) is set to 69.0 MHz for the MOT loading step and 50.5 MHz for the sub-Doppler cooling step. We use a 80 MHz center frequency AOM purchased from IntraAction (all of our AOMs are purchased from this company). The purpose of this AOM is to enable variable detuning of the slave laser: for a fixed slave laser diode current, the slave follows the frequency of the injected laser (within a ~ 100 MHz bandwidth) and double-passing the AOM allows variable detuning without changing the injected beam's path. The total detuning imprinted on the injected beam by AOM #1 is +138 MHz for MOT loading and +101 MHz for sub-Doppler cooling. After the slave laser, the beam passes through AOM #2 which controls the beam's intensity and decreases the beam detuning by 30.5 MHz. The total detuning for the MOT trapping beams is +108 MHz, and this is 3.4γ red-detuned from resonance with the cycling transition ($\gamma = 2\pi \cdot 5.2$ MHz). For sub-Doppler cooling, the total detuning is +70.5 MHz which is 13.6γ red-detuned from resonance.

The slave laser is injection-locked by spatially mode-matching the master laser beam into the slave's diode. The temperature and current of the diode are adjusted to capture the lock, and for injection input powers of ~ 1 mW, the slave laser follows the frequency of the master over the detuning range of AOM #1. Even though double-passing AOM #1 should eliminate beam pointing variation between the two detuning set-points, we noticed that the beam shifts enough to degrade the injection locking performance. We filter the beam with a short single-mode fiber to map pointing instability into intensity variation. This solves the problem since the frequency and output power of the slave laser are very insensitive to injection lock power fluctuations at the 10 to 20% level. Misaligning the output polarizer of the slave laser's optical isolator allows one to pass the injection lock beam through the isolator and into the laser diode. The optical isolators from Optics For Research have a

beam rejection port at the output polarizer, and a small fraction of the slave laser beam is redirected through this port when the polarizer is misaligned. Injection beam mode-matching is achieved by overlapping this beam with the rejected beam and matching the two beam's widths. The slave laser lock quality is greatly affected by this mode-matching and this needs to be re-optimized every few weeks. Our slave laser outputs 140 mW when the current controller dial is set to ~ 8.25 , which is the optimal setting for locking the master laser's injection beam. The slave diode's temperature needs to be adjusted to make the slave operate at a different output power. The slave laser's output beam is shown in purple. We monitor the injection lock quality with a saturation-absorption set-up, but do not use the set-up for any feedback control of the laser. This beam then passes through AOM #2 which uses the -1 order to decrease the beam's detuning and to control the beam's intensity. This is a 40 MHz center frequency AOM. For sub-Doppler cooling, AOM #2 attenuates the beam power by a factor of 10 and extinguishes the beam during the magnetic trapping steps.

After AOM #2, the MOT trapping beams pass through a polarization-maintaining (PM) fiber. This fiber serves four purposes: 1) Decouples the laser control system from the chamber optics. This prevents misalignment of the control system from propagating all the way to the MOT alignment optics; 2) Filters the spatial mode of the beam into a Gaussian profile (the AOM's distort the beam shape); 3) Converts pointing instability into intensity noise that is subsequently filtered-out using a Thorlabs noise eater. This is an LCD polarizer that variably attenuates the beam to maintain a constant power. 4) Rejects retroreflected MOT beams, preventing the slave laser lock from being disturbed by an excess of reflected light. The power of the beam pre-fiber is 54 mW (the zeroth order beam from AOM #2 has 42 mW of power). The post-fiber power is 25 mW when set for the MOT beams, and 2.2 mW for the sub-Doppler cooling beams. The AOM #2 and fiber allow 1 μm of power to pass during the magnetic trapping stages, and a mechanical shutter can be used to completely extinguish this light if needed. Post-fiber and noise eater, the beam is further filtered by a pinhole. This is overkill, but serves as a compact beam expander. The beam is expanded to a 1 cm diameter and split into four beams using a series of half-waveplates and PBCs. The intensities are controlled by the waveplates so that counterpropagating MOT beams at the position of the chamber are intensity balanced. Irises are used as aids in overlapping the beams. Beams "(a)" are used for the two 45° mirror MOT beams and have slightly more power than the two horizontal mirror MOT beams, "(b)."

The repumping laser is constructed and frequency-locked in the same manner as the master laser. The repumping laser is locked to the $F'=3$ to $F'=4$ crossover transition in the $F=3$ ground-state manifold, which is 100.6 MHz red-detuned from the $F=3$ to $F'=4$ transition. The beam passes through a 100.6 MHz AOM (#5, 110 MHz center frequency) to bring this beam into resonance and to allow controllable beam attenuation. The repumping beam then passes through neutral density filters before being expanded and overlapped onto the 45° MOT beams.

The optical pumping and absorption imaging beams are produced and controlled by AOMs #3 and #4, respectively. These AOMs are detuned by +125.5 MHz to bring the master laser's beam into resonance with the cycling transition. The absorption beam is put through a PM fiber for spatial filtering. The output of the fiber is expanded and collimated at a width of 2 cm by an achromatic lens to improve image quality. Due to constraints imposed by the chamber, the optical pumping beam happens to be aimed at a CCD camera that is used for fluorescence imaging. AOM #3 cannot extinguish the beam below the sensitivity of the CCD camera, and we must use a mechanical shutter to completely extinguish the beam when it is not in use. We employ a shutter fashioned from an ordinary computer speaker [88]. An opaque flag is attached to the solenoid of the speaker, and flipping the direction of the drive current makes the flag pop up and down. We can extinguish a tightly focused beam in $20 \mu\text{s}$ and produce pulses of 1 to 2 ms. This device is simple to build, and at a cost of $\sim \$20$, is much cheaper than commercial alternatives.

We intend to use our lab's Coherent MBR 110 Ti:Sapphire laser (pumped with a Coherent Verdi V8) for the cavity probe beam. Using an electro-optical modulator (EOM) we lock the laser to transitions in either the $F=3$ or $F=4$ ground-state manifolds. For the GHz detunings that will be required for probing the vacuum Rabi structure of the photonic bandgap (PBG) and microdisk atom-cavity systems, we could either let the laser free-run at these detunings or use a beat-note lock [89]. The laser is passively stable to within a few hundred MHz—as measured by an in-house Burleigh wavemeter—and has a linewidth of ~ 100 kHz. The cavity linewidths are larger than a GHz, so detuning without locking for initial experiments is not inconceivable.

The optimal drive power for the PBG and microdisk cavities is in the 10 nW range which is inconvenient for shot-noise limited detection. This power is too large to use Geiger-mode APD's and too low for Si photodiodes. We plan to use either heterodyne detection with a

Si 125-MHz Photodetector from New Focus (1801), or photon counting with a cw-mode Si APD from Analog Modules (712A-4). The Analog Modules APD has a gain of 7.7×10^7 V/W, an AC-coupled bandwidth of 80 MHz, and a NEP of $20 \text{ fW}/\sqrt{\text{Hz}}$. At a power of 10 nW [30 nW], the photon counting detection would be shot-noise limited by 3.8 dB [6.2 dB] with a $1 \mu\text{s}$ detection integration. The signal size would be 0.77 V for 10 nW of optical power.

The New Focus Si photodetector has a gain of 2.5×10^4 V/W, a NEP of $2.8 \text{ pW}/\sqrt{\text{Hz}}$, and a measured saturation power of $140 \mu\text{W}$. For a signal power of 30 nW, an integration time of $1 \mu\text{s}$, and a local oscillator (LO) power of $140 \mu\text{W}$, we could achieve 3.1 dB shot-noise limited detection. The signal-to-detector noise would be ~ 300 , and the signal size would be 50.8 mV. We plan on using an all-fiber heterodyne mixing set-up with a 97.5/2.5 fiber-based beamsplitter. The LO will be 10 MHz detuned from the signal to ensure the beat signal is at the low-noise frequency band of the detector but high enough to have a 1 MHz detection bandwidth centered about the beat signal frequency. The 10 MHz beat signal will be captured by a 14-bit, 100 MHz bandwidth digitizing board from Gage Applied Technologies and we will digitally mix-down the signal. This LO detuning will be generated by the difference frequency between two AOMs to ensure that no RF noise will leak into the beat signal bandwidth.

We would eventually like to be able to lock the microcavity. Since both the cavity-laser detuning and the cavity linewidth are so large, we will need a high bandwidth detector to derive the error and/or beat-lock signals. Menlo systems makes a Si APD that could be used with low probe power: their APD210 has an AC-coupled bandwidth of 1 GHz, gain of 2.5×10^5 V/W, and an NEP of $280 \text{ fW}/\sqrt{\text{Hz}}$.

3.4 Computer control system

A computer is required for $100 \mu\text{s}$ time-scale control of the laser system, microwire currents, magnetic fields, and CCD camera image acquisition and processing. We use a combination of Labview and Matlab to control various DAQ boards and pulse-delay generators. The system is divided into a master computer and three slaved instruments consisting of one computer and two pulse-delay generators. The master computer contains

the main control programs, sets the master trigger, and houses the National Instruments (NI) DAQ boards. The slave computer captures the CCD camera images with a NI IMAQ board and uses Labview to process, display, and record them. A Stanford DG535 pulse-delay generator receives the master trigger and sends out two secondary triggers, and a BNC 555 pulse-delay generator sends-out four TTL pulses that control various elements of the experiment.

The experiment is governed by five programs in the master computer—one Matlab-based and four Labview-based. The Matlab program generates data files which contain a synchronized sequences of voltage outputs for the analog and digital output channels. The analog channels output the control voltages for the magnetic coil and microwire power supplies. The digital outputs send TTL pulses to various triggers, AOM RF switches, shutter triggers, etc. This program is modular in that each trapping step is self-contained and may be added or subtracted from the experiment or modified without affecting the others. The modules contain the voltages and durations for each output channel, and the program produces plots that chart the timing of all of the channel voltages. This is useful for debugging and diagnosing the microwire and coil currents. The code for this program is listed in Appendix B, Section B.4. The program outputs a series of .dat files for each channel, and two Labview programs—one for the analog channels and one for the digital—read these files and upon a TTL trigger, synchronously outputs the files' voltages.

The analog channel Labview VI (such as `Waveguide.vi`) is run in loop mode: the experiment trigger causes the program to read the data files into the NI DAQ boards for output. There are two NI boards used for analog output. The PCI-MIO-16E-4 has only two analog outputs, and we also use it for generating the Stanford DG535 trigger. The NI PCI-6713 has 8 analog outputs. Note: neither of these two boards can be used for outputting a binary waveform due to the lack of an on-board internal clock. The digital output VI, `D0.vi`, functions in a similar manner. In loop mode, the program waits for a trigger before sending a digital binary waveform to one or more ports of the NI PCI-6533 board which contains 32 channels partitioned into 4 ports⁵. The VI, `Trigger - DG535.vi`, sets—via GPIB—the trigger mode of the Stanford DG535. During normal experimental operation, this should be set to external trigger, but in debugging the MOT it is often useful to set the trigger to

⁵Due to a bug in this VI, to terminate this program one must release the loop-mode button and allow one more trigger to pass. Otherwise Labview crashes.

internal mode with a sub-Hz repetition rate.

The VI that starts each experiment is **Series – DG535.vi**. The experimental duration and repetition delay time are set in this program. The duration is the time at which after $t = 0$ the CCD camera acquires an image, and this time should be set according to the timing schedule produced in the Matlab program. The repetition delay time, usually around 5 s, is approximately the MOT loading time. This program also has the ability to scan the experimental duration in time steps of one's choosing. With a GPIB connection, **Series – DG535.vi** sets the trigger delay times of the Stanford DG535. The DG535 has two TTL pulse outputs defined by channel A plus B and by channel C plus D. The AB pulse triggers the CCD camera's asynchronous reset and the image acquisition board. The CD pulse triggers the experiment—the DAQ boards and the BNC pulse-delay generator—and **Series – DG535.vi** calculates the delay of this pulse to ensure that the image is acquired during the intended moment. The BNC 555 outputs four TTL pulses: channel A (active low) is the macro U-wire shut-off and has a width and delay of 2 s and 20 ms, respectively; channel B (active high) opens the optical pumping shutter for a 10 ms pulse centered about the optical pumping step; channel C (active low) turns-on AOM #3 for the optical pumping beam; and channel D (active low) activates the sub-Doppler cooling detuning and intensity attenuation. The width of the channel A pulse sets the MOT duty cycle, and this can be tuned to adjust the atom chip temperature since the macro U-wire dissipates heat. The BNC 555 instrument should be replaced by the digital output DAQ channels.

The experimental control sequence proceeds in the following manner. After setting the experimental timing sequence and analog voltage magnitudes with the Matlab program, the VIs **Waveguide.vi** and **D0.vi**, are initialized in loop-mode. The experiment duration is set in **Series – DG535.vi** and the experiment begins when this VI is run either in loop-mode—for continuous operation—or single-shot mode. **Series – DG535.vi** sets the correct delay times in the Stanford DG535 via GPIB and then sends a TTL pulse from a digital output of the PCI-MIO-16E-4 board to trigger the DG535. The DG535 then triggers the image acquisition board in the slave computer to wait for a CCD camera signal, sends an asynchronous reset pulse to the camera, and outputs a trigger pulse to the PCI-MIO-16E-4, PCI-6713, PCI-6533, and BNC 555. The VIs then begin to synchronously output the data points in the Matlab generated files to the PCI-MIO-16E-4, PCI-6713, and PCI-6533 boards. The output time resolution can be adjusted in the software, and 100 μs is sufficient

for present purposes.

We use a separate slave computer for the image acquisition because we found it difficult to get the NI IMAQ and NI DAQ boards to work simultaneously on the same computer. Regardless, the master computer does not have enough PCI slots for all the boards, and an extra computer is necessary regardless. The slave computer has two VIs for image acquisition, one for fluorescence imaging and one for absorption. The programs display the image in real time as well as the atom number and cloud width, and automatic image saving is incorporated. Both VIs run a Matlab script to calculate the atom number calibration and to fit Gaussians to slices of the atom cloud. The gain and integration time of the CCD cameras (Cohu models 2100 and 4920) can be adjusted with on-board switches. Asynchronous resetting is required of the analog cameras because they are triggered off the 60 Hz power signal. We use the cameras in “frame” mode to maximize the image resolution in each integration time. However, in this mode only every other horizontal line is captured and the image resolution is reduced by half in the vertical direction. The first image acquired after an asynchronous pulse is garbage, and the image capture VIs wait to record the second, ~ 16 ms delayed image. The triggers from the Stanford DG535 are designed to account for this delay. For absorption imaging, a background image is captured ~ 32 ms after the first.

We use Kepco power supplies for all the microwire and magnetic coil currents except for the macro U-wire which is powered by a PowerTen 60 V/55 A supply (0-5 V current control). Crydom solid-state relays are used for fast termination of the wire currents, and fast-acting protective fuses are installed at both the positive and negative terminals⁶ for each microwire power supply. All Kepco supplies are operated in “fast” mode to allow rapid current switching. Current-controlled Kepco 20-20 (20 V/20 A) BOPs are used for the \hat{x} and \hat{y} axis bias coils to enable fast switching and bi-polar current operation. These are operated in current mode and controlled by a ± 10 V voltage from the computer DAQ boards. The microwires are supplied by Kepco 25-4 (25 V/4 A) supplies. These can be wired to be either in current or voltage mode and are controlled by a voltage from the computer DAQ boards. Current mode is preferable, but if a circuit relay is installed, voltage control must be used. In current mode, a 0-1 V control signal spans the 4 A range and in voltage mode, 0-10 V spans the 25 V range. If the supply is to be grounded, the *positive* terminal must be shorted to ground. To float the supply, disconnect this short and install an opto-electrical grounding

⁶We protect both terminals since the power supplies are linked by the crossing microwires.

isolator between the computer DAQ voltage signal and the supply's control circuitry: if the latter is not done, the computer ground will propagate through the control circuitry into the power supply. A Burr-Brown ISO124P chip is suitable for this optical isolation.

Future experiments should use a digital CCD camera for absorption imaging. These 12-bit (and higher) resolution cameras will be a great improvement over the 8-bit analog cameras. Moreover, they can be triggered at will, removing the complications that arise from synchronizing the experiment timing to the 60 Hz power line signal and would allow more rapid capture of background images. Although the computer control system is a hodgepodge of several programs and instruments, it is quite robust and transparent to debug. However, future experiments might want to consolidate all the NI board control into Matlab: the newest version of Matlab—Version 7—contains NI drivers that could allow a more compact programming interface than using several Labview VIs. Another alternative is to use the ADwin-GOLD system that we recently purchased. It has a dedicated CPU that can be easily programmed to produce all the analog and digital input and output.

Chapter 4

Fabrication of Micro-Magnetic Traps for Cold Neutral Atoms

This is an updated and expanded version of Reference [8].

Many proposals for quantum information processing or quantum feedback experiments require precise control over the motion of neutral atoms, as in the manipulation of coherent matter waves or the confinement and localization of individual atoms. Patterns of micron-sized wires, fabricated lithographically on a flat substrate, can conveniently produce large magnetic-field gradients and curvatures to trap cold atoms and to facilitate the production of Bose-Einstein condensates. The intent of this chapter is to provide the researcher who has access to a standard clean-room enough information to design and fabricate such devices.

The clean rooms of Professors Michael Roukes and Axel Scherer were used extensively in this research. Eyal Buks, Darrell Harrington, and Marko Loncar were generous with their time in teaching the author the basics of photolithography and ion milling.

4.1 Fabrication Challenges and Constraints

Fabrication of atom chips poses several challenges in addition to those encountered in standard photolithography [90]. Many applications require the wires to be a couple microns wide by a few microns tall and spaced only a few microns from one another. One micron resolution is near the limit of standard photolithography, and much care must be taken to accurately produce these micron-sized wires. Wires with widths much less than a micron—though perhaps important for realizing potentials with sub-micron scale features—are of limited usefulness for creating large magnetic field gradients and curvatures since they be-

come limited to the same maximum current density as micron-sized wires [68]. Further fabrication complications arise from the need to trap the atoms near the substrate's surface, and the need to connect the microwires to macroscopic leads without blocking optical access. A common technique for trapping atoms near the substrate surface, the mirror magneto-optical trap (mMOT), requires that this surface be an optical mirror as well as the support surface for the microwires (see Figure 2.8) [2]. The substrate surface needs to be larger than 5 to 10 cm² to accommodate the reflected trapping beams as well as to allow the pads for macroscopic wire contacts to be outside of the mirror area and not blocking the optical access needed for the trapping, imaging, and pumping beams. Consequently, the wire pattern must be flawless over an exceptionally large surface area: during fabrication one must be extremely careful that no dust or surface defects break or short the wires.

The major fabrication challenge lies in increasing the height of the wires to a few microns. Even the smallest wires need to support up to an amp of current, and consequently, the cross-sectional area of the wire must be maximized. This reduces wire resistance and limits the heating that causes wire breakdown. Moreover, attention must be paid to the thermal conductivity of the substrate and mounting system to ensure sufficient power dissipation. Sapphire or polished aluminum nitride (AlN) substrates provide sufficient thermal conductivity, but are slightly trickier to use for fabrication than more standard substrates. The group of Joerg Schmiedmayer in Heidelberg has recently found that Si coated with an insulating oxide layer is the optimal substrate in terms of thermal conductivity [91].

The use of microwires to create an Ioffe trap illustrates these challenges. The wire pattern shown in Figures 4.1(a) and (b) creates a 3D harmonic trap when combined with a perpendicular homogenous bias field [3]. Unlike a quadrupole trap, the Ioffe trap has a non-zero field at the trap center and thus does not suffer from Majorana spin-flip losses. An atom is confined within the Lamb-Dicke regime when its recoil energy is less than the trap's vibrational level spacing ($\eta = (E_{recoil}/E_{vib})^{1/2} < 1$), and for a cesium atom this occurs when the trap curvature exceeds 2×10^6 G/cm². To achieve this extremely large field curvature in all three dimensions, the radius of the wire pattern in Figure 4.1(a) must be smaller than ~ 30 μm . For a trap of inner radius 10 μm , outer radius 15 μm , and wire current $I = 1$ A, the curvature and Lamb-Dicke parameter, η , at the center of the trap in the axis perpendicular (plane parallel) to the substrate is 2×10^8 G/cm² (2×10^{10} G/cm²) and $\eta = 0.38$ ($\eta = 0.11$). The closely spaced wires can only be a few microns wide, and even

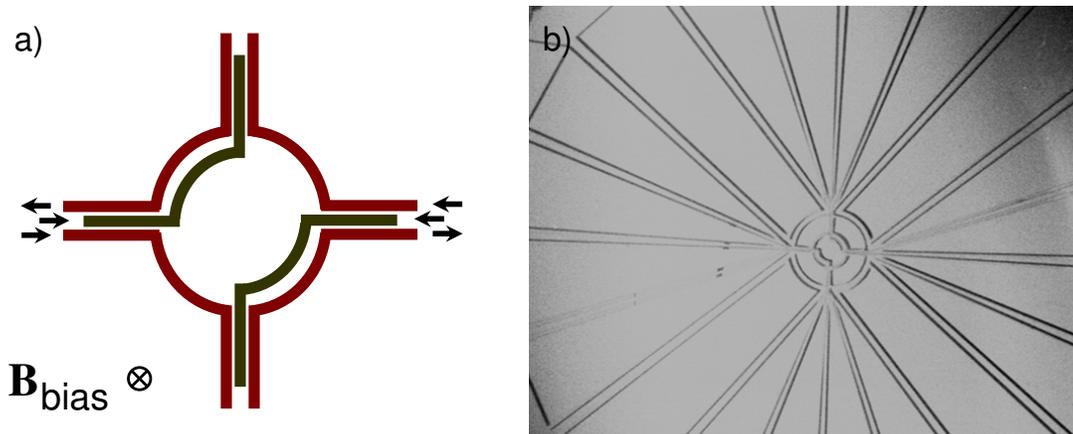


Figure 4.1: The planar Libbrecht-style Ioffe trap. a) When combined with an opposing bias field, this wire pattern produces a 3D harmonic potential above the substrate with a non-zero field at the trap center [3]. b) A planar Ioffe trap with an on-chip bias coil fabricated with gold on sapphire using the lift-off method. In the sample shown here, the wire height is $1.5 \mu\text{m}$ and the minimum wire width is $10 \mu\text{m}$. The gold between the wires forms a mirror for creating a mirror MOT.

if fabricated to a height of 2 to 4 microns, the wires would need to support the large current density of $\sim 10^{11} \text{ A/m}^2$. The accommodation of laser beams for atom cooling, loading, and imaging constrains and complicates the atom chip's design. The trap minimum is only $7 \mu\text{m}$ from the substrate's surface, and the mirror patterned on the surface for use with the mMOT must neither short the Ioffe wires nor extend more than $\sim 5 \mu\text{m}$ from the surface. The following sections describe the necessary fabrication tools and the techniques we use to overcome these challenges.

4.2 The elements of atom chip fabrication

Microfabrication is a labor intensive process, often involving several weeks of trial and error to perfect the fabrication recipe. However, once the process works, five to ten atom chips can be produced over a span of two to three days. The intent of this paper is to provide the researcher who has access to a standard clean-room enough information to design and fabricate an atom chip. We will describe the use of fabrication instruments and techniques only insofar as they are relevant to atom chips. Fabrication is not an exact science, and the techniques described here may not be optimal, but nevertheless have proven successful for

the chips we have fabricated.

In photolithography, UV light shone through a photomask casts shadows onto photoresist, a light sensitive polymer, which is coated on the surface of the substrate. Either positive or negative photoresist may be used, with the primary difference being that exposed areas of positive photoresist are removed after developing whereas exposed areas remain in a process using negative photoresist. The various fabrication techniques differ in how the wire metal and photoresist are used to create the wire patterns. For instance, the wire metal may be either thermally evaporated into the trenches created in the photoresist, or grown upward through the trenches by electroplating onto a seed metallic layer underneath the photoresist. The photoresist and unwanted metal are removed leaving only the desired wire pattern. Generally, chip fabrication consists of six steps: creating a photomask containing the desired wire pattern, using photolithography to transfer the wire pattern to photoresist on a substrate, thermally evaporating wire material, increasing the wire height, preparing the surface mirror, and making contacts to macroscopic wires. The details and exact order of these steps vary depending on the specific requirements of the microwire pattern to be fabricated. For instance, wires wider than $30\ \mu\text{m}$ or less than one micron in height may be fabricated with a much simpler technique than thinner or taller wires. This section discusses the steps common to all techniques. Procedures required to increase the wire's thickness pertain to individual fabrication techniques and will be discussed in the next section.

4.2.1 The photomask

The photomask is typically a 10 cm square piece of glass or transparent plastic on which is printed a positive or negative 1:1 image of the wire pattern. Wire patterns with widths or spacings less than $\sim 30\ \mu\text{m}$ require a professionally made chrome mask: one in which the pattern is written with chromium on a glass plate. We have used the company Photronics, Inc. (telephone 619-992-8467) to make photomasks from AutoCAD drawings. Much care must be taken in producing the AutoCAD files since not all functions are properly converted to the company's file format. In particular, all shapes should be drawn as closed, unfilled polylines. These masks are quite expensive, costing between \$600 and \$800, but have sub-micron resolution and are typically shipped within a week. It is possible to purchase a laser writer to produce in-house photomasks with resolution down to $0.8\ \mu\text{m}$. This can be a cost-effective alternative to purchasing individual masks from companies.

Many commercial printing shops are capable of printing overhead transparencies with high enough resolution to serve as photomasks for wire patterns with features larger than $\sim 30 \mu\text{m}$. The line edges are granular on a scale of a few microns, and the UV exposure time must be adjusted to account for the ink not being perfectly opaque. However, the one day turn-around, low cost of $\sim \$20$, and ease of file preparation—only an .eps file is typically needed—make the transparency photomask quite an attractive alternative for large features.

4.2.2 The substrate

As mentioned earlier, the substrate material for the atom chip should be carefully chosen: it must be electrically insulating, highly polished, not susceptible to fractures upon localized heating, and an excellent thermal conductor. We found that both sapphire and AlN substrates satisfy these requirements. Sapphire substrates 0.5 mm to 2 mm thick with surface areas of several cm^2 may be purchased from companies such as Meller Optics, Inc. (telephone 800-821-0180) for \$30 to \$50 apiece. A surface quality of 80-50 scratch-dig is sufficient for fabrication. The thermal conductivity of AlN, $170 - 180 \text{ Wm}^{-1}\text{K}^{-1}$ at 20°C , is ~ 4.5 times higher than that of sapphire [68]. We measured that the max current density supported by microwires on AlN, $\sim 2 \times 10^{11} \text{ A/m}^2$, is a factor of two greater than for microwires patterned on sapphire. This was measured using electroplated gold wires of varying cross-sections patterned exactly the same way on both AlN and sapphire substrates. Specifically, we used several $3 \mu\text{m}$ and $20 \mu\text{m}$ wide wires whose heights ranged from one to three microns. The substrates were glued to room temperature copper blocks using Epotek H77 (Epoxy Technology, telephone 978-667-3805), a thermally conductive epoxy. Reference [91] finds that substrates of oxide on silicon have superior thermal properties to sapphire and AlN, and this silicon substrate is in some cases more amenable to standard microfabrication techniques.

Compared to AlN, sapphire substrates are easier to use for fabrication because their transparency allows one to detect and avoid defects and dust during the photolithography process. Moreover, with a transparent substrate, it is easy to align features on the substrate to devices on the surface underneath. Polished AlN substrates may be purchased in bulk for less than $\sim \$75$, and unlike sapphire, AlN substrates can be cleaved with a diamond scorer to any shape desired. We were able to dice a $\geq 1 \text{ mm}$ thick sapphire substrate using a diamond saw, but on occasion the substrate cracks in undesirable ways. The polished

AlN still has a considerable amount of surface roughness—one micron wide plateaus a few hundred nanometers tall are typical—but we found that it is nevertheless possible to fabricate on this surface perfect three micron wide wires spaced less than three microns from one another. The surface bumps simply map directly onto the upper surface of the wires.

4.2.3 Substrate cleaning

Before the photolithography process may begin, the surface of the substrate must be cleaned to remove all organic material and dust. Although some of the following steps may seem unnecessary and “overkill,” investing the time to thoroughly clean minimizes the chance that after many hours of work, one discovers that a piece of dirt has broken or shorted a wire. The first step is to immerse the substrate in a beaker of “piranha etch,” sulfuric acid and hydrogen peroxide in a 10:1 volume ratio brought to 100° C on a hot plate for ~5 min. Teflon coated, flat-tipped tweezers are ideal for manipulating substrates. After the etch, the substrate should be placed in a beaker of acetone, heated again to 100° C for a few minutes, and finally inserted into an ultrasound cleaner for few more minutes. In extreme cases of substrate grime, a cotton tipped dowel can be used to manually wipe away the dirt. Acetone leaves a thin film—and sometimes even particulate—when allowed to dry on a substrate’s surface. It is imperative that one spray isopropanol (IPA) onto the substrate as it is removed from the acetone bath. This rinses the surface of acetone and wets it with IPA which does not dry quickly. The substrate must then be rinsed with methanol, which is relatively clean and does not leave a film, and immediately blown dry with an air or nitrogen gun. It is crucial that the air jet is aimed almost parallel to the surface so that the methanol is blown-off rather than dried on the substrate. When done correctly, the only remaining dirt particles will be along the edge of the substrate that is downwind of the air jet, and not in the center fabrication region. If the substrate is reasonably clean after the piranha etch, then the acetone step (which may actually add some dirt particulate) may be skipped, and the substrate should instead be immersed in IPA and placed inside an ultrasound cleaner.

4.2.4 Thermal evaporation

Certain fabrication techniques, to be discussed below, require that a 100 nm metal layer be thermally evaporated before coating the surface with photoresist. We take this opportunity to discuss the thermal evaporation process. We use gold for the wires because of its high electrical conductivity, resistance to corrosion, and ease of evaporation, electroplating, and wet etching. To successfully deposit gold on a substrate's surface, one must first evaporate a 50 Å metallic layer that promotes adhesion between the gold and the sapphire or AlN. We typically use chromium, but titanium may also be used. At the level of our current experiments, the magnetic effects from the thin layer of chromium are negligible. In a thermal evaporator, the substrate is mounted in a vacuum chamber facing a tungsten crucible positioned a few tens of centimeters below. The crucible, known as a boat, can hold 10 to 20 pieces of ~2 mm long and 0.5 mm diameter gold wire. Current flows through the boat, melting the gold and spewing it upwards toward the substrate. A calibrated crystal monitor measures the deposition rate. One to two boats are sufficient to deposit 100 to 200 nm of gold, and this costs \$10 to \$15 per boat. There are typically only four sets of electrical feedthroughs in the evaporator's vacuum chamber, and to deposit more gold, one needs to bring the chamber up to atmosphere, reload the boats with gold, and pump back down to base pressure ($\sim 1 \times 10^{-6}$ Torr)—a process that takes about an hour. The substrate mounting area allows several substrates to be coated at once. Evaporating less than 1 μm of gold is reasonable, but depositing more than 1 μm becomes too expensive and time consuming, and the quality of the gold surface begins to diminish. Moreover, the vacuum chamber eventually becomes hot which may result in the failure of the crystal monitor or the burning of photoresist. Sputtering the gold is an option that we have not explored, but may be more efficient.

Some groups have reported an intermittent difficulty with getting the adhesion layer to “stick” regardless of whether Cr or Ti is used, and have not found a consistent culprit. This results in the peeling away of the gold layer after evaporation. We have only had one episode of this occurring (in what is known as the “left” evaporator in Roukes' lab), and we believe it was caused by the combination of a leaky and dirty vacuum chamber. The Cr or Ti became corroded either as it evaporated onto the substrate or once it was attached. The vacuum pressure would rise abnormally upon the melting of the Cr or Ti boats. The

problem was solved by simply switching to the “right” evaporator, but this is hardly a long-term fix and an inspection and thorough cleaning of the “left” evaporator should be done.

4.2.5 Photoresist spinning and baking

Photoresist does not always adhere well to the substrate’s surface. Before coating with photoresist, the substrate should be baked on a hot plate at $\sim 150^\circ$ C for a few minutes to remove surface moisture. However, caution must be taken with custom-cut sapphire substrates. A few of these have cracked after being placed directly on a 120° C hot plate. Slow ramping of the hot plate temperature may be required.

Hexamethyldisilazane (HMDS) should be used with sapphire and quartz substrates to promote adhesion (this is unnecessary for AlN). Only a few monolayers of HMDS are required: after baking, place the sapphire in a dish next to several drops of HMDS and cover for a few minutes. Note that both HMDS and photoresist are carcinogenic and should be handled with care.

Spinning photoresist onto a substrate is a relatively straightforward process. The substrate, with beads of photoresist dripped onto its surface, is spun by a vacuum chuck to a few thousand rpm for several tens of seconds. A faster rotation results in a thinner film of photoresist. Typically, a film thickness of a few microns is possible with standard photoresists, and there exists special resists that are four to twenty microns thick. These thicker resists are often important for making tall wire structures. The thickness of a photoresist may be increased beyond its specification by dripping resist onto its surface during rotation. After spin-coating, the photoresist needs to be baked on a hot plate to prepare the polymer for UV exposure. The exact temperature and bake duration are often crucial to the success of the fabrication. We would like to note that it is possible to layer microwire patterns on top of one another by fabricating each new wire layer on top of a spin-coated insulator such as polyimide [92]. This will be discussed further in Section 4.2.10.

4.2.6 UV exposure

The central step in photolithography is the UV exposure of the photoresist. An instrument known as a mask aligner allows one to accurately position the photomask flush to the substrate’s photoresist-coated surface, and a built-in UV lamp exposes the photoresist for a

specified amount of time. Essential for photomask and substrate registration is an optical microscope mounted on the mask aligner. This enables one to simultaneously view the wire patterns on the mask and on the underlying substrate. Dust particles or scratches often remain on the substrate even after a thorough cleaning. If these defects are sparse, then the substrate may be translated so that the wires avoid all defects. Aligning the chip's wire pads along one or more edges of the substrate further constrains the relative position of the photomask to the substrate. It should be noted that it is difficult to properly develop the pads (or other wire features) less than a millimeter from the edge due to photoresist beading. Certain fabrication recipes require the photoresist to be baked and exposed again before developing. For periodic micron-sized features such as those used for making an atom mirror (see Chapter 8), it may help to remove the beaded photoresist at the edge of the substrate to allow the substrate to lie flush against the photomask.

It is good practice to clean the chrome photomasks after every use. Photoresist can stick to the surface, and if left for days, will produce hard-to-remove specs that can block the UV light, creating unwanted features or breaks in the patterned wires. Immersing in a dish of acetone and rinsing with IPA and methanol is sufficient for routine cleaning. Some chrome masks can withstand ultrasound cleaning as well as being wiped with a soft, lint-free cloth, and this seems to be the only way to remove encrusted grime or particulate.

4.2.7 Developing

To remove the photoresist regions defined by the UV exposure, the substrate must be immersed and slightly agitated in a beaker of developer for a few tens of seconds followed by a water rinse. The exact developing time depends on the previous fabrication steps, but it is generally possible—especially with the transparent sapphire substrates—to see a characteristic change in opacity of the photoresist as it becomes fully developed. For instance, when using a positive process, one first sees the exposed photoresist turn hazy, revealing the wire pattern. After a few seconds, the hazy region sloughs off exposing the bare substrate and leaving darker, patterned regions of photoresist. If a mistake is made at any point in the photolithography process, the substrate can be reused by removing the photoresist in a beaker of acetone and cleaning the substrate as mentioned above, starting with the ultrasound.

4.2.8 Ozone dry stripping

Certain fabrication processes require the substrate surface to be etched in an ozone dry stripper. This uses UV light, ozone, and heat to remove thin films of unwanted organic material, photoresist, or HMDS that may prevent the deposition of thermally evaporated or electroplated gold. The time and temperature of the process may be adjusted to optimally remove organics without over-baking the photoresist.

4.2.9 Wire contacts

Wire bonding and ultrasonic fluxless soldering are useful methods for attaching macroscopic wires to the substrate's contact pads. Wire bonding is the standard method for making contacts to micro- or nanofabricated devices. The wire bonder attaches each end of a thin thread of gold wire to a pad using a heated, ultrasonically vibrating tip. The thin wire may be stretched over several millimeters between the pad on the substrate and a pad on the substrate support structure. The pads on the support structure may then be connected to standard wire contact pins. Because the wire threads are prone to break and cannot individually support more than a few hundred mA of current, it is necessary to make several redundant bonds per pad. This process can be quite time consuming. As an alternative, ultrasonic soldering irons are capable of attaching regular wires to sapphire or AlN using fluxless solder. Attaching wires is nearly as simple as standard soldering, and the fluxless solder is vacuum compatible to at least 10^{-9} Torr. Unfortunately, the solder material forms mounds on the substrate's surface that can limit optical access.

4.2.10 The mirror

Finally, we would like to discuss methods for making the atom chip's surface mirror-like. This is a crucial step for ensuring good mirror MOT performance and for minimizing scattered light. The latter is important for imaging the atoms with the substrate as a background.

The most compact method involves simply patterning gold on the entire chip's surface except for thin, $\leq 10 \mu\text{m}$ wide gaps around the actual wires [7]. This technique does not add any additional steps to the fabrication procedure, but it does increase the likelihood that surface defects will short the wires through contacts to the large mirrored areas. The

mirror gaps that define the wires imprint defects onto the reflected mirror MOT beams, but we have nevertheless been able to trap more than a million cesium atoms with this less than perfect mirror. It is important that the wires themselves are highly reflective. Thermally evaporated or sputtered wires work well for this, but electroplated wires do not. The gold electrodeposition process grows wires with a granularity large enough to scatter light non-specularly. Although we have been able to form good mirror MOTs with electroplated wires, it is impossible to fluorescence image atoms looking down onto the substrate when the atoms are above these wires. This limits diagnostic capability.

Coating the chip's surface with an insulator and then applying a mirror coating can produce a more specular mirror, though at the expense of additional material between the atoms and the wires. Since the atoms can no longer be trapped immediately above the wires, this limits the maximum attainable trap gradient which scales inversely with trap height. Nevertheless, we have found the various forms of this technique quite useful.

Epoxying a silver mirror (with Epotek 353) to the surface forms a good mirror, and it eliminates any corrugations on the mirror surface caused by the underlying wires [50]. Unfortunately, the minimum distance between the atoms and the wires is set by the mirror and epoxy thickness. An improved mirror can be made by epoxying a dielectric mirror onto the surface. The mirror was grown on a sacrificial glass substrate with a detachable layer between the mirror and the glass. The device to be mirrored is glued onto the mirror with a thin layer of Epotek 353, and after curing the device plus mirror easily peel-away from the glass substrate. Anti-reflection coatings can be attached to difficult-to-coat windows in the same manner¹. Vacuums of 2×10^{-10} Torr, in a chamber baked to 150° C, have been achieved despite using this glue and dielectric coating.

An alternative technique, which hasn't been completely successful, is to spin one or more layers of photoresist onto the substrate. Swabbing with acetone removes the photoresist covering the wire pads near the substrate's edge, and the mirror is created by using a mask to thermally evaporate gold only onto the coated region. This technique is simple and works well as long as there are no vertical protrusions of gold from the wires to short to the gold mirror layer. We find that one micron tall wires are fine, but wires 5 to 15 microns tall can on occasion short to the mirror even when the surface is coated with three

¹This coating is produced by the German company OIB (Optical Interference Components) <http://www.oib-jena.de/firmenpreng.html>

layers of photoresist. We have tried to wet etch the protrusions away before reapplying the photoresist without success. Another drawback of this method is the fact that the photoresist shrinks after a hard baking. A mirror placed on the photoresist before baking will wrinkle terribly. However, one placed on the resist after a hard bake will not wrinkle too badly after subsequent hard bakes. Unfortunately, the photoresist does not planarize the wires, and the wire pattern and wire surface roughness is mapped onto the mirror. The photoresist is removable with acetone in an ultrasonic bath, and once baked is compatible with vacuums down to at least 10^{-9} Torr and perhaps slightly lower.

A similar, but much better technique—no shorting problems—is to use polyimide (Kapton) coatings in place of the photoresist. Polyimide is extremely viscous, and when spun onto a substrate and hard baked, forms a tough yellowish protective coating. With a single spin, coatings up to 10 microns can be obtained, and with hard baking the film is UHV compatible and acts to planarize wires or other protrusions on the substrate. It is this last attribute that makes polyimide widely used in the semiconductor industry. We purchased our polyimide, Pyralin PI2560, from HD Microsystems. The thickness for a spin of 2000 rpm is 8 μm and increases to 11 μm for a 3000 rpm spin. (The company sells a similar polyimide, PI2562, that coats up to 2 μm .) Application requires an adhesion promoter, VM-652, bought in addition to the polyimide. The coating procedure is as follows. Pipet the VM-652 onto the substrate as it is held in the vacuum chuck of the spinner. Wait for 20 seconds and spin dry for 30 seconds. The spin speed can be the same as used in the polyimide application step. Bake the VM-652 coated substrate on a hot plate for a minute at 120°C . Be careful not to crack a custom-cut sapphire substrate—a temperature ramp might be required. Placing the substrate back on the vacuum chuck, pipet enough polyimide to cover roughly half the surface. The spinning will cause it to cover the rest. The rotation should start at 500 rpm for 5 seconds before rotating at final speed—2000 to 3000 rpm—for 30 seconds. Remove substrate, and with a towel soaked in acetone, quickly wipe-off any polyimide coating the bottom of the substrate. To prevent flowing, immediately place the substrate on a $\sim 100^\circ\text{C}$ hot plate for an initial cure of roughly 5 minutes. One will notice that the polyimide surface becomes smoother as it bakes on the hot plate. To fully cure the polyimide and prepare it for UHV chambers, it must be hard baked to 350°C for an hour. A slow ramp is required to prevent substrate cracking. The Munich group has noticed bubbles forming in the polyimide, but we have never seen this problem. It might have been

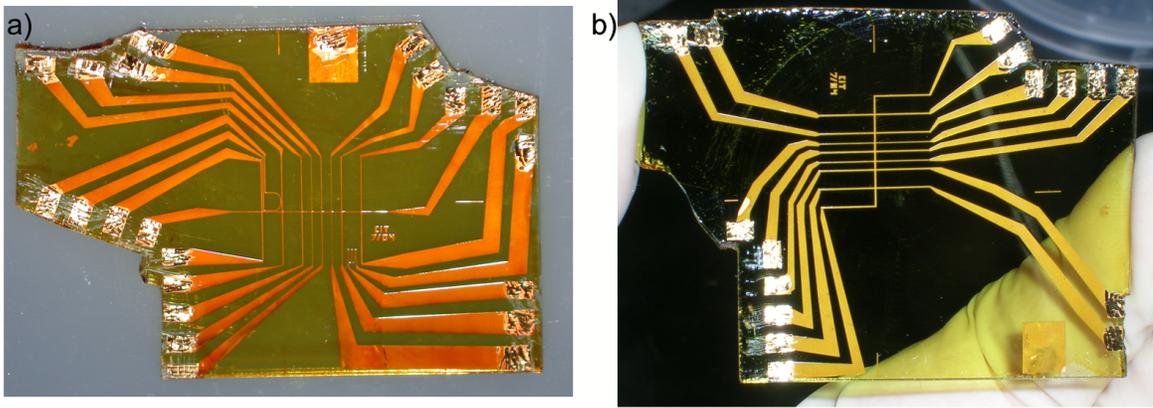


Figure 4.2: Polyimide-coated waveguide atom chips fabricated with electrodeposited gold on custom-cut sapphire. The polyimide has been removed above the wire pads revealing the shiny gold underneath. The substrates have a maximum width of 5 cm in the horizontal direction, and the wires are 12 to 14 microns tall. a) Waveguide with P-trap loading. Minimum wire width is $50\ \mu\text{m}$. Three coats of polyimide. b) Directly loaded waveguide from U-trap. Minimum wire width of $100\ \mu\text{m}$. A single coat of polyimide.

due to an expired polyimide sample: they bought a new batch of PI2560 and the problem never arose again. The main vat of Pyralin PI2560 must be kept refrigerated, and a sample in a room temperature bottle will expire in a month or so and should be discarded.

We found that for 12 to 14 μm tall wires, the polyimide spun at 2000 rpm produced a 50% planarization (only a $\sim 6\ \mu\text{m}$ bump remained). A soft bake followed by another coating resulted in a 40% further planarization, but a third soft bake and application didn't change the height of the bumps. Rather, the bumps simply became wider. After a hard bake, the bump height returned to $\sim 6\ \mu\text{m}$: the polyimide shrunk slightly. Hard baking between each polyimide application would work much better. Wires much shorter than 10 microns would be much easier to planarize, and the Munich group has found this to be the case. The polyimide surface forms a suitable surface for additional microwires, allowing the creation of multilayered chips. Cured polyimide is easy to flake-off the surface of the substrate, which is necessary for uncovering the wire pads, but seems impervious to acetone. Coating with polyimide forms a UHV compatible and easily cleanable protective layer for the delicate microwires and should, if possible, be used. Figure 4.2 show two sapphire substrates coated with polyimide.

Our first attempt at making a mirror MOT using a gold-coated polyimide and 12 μm tall wires proved unsuccessful. Although it formed a smooth, specular surface, the wire bumps—

several 100 microns in width—misdirected enough laser power to prevent the MOT from forming. Further improvements in planarization will mitigate this problem (for instance by hard baking between each layer application), and enable the use of polyimide as a desirable technique for making atom chip mirrors.

Recently, we have been making atom chip mirrors by simply gold coating a thin—130 to 170 μm —glass cover slide, and attaching it to the chip either face up or down with a small drop of photoresist. The mirror is only 90% reflective when placed upside down due to the Cr adhesion layer. A mirror MOT was made with such a mirror and the low laser scattering allowed nice fluorescence images of the atoms to be taken. We have noticed that Cs from our oven begins to discolor (through accumulation, presumably) the gold mirror after a few months of exposure. A mirror MOT may still be formed, but the mirror scatters more light into the imaging CCD cameras. Using a Cs dispenser might help to improve this situation.

4.3 Specific fabrication techniques: wet etching, ion milling, lift-off method, and electroplating

The minimum required wire dimensions vary significantly depending on the atom chip's application, and an optimal fabrication technique should be chosen accordingly. This section gives the recipe and discusses the relative merit of each fabrication method.

4.3.1 Wet etching and ion milling

The simplest chip to fabricate has wire widths no smaller than 30-40 μm and wire heights less than 1 μm . A transparency mask should be used for the photolithography (see Section 4.2.1). The wire height is set by a thermally evaporated or sputtered gold layer and the photoresist masks the gold intended for wires from the wet etch solution (see Figure 4.3[a]). To begin the procedure, the cleaned substrate should be placed in the ozone dry stripper for five minutes at 65° C to ensure that no organic material will prevent the adhesion of chromium and gold. The thermal evaporation step follows, with the thickness of the gold layer determined by chip's current density requirements. Because the photoresist adheres well to gold, only a 5 min bake at 180° C is necessary for adhesion (this temperature might be too high for cut sapphire substrates). Wet etching removes exposed gold, and the pho-

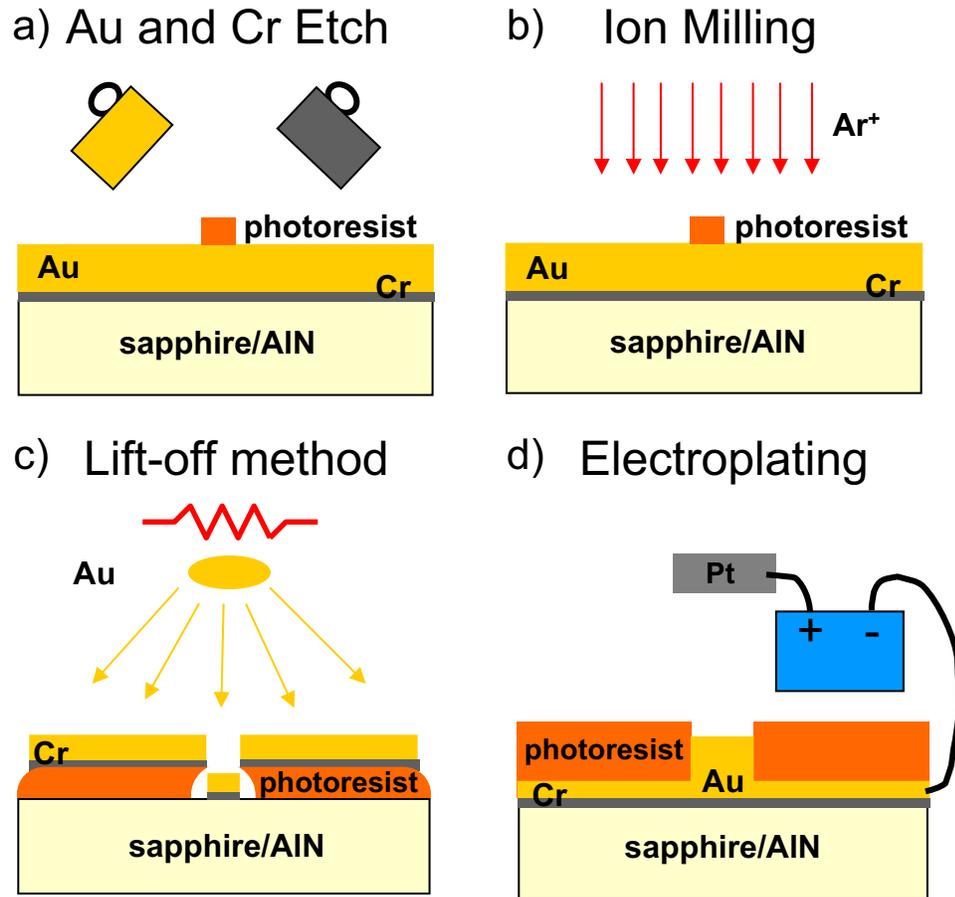


Figure 4.3: Fabrication techniques. a) Patterned positive photoresist masks the gold layer from the gold and chromium wet etch. b) The argon ions mill away the gold not covered by positive photoresist. c) Gold is thermally evaporated into the trenches patterned in the negative photoresist. The undercut allows the photoresist and unwanted gold to separate from the substrate without peeling away the gold in the trenches. d) Wires are defined by gaps in the positive photoresist, and the walls of the photoresist guide the wires as they are electroplated. After electroplating, acetone removes the photoresist and gold and chromium wet etches remove the seed layer.

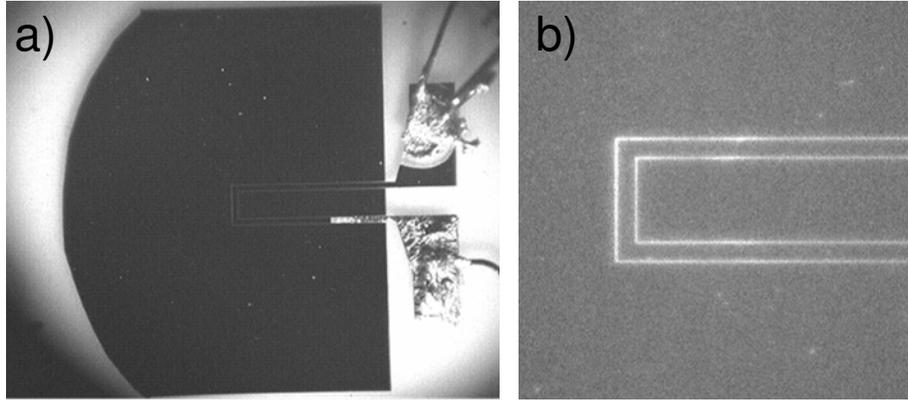


Figure 4.4: Gold wire patterned using the wet etch technique. (a) This atom chip contains a quadrupole trap in the U configuration. The gold wire, patterned on sapphire and surrounded by a gold mirror, is $300\ \mu\text{m}$ wide and $1\ \mu\text{m}$ tall. (b) Close-up of the wire region. The gold appears darker than the uncovered sapphire substrate.

toresist should be patterned so that it covers the areas intended for wires, i.e. the photoresist should be a positive image of the wire pattern. A photomask on which the wires are opaque, used in conjunction with positive photoresist, will produce a positive image of the wire pattern. We use the photoresist AZ5214 (Clariant), which can serve as both a negative and positive photoresist depending on the bake and exposure procedure. The positive process recipe is as follows: spin coat at 5000 rpm for 50 s, bake at 95°C for 2 min, expose for 10 to 20 s, and develop in AZ327 MIF (or some similar developer) for 30 s. All of the above times are approximate and will vary depending on the UV light intensity of the specific mask aligner and on various environmental conditions such as humidity. It may be necessary to try various exposure and bake times to find the optimal recipe. These exposure times are based on the $16\ \text{mW}/\text{cm}^2$ UV intensity of our mask aligner. To remove the gold not covered by photoresist, submerge the substrate in gold etch solution (Gold Etchant TFA, Transene Company, Inc., telephone 978-777-7860) for a few tens of seconds until only the dull gray of the chromium layer remains. Finally, remove the chromium layer with chrome etchant (CR-7S, Cyantek, Co., telephone 510-651-3341). Figures 4.4 (a) and (b) show a substrate patterned in this manner. The wet etch dissolves the gold isotropically, and the decrease in wire width is insignificant for wires and wire gaps larger than 10 to $20\ \mu\text{m}$. Of course, transparency masks cannot be used for features smaller than a few tens of microns.

Ion milling can be a useful alternative to wet etching. Instead of removing the unwanted gold with an etch solution, argon ions anisotropically bombard the surface, removing the gold not covered by photoresist (see Figure 4.3[b]). This method can produce very narrow features, limited only by photoresist resolution, with heights determined by the thermally evaporated gold layer. The photoresist is also milled, but this is of no consequence as long as it is thicker than the gold layer. The substrate may become quite hot during the ion etching, and one needs to be careful that the substrate does not overheat, causing the photoresist to become hard and difficult to remove. We have used ion milling to make atom chips as well as to etch a common hard drive for use as a magnetic atom mirror (see Chapter 8 and Reference [12]).

4.3.2 The lift-off method

The quick and easy wet etch technique is unfortunately not suitable for wire widths smaller than $20\ \mu\text{m}$, and ion milling machines are not readily available. The lift-off method should be used for the case in which the wires need not be taller than $1\ \mu\text{m}$ but less than $20\ \mu\text{m}$ wide. However, if the surface quality of the sub-10 micron wires is important to the application (e.g., for BEC experiments), then the lift-off method will be worthwhile regardless of the height of the wires (see Section 4.4 for more details about these constraints).

In contrast to the wet etch technique, the photoresist in this method is used as a mask for the deposition of thermally evaporated gold. Trenches are created in a negative photoresist using a photomask with opaque wires, and evaporated gold deposits both into the trenches, adhering to the substrate, and onto the surface of the photoresist (see Figure 4.3[c]). If done properly, the walls of the trenches have an overhang—which looks like an undercut when viewed from above—that prevents the unwanted gold on the photoresist from connecting to the gold in the trenches. An acetone bath dissolves the photoresist, allowing the unwanted gold to lift-off leaving the wire pattern formed from the gold in the trenches.

After cleaning the substrate, the AZ5214 is spun on the substrate for 45 s at 5000 rpm. The maximum height of the thermally evaporated wires is set by the thickness of the photoresist since lift-off will not work once the top of the gold connects with the gold on the overhang. We have been able to achieve lift-off with wires $1.5\ \mu\text{m}$ tall by spinning the photoresist on at 2000 rpm and thermally evaporating many boats of gold over a period of three to four hours. The photoresist should then be baked for 45 s at 100°C , UV exposed

with the photomask for 10 s, baked again for 45 s at 123° C, UV exposed with no mask for 2.1 min, and developed for 25 to 35 s. Developing is finished when one can see the wire pattern in the photoresist. A successful undercut can be seen in a microscope as a bright outline of the edges of the trenches. Before thermal evaporation, the substrate should be placed in the ozone dry stripper at 65° C for 5 minutes. This removes unwanted material that could prevent gold adhesion, and does not seem to hamper photoresist removal as in the electroplating process described below. To promote lift-off, the acetone bath should be heated on a hot plate, and the substrate, while inside the beaker, should be sprayed with an acetone squirt bottle. It is very important that all of the gold-coated photoresist be peeled away before the substrate is removed from the acetone. Otherwise, once dried, the unwanted gold flakes become extremely difficult to separate from the surface. Difficulty in achieving lift-off may be overcome by briefly exposing the substrate to ultrasound. This is risky, however, since the gold wires might be stripped-off as well. Figure 4.1 (b) shows an atom chip fabricated with the lift-off method.

4.3.3 Electroplating

The above methods rely on thermal evaporation to achieve the required wire thickness. This limits the wire heights to $\sim 1 \mu\text{m}$. Electroplating the wires can increase the wire height considerably: for example, we have made $3 \mu\text{m}$ wide wires $4 \mu\text{m}$ tall, and $50 \mu\text{m}$ wide wires $14 \mu\text{m}$ tall. Thick photoresist spun and patterned on a thin gold seed layer provide a template for the growth of the wires. The walls of the photoresist maintain a constant wire width as the wire height increases (see Figure 4.3 [d]). An acetone wash followed by a brief wet etch removes the photoresist and gold seed layer. Electroplating is a tricky process that does not always produce reliable results. We provide here a general guideline for the process, and with this process we have typically been able to achieve a wire height accuracy of $\pm 0.5 \mu\text{m}$.

Fabrication begins with cleaning and ozone dry-stripping the substrate, followed by the thermal evaporation of a 100 to 150 nm seed layer of gold along with a 50 \AA chromium adhesion layer. For proper vertical wall guiding of the wires, the photoresist must always be taller than the electroplated wires, and a photoresist thicker than that one used in the aforementioned techniques is necessary. Clariant's AZ9200 series photoresists are 4 to 24 microns thick, and can achieve aspect ratios of 5 to 7 with resolutions of $< 1 \mu\text{m}$ to $3.5 \mu\text{m}$

depending on the resist thickness. After spin coating, it should be baked on a hot plate at 110° C for two minutes, and then the photoresist should be UV exposed for 60 s (or longer depending on the photoresist thickness) using a photomask with transparent wire patterns. The resist is developed in a 1:4 solution of AZ400K and water for 10 seconds to a minute depending on the exact solution concentration: the exposed photoresist will turn hazy before dissolving away. The gold seed layer also acts as the cathode in the electroplating process, and some of the photoresist must be wiped away with acetone—or a blank spot should be designed in the photoresist—to serve as a contact for the cathode lead. An ozone dry etch is then used to remove any layers of HMDS, photoresist, or organics that might mask regions of the gold from the electroplating solution. The time and temperature of this process is crucial: too long an exposure at too high of a temperature will make the photoresist difficult to remove between closely spaced wires, and too short an exposure will not remove enough unwanted masking material. For example, we found that an 18 s room-temperature ozone dry etch was optimal for removing unwanted material while also enabling the removal of photoresist between wires spaced by 3 μm .

Alternatively, if a rectilinear wire cross-section is not desired, then the cathode may be connected to the wires on the substrate post-acetone removal. The wires may be formed from performing lift-off, wet-etch, or ion milling and should all be shorted together with connections that are subsequently scratched-out. These electroplated wires grow horizontally as well as vertically with roughly semicircular cross sections.

We use an ammonium gold sulfite solution from Metakem GmbH for the electroplating. A sodium gold sulfite solution from Technic, Inc. did not work as well. Specifically, it would turn brownish (from clear) during a deposition, preventing its reuse and causing large towers of gold to form on the wire. The Metakem solution, with 15 g of Au/liter, is poured into a roughly 1200 mL container and heated to $\sim 65^\circ\text{C}$. The anode, also purchased from Metakem, is platinized titanium (type B mesh, size 10 \times 10 cm). Place the anode mesh vertically into the container so that a part of it is above the solution. The depth of the container should be so that the sample, when suspended vertically in the solution, is completely submerged: any excess solution will be unnecessarily subject to evaporation. Attach the positive alligator clip to the anode mesh outside the solution and the negative clip to the substrate's gold seed layer (the cathode). The cathode alligator clip usually has to be in contact with the solution for the substrate to be completely submerged. The

current should be off when the substrate is submerged and turned on or off gently thereafter. We have found that a 1 mA current does not activate the deposition, but a 20 mA current deposits 6 to 7 microns of gold per 20 min. With a 40 min deposition, 12 to 14 μm tall wires can be made. The substrate should be gently agitated while electroplating to promote even plating and suppress the formation of $\sim 5 \mu\text{m}$ tall towers of gold. After electroplating, the substrate, anode, and container should be rinsed with water. The gold solution can be reused, but should be filtered if flakes of material appear.

Both our group and the Munich group noticed an uneven, shadow-like effect on the cross-sectional height of electroplated wires using the Metakem solution. The section of an exposed wire next to a steep photoresist wall will not electroplate as well if the solution is stirred so that the wall creates a fluidic shadow on the wire: the wire's cross-section is no longer rectilinear, but rather thins on one side. Reversing the direction of fluid rotation creates wire thickness shadows on the wire's opposite side. Agitating the substrate in random directions helps to mitigate this effect, and the Munich group saw that it might be preferential to not stir the solution at all. They also noticed that the use of a fresh batch of solution seems to mitigate this problem.

To prepare for wet etching the seed gold layer away, the photoresist should be removed in a room-temperature acetone bath. Sometimes it is difficult to remove the photoresist between wires spaced only several microns from one another, and in these cases the substrate—while in the acetone—should be placed in an ultrasound for a few minutes. The gold should not peel away since it is attached to the entire substrate surface. After rinsing the acetone away with IPA and methanol, the gold seed layer is removed with a ~ 15 s wet etch. The chromium adhesion layer should also be wet etched away. Occasionally, the air jet does not remove all of the methanol from the substrate, and tiny drops of methanol can sometimes dry on the leeward side of the wires. This dried methanol acts as a mask for the gold etch, leaving small puddles of seed layer that can short adjacent wires. These puddles can be removed by rinsing with methanol, blow-drying from a different angle, and briefly wet etching a second time. The surface reflectance of the gold is typically diminished after the wet etch, and a mirror fabricated with this gold may not be ideal. A second photolithography step can add photoresist on top of the wires and/or areas of the seed layer to protect them from the wet etchant. The protected seed layer is suitably reflective for forming a mirror MOT.

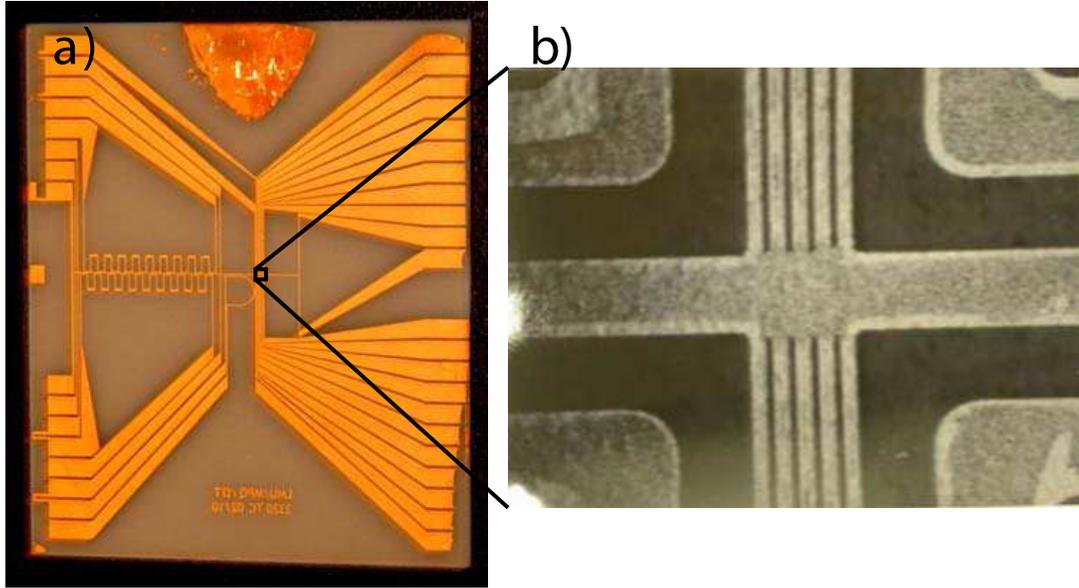


Figure 4.5: An atom chip-based BEC interferometer fabricated by electroplating onto an AlN substrate. a) The chip can produce a BEC and transport it to the center region where b) five wires $3\ \mu\text{m}$ wide, $4\ \mu\text{m}$ tall, and spaced by $3\ \mu\text{m}$ can split the BEC in a double-well potential.

A surface profilometer, commonly known as an alpha step machine, is quite useful for quickly measuring the height of the wires. Inevitably, a few substrates must be spent optimizing the electroplating process for a specific wire height. Figures 4.5 (a) and (b) show an atom chip-based BEC interferometer that we fabricated by electroplating on an AlN substrate². The smallest features are five, 1 mm long wires that are each $3\ \mu\text{m}$ wide, $4\ \mu\text{m}$ tall, and spaced less than $3\ \mu\text{m}$ from one another.

4.4 Trap fragmentation

In 2002, it was discovered that BECs in an elongated, cigar-shaped trap held within tens of microns from the wire surface can fragment into “sausage-link” sections [93, 94, 95, 62, 96]. This is due to the roughness and meandering of the wire surface which results in the deviation of the wire current away from linear propagation along the wire. An errant magnetic field 10^3 to 10^4 times smaller than the trapping field causes the BEC to fragment into local potential valleys. In the last two years, groups determined that

²Wire pattern designed by T. Steinmetz and P. Hommelhoff at the MPQ/LMU in Munich. See Chapter 10 for more details.

electroplating wires is much worse for making smooth wires and is the main culprit behind this fragmentation problem [97, 91]. Thermal evaporation and sputtering are superior for creating smooth wires and should be used in experiments that confine BECs in tight traps near the wire surface.

4.5 Conclusion

The techniques described in this paper provide a basic starting point for the design and fabrication of atom chips. The precise control of atomic position enabled by these chips is quite crucial to many areas of research. Moreover, these devices allow an incredible miniaturization of experiments involving cold atoms. From constructing atom optical elements to studies of BECs and cavity QED, atom chips are proving invaluable to the fields of atomic physics, quantum optics, and quantum computation.

Chapter 5

Cold Collisions and Cesium Microtrap Losses

Magnetically trapped cesium has an exceptionally large susceptibility to collisional losses, greatly hampering the ability to perform evaporative cooling long enough to achieve a BEC. In fact, a cesium BEC was first formed only after the employment of an optical dipole trap in 2003, eight years after the first dilute gas BECs in rubidium and sodium [98].

When we tightly confine cesium in magnetic microtraps, we must pay attention to its collisional loss mechanisms to ensure efficient trapping and loading. This chapter is taken—with additional references and comments—from an unpublished set of notes the author wrote in late 2002 which describe the collisional loss mechanisms we face when trapping and cooling alkali atoms. Particular emphasis is paid to cesium and its peculiarities. We also document an analysis of U-MOT and U-trap decay dynamics in a cesium-based experiment performed in the fall of 2002. To the best of our knowledge, this constitutes the first collisional rate measured with an atom chip, and the measured value is consistent with those found from previous experiments from other researchers using different experimental systems. These measurements greatly aided our understanding of U-MOT performance and helped us to optimize trap loading efficiencies.

5.0.1 Collision Basics

Since the late 1980's, the ability to trap, laser cool, and evaporatively cool alkali atoms to the μK regime and below has enabled researchers to study collisional physics in this hitherto unexplored regime [99, 46]. In turn, a better understanding of cold collisions has led to many other achievements. A clear understanding of trap losses led to the production of an atomic

BEC by way of evaporative cooling, and collision studies uncovered the Feshbach resonances of cesium which explain why a cesium BEC has been so elusive until recently [100, 101, 98].

The method of partial waves is often used to describe collisions. The Hamiltonian for a collision between two particles is

$$\mathcal{H} = -\frac{\hbar^2}{2\mu R^2} \frac{d}{dR} \left(R^2 \frac{d}{dR} \right) + \frac{\hbar^2 l(l+1)}{2\mu R^2} - \frac{C_n}{R^n}, \quad (5.1)$$

where μ is the reduced mass, the second term is the centrifugal energy for a partial wave l , and the last term is the interaction potential [46]. Using the above Hamiltonian, the Schrödinger equation can be solved for each partial wave. The cross section is

$$\sigma = \frac{4\pi}{k^2} \sum_{l=0}^{\infty} (2l+1) \sin^2 \delta_l, \quad (5.2)$$

where δ_l is the phase shift between incoming and outgoing wave functions. At the low temperatures of cold collisions, only the s-wave, $l=0$, partial wave contributes to the cross section. In this case it is useful to define a scattering length,

$$a = -\lim_{k \rightarrow 0} \frac{\delta_0(k)}{k}, \quad (5.3)$$

and the cross section becomes $\sigma = 8\pi a^2$ [46].

In the last few years, much interest has been paid to the variance of the scattering length due to Feshbach resonances. These resonances occur when the scattering excited state is nearly degenerate with a bound molecular state with different quantum numbers [102]. An external magnetic field can shift the energy levels of these states with respect to one another, which results in a tunable scattering length [103]. Cesium has many low magnetic field Feshbach resonances, and these cause its scattering length to vary dramatically (to even negative values) over a magnetic field range of zero to a couple hundred Gauss [101, 104, 102, 105, 106]. Reference [107] provides a thorough and up-to-date discussion of cesium Feshbach resonances. The tunability of the scattering length has enabled research groups to manipulate the atomic interactions in a BEC. Specifically, in tuning the interactions from repulsive to attractive, C. Weiman's group has seen the collapse and explosion of rubidium condensates. Reference [108] and [109] provide information on these resonances in rubidium.

Manipulation of BECs via Feshbach resonances has led to the formation of molecular

BECs and BCS gases [110, 111]. Unfortunately, these Feshbach resonances are the cause of the exceptionally high ground-state collisional rates in cesium. By understanding these resonances and exploiting them using external magnetic fields, researchers have been able to suppress ground-state collisions, enabling them to achieve Bose-Einstein condensation in cesium [112].

5.1 Cold collision processes

The terms cold collisions and ultra-cold collisions are often used interchangeably in the literature. One can, however, distinguish between the two regimes [99]. Cold collisions occur between 1 mK to 1 μ K which is accessible with laser cooling, and the ultra-cold regime begins below 1 μ K and extends to the lowest temperatures achievable by evaporation cooling, approximately 10 nK. In the ultra-cold regime, research has focused on elastic collisions since these rates determine the efficiency of evaporative cooling. This section will instead focus on trap losses in the cold-collision regime. In general, trap losses may be categorized as follows: excited-state collisions, diffusive loss, background collisions, ground-state collisions, and Majorana spin-flip losses.

5.1.1 Excited-state collisions: MOT loss mechanisms

Excited-state collisions are the main decay channel for atoms trapped in a MOT [46, 99]. These light-assisted collisions occur when two atoms collide while one atom is in the P excited state, and in this case the potential is dominated by a $-C_3/R^3$ dipole-dipole interaction as opposed to the shorter range $-C_6/R^6$ van der Waals ground-state potential. At the Condon point,

$$R_c = \left(\frac{C_3}{\hbar|\delta|} \right)^{1/3}, \quad (5.4)$$

the quasimolecule is excited with the red-detuned, $-\delta$, trapping light (see Figure 5.1). The Condon point is typically 1000 to 2000 a_0 for δ equal to a few γ .

Radiative escape and fine-structure changing collisions are two exoergic processes that may lead to MOT losses. At low temperatures (less than 1 mK), the collision time of the atoms is comparable to their decay lifetime. After excitation at the Condon point the quasimolecule may decay while $R < R_c$, and the emitted photon is lower in energy by

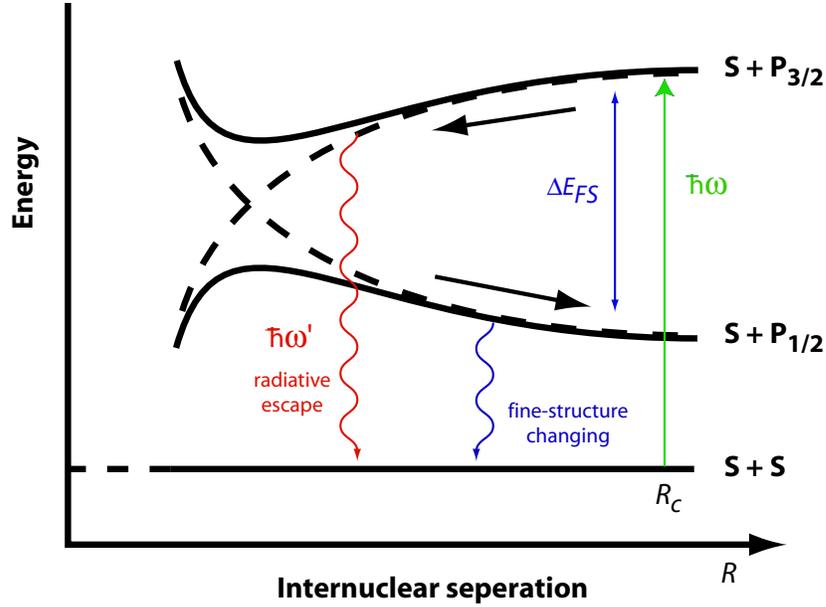
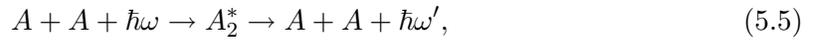


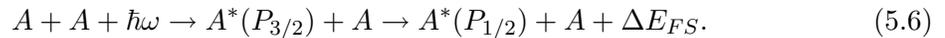
Figure 5.1: Excited-state collisions. Figure adapted from References [99] and [46].

$C_3/R_c^3 - C_3/R^3$. This radiative escape (RE) process:



may release enough kinetic energy to expel the atoms from the MOT.

As the atoms near their closest approach, the asymptotic $S_{1/2} + P_{3/2}$ state mixes with the asymptotic $S_{1/2} + P_{1/2}$ state. If the quasimolecular state does not decay before this occurs, then atoms can exit the collision in the lower energy fine-structure level:



This is known as a fine-structure changing collision (FCC), and the energy imparted to the atoms is well over tens of Kelvin—more than enough to eject the atoms from the trap.

Measurements of the excited-state loss rate in cesium vary from $\beta_{es} = 1$ to $20 \times 10^{-11} \text{ cm}^3/\text{s}$ [113, 114, 115, 116, 117]. The first measurement of these collisions, made by Weiman's group, showed that β_{es} rises from $2 \times 10^{-12} \text{ cm}^3/\text{s}$ at $4 \text{ mW}/\text{cm}^2$ of laser power to $1 \times 10^{-11} \text{ cm}^3/\text{s}$ at $12 \text{ mW}/\text{cm}^2$. Below $4 \text{ mW}/\text{cm}^2$ they observed a sharp increase in the loss rate to $\sim 8 \times 10^{-11} \text{ cm}^3/\text{s}$ at $2 \text{ mW}/\text{cm}^2$. They attribute the low intensity

loss mechanism to hyperfine-changing collisions (HCC) between ground-state atoms: when the optical intensity is lowered, the atoms spend more time in their ground states and these collisions release enough energy to allow the atoms to overcome the lower trap depth (see Section 5.1.4). Saturation intensity for the $F=4$ to $F'=5$ cycling transition in cesium is 1.09 mW/cm^2 , and typically the total power from the six MOT beams is at least 12 mW/cm^2 —well into the FCC and RE loss regime. Subsequent measurements by Libbrecht, Julienne, and Gomer’s groups find β_{es} in the range mentioned above, and Julienne’s group in Reference [114] measured the FS rate to comprise 25% the total excited-state loss rate. Other alkali atoms generally have a lower β_{es} . For instance, the values of β_{es} at 10 mW/cm^2 for ^{85}Rb and ^{87}Rb are $3.4 \times 10^{-12} \text{ cm}^3/\text{s}$ and $1.0 \times 10^{-12} \text{ cm}^3/\text{s}$, respectively, and the ^7Li loss rate is $\sim 1.0 \times 10^{-12} \text{ cm}^3/\text{s}$ [99].

Optical shielding is a technique to suppress excited-state collisions. A blue-detuned laser can prevent the colliding atoms from reaching the attractive potential Condon point. See Reference [104] for a collision suppression technique based on tuning cesium’s Feshbach resonances.

The earliest description of cold, excited-state collisions is the semi-classical GP-model proposed by Gallagher and Pritchard in 1989. The JV-model—a more quantum mechanical extension to the GP-model—was introduced by Julienne and Vigué two years later [118, 119]. In both models, the collisional rate constant is factored into two terms: a probability for quasimolecular excitation and survival of this excitation to short distances; and the probability for a trap loss collision occurring given that the atoms are in the quasimolecular state at close range. These models form the basis for the interpretation of many experimental results as well as the basis of more recent and sophisticated models. Please see References [99, 46] for details regarding the GP- and JV-model and for current theoretical techniques.

5.1.2 MOT diffusion losses

Libbrecht’s group measured the MOT loss rate due to stochastic diffusion of atoms out of the trap at high quadrupole field gradients [116]. The diffusive losses are incorporated into the atom number rate equation by adding a $-N/\tau_{diff}$ decay term. Diffusive loss dominates for above gradients of 1 kG/cm , and this loss limits MOT sizes to $> 5 \text{ }\mu\text{m}$. Specifically, $\tau_{diff} = 1 \text{ s}$ for $\nabla B = 1.5 \text{ kG/cm}$ and $I/I_{sat} = 1.5$.

5.1.3 Background losses

Residual atoms in the vacuum chamber collide and expel trapped atoms, and these losses affect both the steady-state number of atoms in a MOT and the decay rate of a magneto-static trap. In a vacuum typically used in trapping experiments, 10^{-9} Torr or lower, these residual atoms mainly consist of the background vapor of alkali atoms being trapped (in our case, cesium). In this regime, we can neglect the non-cesium atoms, and the MOT state equation becomes

$$dN/dt = R - N/\tau_{cs}, \quad (5.7)$$

where R is the capture rate and $1/\tau_{cs}$ is the loss rate due to Cs-Cs background collisions. The capture rate and the loss rate are

$$R = 1/2nV^{2/3}v_c^4(m/2kT)^{3/2}, \quad (5.8)$$

$$1/\tau_{cs} = n\sigma_{cs}(3kT/m)^{1/2}, \quad (5.9)$$

where V is the trapping volume, n is the cesium vapor density, v_c is the capture velocity, σ_{cs} is the cross section for trap ejection due to background collisions, and T is the cesium vapor temperature [120]. Typical values for a cesium MOT are: $V = (0.1 \text{ cm})^3$; $v_c \sim 15$ m/s; and $\sigma_{cs} = 2 \times 10^{-13} \text{ cm}^2$. Solving Equation 5.7 gives $N(t) = N_{ss}(1 - \exp^{-t/\tau_{cs}})$, and the steady state number of atoms, N_{ss} , is independent of the cesium vapor pressure

$$N_{ss} = R\tau_{cs} = V^{2/3}/(\sqrt{6}\sigma_{cs})v_c^4(m/2kT)^2. \quad (5.10)$$

One can measure the lifetime, τ_{cs} , by recording the ramp-up time of the MOT. This, in turn, provides a good measurement of the cesium background pressure, P_{cs} , using the following expression,

$$P_{cs} = \frac{\sqrt{kTm}}{\tau_{cs}\sigma_{cs}\sqrt{3}}. \quad (5.11)$$

As long as P_{cs} is the largest contribution to the vacuum chamber pressure, than the P_{cs} independence of N_{ss} allows one to lower P_{cs} while maintaining a high MOT atom number. This becomes useful if the atoms are transferred to a purely magnetic trap in which the trap loss is dominated by the $-N/\tau_{cs}$ term. The loading time must increase, however, and this might be prohibitive for experiments requiring the collection of many data points.

If the vacuum pressure is not dominated by P_{cs} , then N_{ss} is no longer independent of P_{cs} . An elevated vacuum pressure is usually caused by light gases unremoved by the turbo and ion pumps during bake-out. Helium has a larger cross section with cesium than the lighter H_2 and is consequently the dominant concern [115]. The He-Cs collisional rate equation is

$$\frac{1}{N} \frac{dN}{dt} = -n_0 \beta_{he}, \quad (5.12)$$

where β_{he} is the loss rate in cm^3s^{-1} , and n_0 is the He background density. Willems *et al.* in Reference [115] calculated that $\beta_{he} \cong 2 \times 10^{-9} \text{ cm}^3\text{s}^{-1}$.

Adding a He-Cs collisional loss term to Equation 5.7, we find

$$dN/dt = R - N/\tau_{cs} - N/\tau_{he} = R - N/\tau^*, \quad (5.13)$$

where

$$\tau^* = \frac{\tau_{cs}\tau_{he}}{\tau_{cs} + \tau_{he}} \quad \text{and} \quad \tau_{he} = \frac{kT}{P_{he}\beta_{he}}. \quad (5.14)$$

The new steady-state atom number, $N_{ss}^* = R\tau^*$, equals

$$N_{ss}^* = \frac{V^{2/3}}{\sqrt{6}\sigma_{cs}} v_c^4 \left(\frac{m}{2kT} \right)^2 \frac{1}{1 + \frac{P_{he}\beta_{he}}{P_{cs}\sigma_{cs}} \sqrt{\frac{m}{3kT}}}. \quad (5.15)$$

Figure 5.2 plots N_{ss}^* for $V = (0.1 \text{ cm})^3$ and $v_c = 8 \text{ m/s}$. High helium background pressures can be a serious limit to MST lifetime: helium not only adds to background collisions but also increases the P_{cs} required to trap a large number of atoms in the MOT. This is a major concern in making a BEC on a chip as discussed in Chapter 10.

5.1.4 Ground-state loss mechanisms

Magnetic traps cannot capture atoms in their lowest energy ground-state. Earnshaw's theorem tells us that it is impossible to create a magnetic field maximum in free space, and consequently, atomic strong-field seeking states are impossible to trap magnetostatically. However, it is possible to trap weak-field seeking atoms in a magnetic field local minimum. For example, cesium may be trapped in the $F=4, m_F=4$ or $F=3, m_F=-3$ Zeeman levels of its hyperfine ground-states. Because the $F=3, m_F=3$ state always has the lowest energy in a magnetic field, magnetically trapped states are unstable. Magnetic dipole-dipole in-

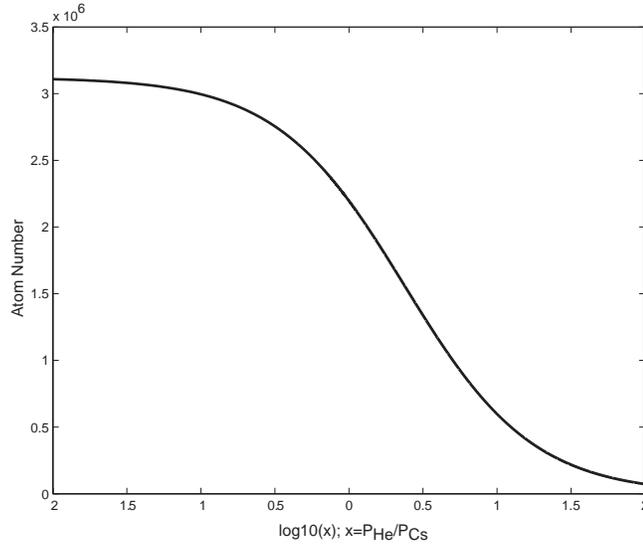


Figure 5.2: N_{ss}^* as a function of the ratio between the He and Cs pressure. Plotted for parameters similar to our experiment.

teractions between two colliding atoms cause spin relaxation to smaller or untrapped $|m_F|$ states. Moreover, the $F=4$ state is susceptible to magnetic dipole-dipole mediated hyperfine changing collisions (HCC) to the $F=3$ state. The energy released is large enough to expel the atoms from the MST. This is particularly severe in the case of cesium: the 9.2 GHz hyperfine splitting transfers 0.22 K worth of energy to each atom.

For sodium and rubidium, the collisional loss rates due to HCC and spin relaxation are in the 6 to $10 \times 10^{-15} \text{ cm}^3 \text{ s}^{-1}$ range, and this is low enough to allow BEC formation in magnetic traps even in the upper hyperfine state. Unfortunately, cesium's ground-state loss rate is a factor of 100 larger, and its hyperfine splitting is also much larger. The collisional loss rate for cesium in the $F=4$, $m_F=4$ ground-state is

$$\beta(T)_{4,4} = (1.5 \pm 0.3 \pm 0.3) \times 10^{-11} T^{-0.63} \text{ cm}^3 \text{ s}^{-1}, \quad (5.16)$$

and the loss rate for $F=3$, $m_F=-3$ is

$$\beta(T)_{3,-3} = 2.2 \times 10^{-12} T^{-0.78} \text{ cm}^3 \text{ s}^{-1}, \quad (5.17)$$

where T is in μK [121, 100]. Using optical dipole traps, collisional loss rates of

$$\beta_{4,4} = (1.1 \pm 0.1 \pm 0.2) \times 10^{-11} \text{cm}^3 \text{s}^{-1} \quad (5.18)$$

and

$$\beta_{3,4} = (1.5 \pm 0.2 \pm 0.3) \times 10^{-12} \text{cm}^3 \text{s}^{-1} \quad (5.19)$$

are reported in Reference [122], the latter for Cs(F=3)-Cs(F=4) HCC collisions. For the type of Ioffe trap used in Reference [100], the loss rate can be expressed as

$$\beta(B, T)_{3,-3} = (38 \pm 12) \times 10^{-12} B^2 T^{-0.78} \text{cm}^3 \text{s}^{-1}, \quad (5.20)$$

where B , the bias field, is in mT and is related to the trap oscillation frequency by $\omega_{xy} \propto B^{-1/2}$.

These high collisional loss rates have thwarted all attempts to attain BEC in cesium with magnetic trapping. Optically trapping cesium in the F=3, $m_F=3$ lowest energy ground-state has recently led to a cesium BEC in the Innsbruck group of Rudy Grimm and Christoph Nägerl [98]. Three-body collisions, in which two atoms form a molecule and the third carries away excess energy, lead to the expulsion of all three atoms and is an obstacle to effective evaporative cooling. Measurements indicate that this rate is $1.5 \times 10^{-25} \text{cm}^6 \text{s}^{-1}$ [100, 123]. Recent experiments have measured loss rates due to heteronuclear loss rates in simultaneously trapped cesium and rubidium [124] and cesium and lithium [122].

5.1.5 Majorana spin-flip losses

Quadrupole fields are easy to produce using a pair of coils with opposite currents, and just a few amperes can generate a large enough field gradient to trap atoms in a MOT. Unfortunately, the quadrupole trap is susceptible to a loss mechanism, and consequently, there are no bound states for the atoms. A quadrupole trap works while the atom's spin adiabatically follows the field as the atom moves: the Larmor frequency must be larger than the rate of change of the magnetic field experienced by the atom,

$$\frac{g\mu_0 r \nabla B}{\hbar} > \frac{v}{r}, \quad (5.21)$$

where v is the velocity of the atom, and r is its distance from the trap center. If this adiabatic condition fails, then the atom's spin can no longer follow the field and may spin-flip to an untrapped state. This occurs when the atom passes near the zero of the quadrupole field. Here the magnetic field becomes so weak that the Larmor frequency approaches zero while the atom's velocity remains relatively constant. This Majorana spin-flip loss mechanism prevents long-term trapping in a magnetic quadrupole trap. Replacing the quadrupole field with an Ioffe-Pritchard trap, which has a non-zero field minimum, solves this problem. For cesium, an Ioffe-Pritchard trap with a 1 G field minimum is sufficient to prevent non-adiabatic spin-flips.

To estimate the Majorana spin-flip loss rate one needs to find the flux of atoms through the surface defined by $r_b \sim v\hbar/g\mu_0\nabla B$. For an only slightly elliptical trap, the density of atoms is roughly $3N/4\pi l^3$, and the velocity of atoms as a function of the cloud radius, l , is found by application of the virial theorem: $mv^2 = g\mu_0 l \nabla B$. Dividing the trap volume by both the atom velocity and the above surface area, we find the decay time to be $\tau = (m/3\hbar)l^2$, which for cesium traps of radii 500 μm and 50 μm is 180 s and 1.8 s, respectively. During evaporative cooling, the slower atoms spend more time near the center of the trap, and this causes the loss rate to quickly increase [125].

5.2 Cesium microtrap losses

As mentioned in the introduction to this chapter, our main purpose for investigating alkali trap loss mechanisms is to determine to what extent this limits our ability to trap cesium in a high gradient magnetic microtrap. The following sections discuss cesium trap loss mechanisms as they relate to the U-MOT and U-trap used in our experiments. A major goal is to explain and mitigate the large U-MOT loss rate we measured in our microtrap experiment.

5.2.1 U-MOT losses

A major difficulty we face in these experiments is to efficiently load atoms from the mirror MOT into tight microtrap waveguides and Ioffe traps (such as the tiny Libbrecht-style traps). The use of a U-MOT is an intermediary step that allows the precise transfer of the mirror MOT atoms into the U-trap. One could load the U-trap directly from the mirror

MOT, but this requires a precise shifting of the MOT's external quadrupole field to the position of the U-trap during the transfer. By comparison, loading atoms from the mirror MOT to the U-MOT is more robust due to the continual presence of the MOT light forces as the quadrupole fields are exchanged. Once the atoms are in the U-MOT, sub-doppler cooling can be performed and the U-trap turned on without much worry about spatial mode matching since the U-wire magnetic field for the U-MOT and U-trap are one and the same.

Our current procedure utilizes a substrate-based U-trap to form a U-MOT in the vicinity of the mirror MOT. The atoms are transferred to the initial U-MOT with 100% efficiency, but they still need to be lowered much closer to the substrate surface to obtain the magnetic field gradients sufficient to form a U-trap. We find that lowering the U-MOT precipitates a rapid loss of atoms. The measurements and analysis in this section were performed to obtain an understanding of this process and collisional processes in microtraps in general. The knowledge gained enabled us to design a more efficient loading procedure that is discussed in Chapter 2.

Figure 5.3 shows a typical U-MOT decay curve as the trap height decreases and the trap compresses. We believe the atom loss is due to excited-state collisions. The trap compression exacerbates the excited-state collisional loss, and the smaller trap surface area decreases the loading rate from vapor atoms. This lower loading rate leads to a much smaller N_{ss} , and the trap cannot recover to its original number of steady-state atoms.

With regards to maximizing microtrap loading efficiency, we learned two lessons from these measurements: 1) MOT light should be extinguished as early as possible to prevent excited-state collisional losses exacerbated by a compressed trap. The cesium atoms should be transferred to the U-trap at a distance higher from the surface to minimize the required U-MOT compression. This is accomplished by using a U-wire capable of supporting 2 A of current, allowing the capture into the U-trap of atoms high above the substrate surface. 2) The mirror MOT should be as spatially mode-matched to the U-MOT as possible to minimize MOT transfer time. The mirror MOT in this experiment was formed using a quadrupole field produced by coils external to the vacuum chamber. This technique was abandoned in subsequent experiments in favor of using a macro U-MOT formed by a half-centimeter scale block of copper in the shape of a U and positioned underneath the atom chip. This has the capability for trapping atoms from vapor while providing a quadrupole profile much more similar to that used by the on-chip U-MOT, thus maximizing the spatial

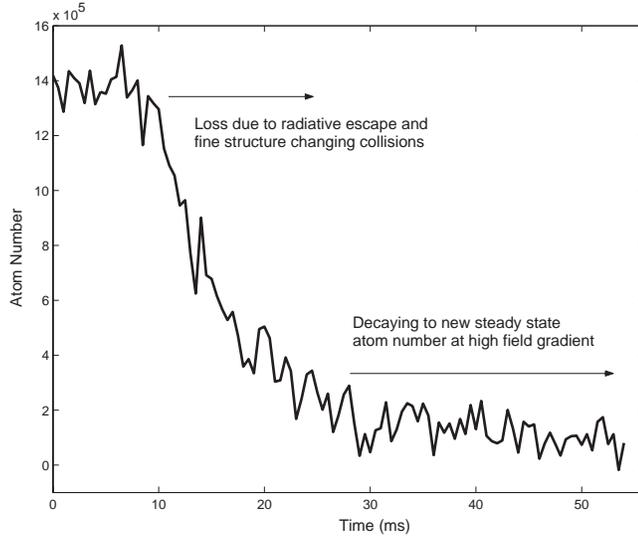


Figure 5.3: U-MOT decay upon compression. At $t = 0$, the microwire is turned on and the bias field ramps to full value in $t = 2$ ms. The external quadrupole field ramps off by $t = 2$ ms. $\nabla_y \mathbf{B}$ varies from 5 G/cm to 90 G/cm.

mode matching during the transfer. Moreover, we no longer had to worry about accurately aligning the center of the large, exterior mirror MOT coils to the $\sim 100 \mu\text{m}$ sized U-wire feature inside the vacuum chamber: the macro U-MOT is pre-registered to the U-MOT microtrap to provide nearly turn-key trap loading.

The following analysis may provide a model for future microtrap-based collisional rate measurements. Though the quality of the data could be improved-upon in a subsequent experiment—by using absorption instead of the less accurate fluorescence imaging, for instance—we believe the conclusions and methodology remain valid.

To test the excited-state collision explanation for the observed U-MOT decay, we measured the decay for different gradients in both tightly-confined directions, \hat{y} and \hat{z} , and in the weakly-confined direction, \hat{x} , for three different final trap heights, $y_{final} = 1, 0.5,$ and 0.33 mm (see Figure 5.4)¹. The gradients and trap minimum in \hat{y} of the U-MOT are not independent of one another and the following are only approximate expressions:

$$\begin{aligned} \nabla_y \mathbf{B}_{approx} &= \frac{2\pi}{\mu_0} \frac{B_{bias}^2}{I} \\ y_{0 approx} &= \frac{\mu_0}{2\pi} \frac{I}{B_{bias}}. \end{aligned} \quad (5.22)$$

¹ $y_{final} = y_0(t \rightarrow \infty)$ is the final resting position of the atoms in \hat{y} .

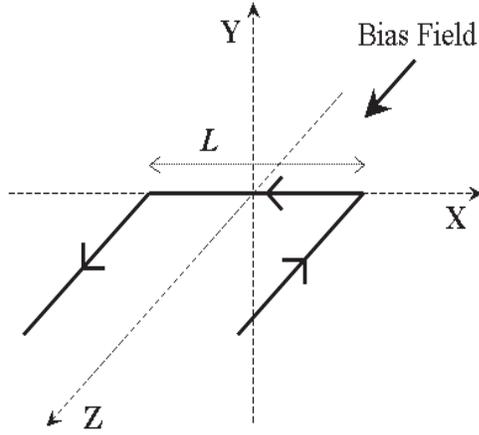


Figure 5.4: U-trap coordinate system.

Generally the gradients in both \hat{x} , \hat{y} , and \hat{z} increase as the final trap height decreases (see Figures A.3, A.4, and A.5 and Appendix A for non-approximate expressions for the \hat{y} gradient, $\nabla_y \mathbf{B}$, and position of the field minimum, y_0). This gradient increase compresses the trap and by accounting for this volume compression in the trap population equation, we arrive at a description of the U-MOT trap loss. The trap population equation is

$$\frac{dN}{dt} = R(t) - \frac{N}{\tau_{back}} - \frac{\beta N^2}{V(t) + V_0}. \quad (5.23)$$

The $R(t)$ term in Equation 5.23 describes the trap loading from background vapor. The capture cross section decreases as the surface area of the trap decreases, and this leads to a diminished loading rate. Comparing a 1-D model to experimental data, References [126, 127] calculate how the capture rate depends on the trap gradient. The capture rate, R , is proportional to the trap surface area, S , multiplied by the fourth power of the capture velocity, v_c [120]. The papers by Meschede's group find that v_c^4 varies as $(\nabla B)^{-2/3}$, and because S is proportional to $(\nabla B)^{-2}$, we find that $R \propto (\nabla B)^{-8/3}$. The gradient varies as a function of time in our experiment, and this adds a time dependence to the capture rate:

$$R(t) = \frac{N_s}{\tau_{back}} \frac{\nabla_y \mathbf{B}(0)^{8/3}}{\nabla_y \mathbf{B}(t)^{8/3}}, \quad (5.24)$$

where at $t = 0$, R equals the MOT steady-state value of N_{ss}/τ_{back} . This model is applicable for an isotropic MOT, but the U-MOT is highly elongated in one direction: we

use the gradient in the tightly-confined \hat{y} direction in the above expression since the gradient in \hat{x} is negligible. Neglecting the gradient in the weakly-confined \hat{x} direction does not significantly change the results of this section because the excited-state collisional term in Equation 5.23 dominates during the time scales associated with the U-MOT loss (see Figures 5.10 and 5.11).

The second term in Equation 5.23 accounts for the background collisions. The background collision rate is equal to $(0.3 \pm 0.1 \text{ s})^{-1}$ for the experiments described in this section.

The excited-state collisional loss term—the last one in Equation 5.23—accounts for the rapid atom loss that occurs as the U-MOT compresses. $V(t)$ is the trap volume as a function of time and β is the collisional loss rate. As mentioned earlier, increasing the trap gradients shrinks the volume of the trap which exacerbates this collisional loss. Experimentally we see that the U-MOT remains cigar-shaped throughout compression and a slice perpendicular to the long axis of the U-MOT remains roughly circular. Therefore, to obtain an expression for $V(t)$ we may replace the gradient in \hat{z} with that in \hat{y} , and write the volume as a function of the gradient in the tightly-confined direction, \hat{y} , and the gradient in the weakly-confined direction, \hat{x} , as:

$$V'(t) \propto 1/(2\pi[\nabla_y \mathbf{B}(0, y(t), z(t), B_{bias}(t), I)]^2 \cdot 2\nabla_x \mathbf{B}(0, y(t), z(t), B_{bias}(t), I)). \quad (5.25)$$

The first three indices of \mathbf{B} are for the centroid of the atoms as a function of time where $x(t)$ has been taken to equal zero. Note: the positions of the atoms, $y(t)$ and $z(t)$, are not necessarily the position of the field minimum, y_0 and z_0 , at time t , since atoms track this minimum with a time delay. The last two indices are for the U-trap's bias field and wire current, respectively. Using $\nabla_z \mathbf{B}(0, y(t), z(t), B_{bias}(t), I)$ in this expression instead of squaring the gradient in \hat{y} is problematic since this gradient vanishes for certain values of y_0 (see Figure A.5 in Appendix A and the corresponding text). At $t = 0$, the experimentally measured volume is $V'_0 \cong 9 \times 10^{-5} \text{ cm}^3$, and the volume as a function of time is

$$V(t) = \frac{V'(t)V'_0}{V'(0)}. \quad (5.26)$$

The actual trap volume does not follow $V(t)$ for large t because of heating. The V_0 constant in Equation 5.23 is an effective final volume added to account for this heating.

The time dependence of the volume and capture rate arises from the variation of the magnetic field gradients that the atoms feel as they settle into the compressed U-MOT as it is brought closer to the substrate surface. Qualitatively, as the bias field ramps-up to B_{bias} in 2 to 3 ms, the trap field minimum shifts to a predetermined position less than a millimeter above the substrate surface. From the mirror MOT height of 3 to 4 mm, the atoms slide down the U-trap magnetic field potential until they are damped by the trapping lasers to zero average velocity at the trap field minimum, $y_{final} = y_0$. To model $V(t)$ and $R(t)$ we need to calculate $\nabla_i \mathbf{B}(t)$: we must account for the time variation of the bias field as we evaluate the trap gradients at each position of the atoms. Specifically, we must measure or calculate $y(t)$, $z(t)$, $B_{bias}(t)$, and $\nabla_i \mathbf{B}(0, y, z, B_{bias}, I)$ to find $\nabla_i \mathbf{B}(0, y(t), z(t), B_{bias}(t), I)$. The expressions for the trap gradients as function of position, current, and bias field are listed in Appendix A.

At $t = 0$, the wire current is instantaneously turned on, and a computer sends a 2 ms linear ramp to the power supply for the bias field coils. The bias coils have significant inductance, and the power supply can only ramp the current with a 0.5 ms time constant. Convoluting the power supply response with the linear ramp, we find

$$B_{bias}(t) = \frac{B_{bias}}{\tau_r} (t - \tau + te^{-t/\tau}), \quad (5.27)$$

for $t \leq \tau_r$ and

$$B_{bias}(t) = (B_{bias} - B_{bias}(\tau_r))(t - \tau + te^{-t/\tau}) + B_{bias}(\tau_r), \quad (5.28)$$

for $t > \tau_r$, where τ_r is the linear ramp time, B_{bias} is the target bias field value, and τ is the power supply response time. Figure 5.5 shows $B_{bias}(t)$ along with the ideal bias ramp.

From trap images we can measure the position of the atoms above the substrate as a function of time, $y(t)$. Figure 5.6 shows several experimental $y(t)$ curves for different values of y_{final} , I , and B_{bias} . Generally, smaller y_{final} and larger final trap gradients increase the rate at which the atoms sink to the bottom of the trap. If atoms exactly followed the position of the trap field minimum $y_0(t)$ (see the Appendix A for the calculation of $y_0(t)$ as a function of $B_{bias}(t)$), then we would expect $y(t)$ to behave as shown in Figure 5.7 after accounting for the time variation of B_{bias} . This clearly does not mimic the experimental

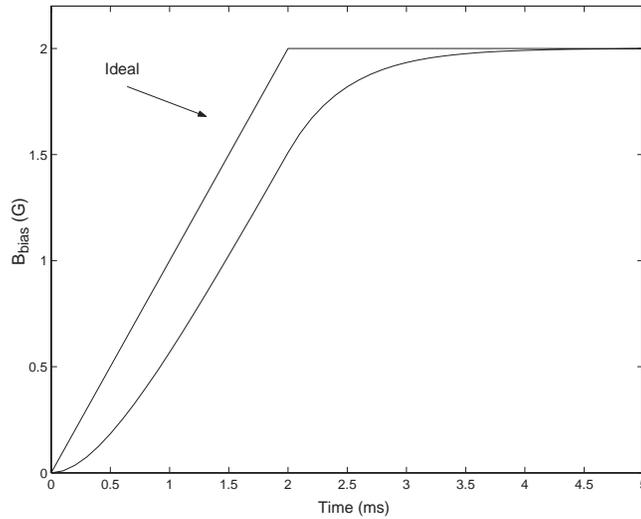


Figure 5.5: B_{bias} response to a $\tau = 2$ ms linear current ramp. The power supply has a $\tau_r = 0.5$ ms power supply response time.

data shown in Figure 5.6, and even though we use the experimental $y(t)$ data for finding $V(t)$ and $R(t)$, we would like to have a better understanding of what determines $y(t)$. A reasonable model for the force in the \hat{y} direction on an atom at the center of the trap includes a potential from the magnetic field gradient and a damping force due to the trapping lasers:

$$\ddot{y} + \frac{\gamma}{m} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-((y-y_0(t))/\sqrt{2}\sigma)^2} \dot{y} + \frac{\mu_b}{m} \nabla_y \mathbf{B}(0, y, 0, B_{bias}(t), I) \cdot \text{sign}(y - y_0(t)) + g = 0, \quad (5.29)$$

where γ is the damping coefficient. The atoms experience a large magnetic field when they are far from the trap minimum, $y(t) \neq y_0(t)$, and this field shifts the Zeeman levels away from resonance. Combining this effect with the misalignment of the lasers with respect to magnetic field when $y(t) \neq y_0(t)$ causes the strength of the damping force to diminish. We model this by multiplying the damping term in Equation 5.29 by a Gaussian that turns-on the damping when the atoms are within a σ of $y_0(t)$. Comparing Figures 5.8 and 5.9, we see that this model qualitatively reproduces the experimental data for the same wire current and bias field. In Figure 5.9 the upper curve is $y(t)$, the lower curve is the velocity in mm/s, $\sigma = 0.1$ mm, and $\gamma = 5.1 \times 10^{-23}$ kg/s (which was calculated using our trapping laser's intensity and detuning).

We now have all the ingredients to solve Equation 5.23: $y(t)$ comes from 4th order

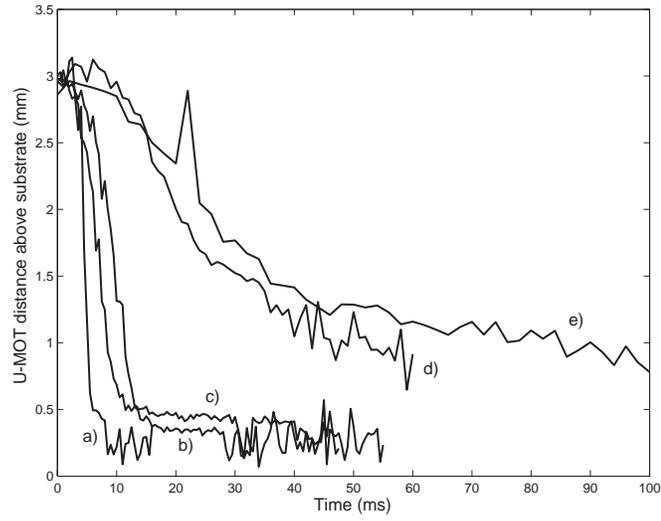


Figure 5.6: Experimental data of the U-MOT distance above substrate versus time for several values of the current and bias field. Curve: a) $y_{final} = 0.33$ mm; $I = 2$ A; $B_{bias} = 12$ G. b) $y_{final} = 0.33$ mm; $I = 0.5$ A; $B_{bias} = 3$ G. c) $y_{final} = 0.5$ mm; $I = 0.5$ A; $B_{bias} = 2$ G. d) $y_{final} = 1$ mm; $I = 2$ A; $B_{bias} = 4$ G. e) $y_{final} = 1$ mm; $I = 0.5$ A; $B_{bias} = 1$ G.

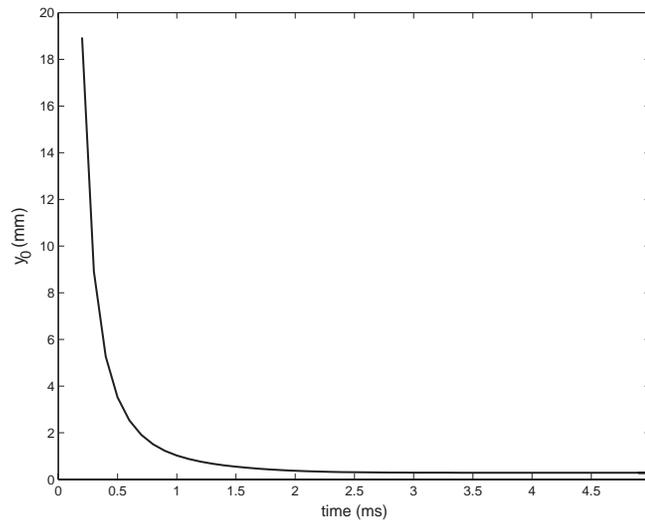


Figure 5.7: y_0 versus time as $B_{bias}(t)$ ramps-up in $\tau = 2$ ms and $\tau_r = 0.5$ ms. $y_{final} = 0.28$ mm; $I = 1$ A; $B_{bias} = 6$ G.

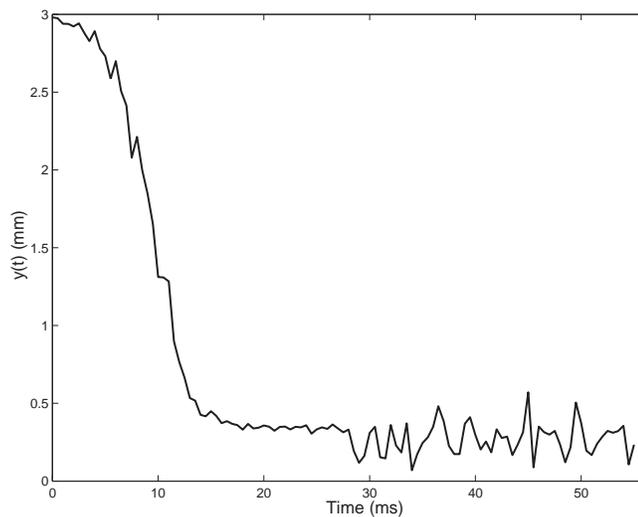


Figure 5.8: U-MOT distance above substrate versus time. The trap parameters are $y_{final} = 0.33$ mm; $I = 0.5$ A; $B_{bias} = 3$ G; $\nabla_{y,z}\mathbf{B} = 90$ G/cm; $\nabla_x\mathbf{B} = 11$ G/cm.

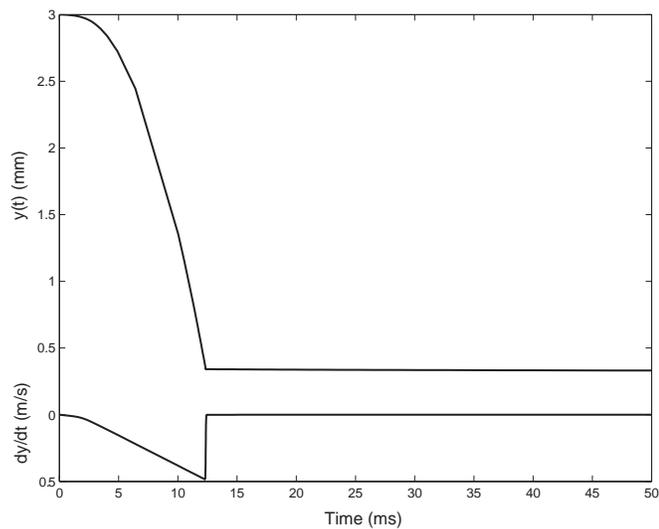


Figure 5.9: Calculated distance and velocity of an atom at the U-MOT center. $I = 0.5$ A; $B_{bias} = 3$ G.

rational fits to experimental data such as that displayed in Figure 5.6, Equations 5.27 and 5.28 express B_{bias} as a function of time, and analytical expressions for trap gradients in \hat{x} , \hat{y} , and \hat{z} for a given $y(t)$, $B_{bias}(t)$, and I are given in Appendix A. The data for each U-MOT decay curve was taken for different combinations of I and B_{bias} , which resulted in final trap gradients ranging from $\nabla_y \mathbf{B} = 10$ to 360 G/cm and final trap heights of $y_{final} = 0.33$ mm to 1 mm.

We fit each decay curve (using the corresponding values for I and B_{bias}) with Equation 5.23, allowing β and V_0 to be free parameters. The bootstrap method is used to obtain error bars for each data set [128]: We assume a model for the errors and adjust its parameters so that the resulting reduced χ^2 's for the fits equal one. The error model contains three terms:

$$\sigma^2(N) = \left(m \cdot 100 \times 10^3 \text{ atoms} \frac{N}{N_0} \right)^2 + \left(s \cdot \sqrt{\frac{N(1 + 4(\Delta/\gamma)^2 + 6I_0/I_{sat})}{\gamma/2 \cdot 6I_0/I_{sat} \cdot \pi/d^2}} \right)^2 + (g \cdot 20 \times 10^3 \text{ atoms})^2, \quad (5.30)$$

where $\Delta = 10$ MHz is the detuning from resonance, $\gamma = 5.2$ MHz is the spontaneous emission rate, and the saturation parameter, I_0/I_{sat} , equals two. The first term in Equation 5.30 describes the mirror MOT atom number variation due to loading. The next term accounts for the variation in detected atom number caused by shot noise, and last term is the noise due to background light and electronic noise in the imaging system. The parameters m , s , and g are free fit parameters, and are varied to satisfy the equation $\chi^2(\sigma)/\nu = 1$ for each data set, where ν is the number of data points minus the number of fit parameters. Averaging the values obtained for each U-MOT decay curve gives $m = 1.4 \pm 0.4$, $s = 0.5 \pm 0.5$, and $g = 2.8 \pm 0.9$.

The U-MOT population equation, Equation 5.23, is in the form of a Riccati equation, $f'(x) = q(x)f(x)^2 + p(x)f(x) + r(x)$, which cannot be solved in general [129]. If we can justify neglecting the $r(x)$ term, then this equation assumes the form of a Bernoulli equation, $f'(x) = q(x)[f(x)]^n + p(x)f(x)$, with $n = 2$. Bernoulli equations are easily solved by using the substitution $u(x) = [f(x)]^{1-n}$. This results in the linear differential equation, $u'(x) = -p(x)u(x) - q(x)$, which has the general solution

$$u(x) = \frac{-\int_0^x \mu(x)q(x)dx + u_0}{\mu(x)}, \quad (5.31)$$

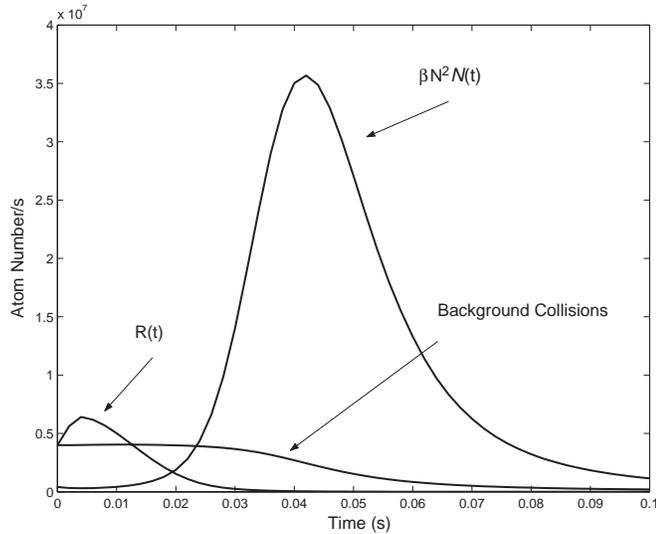


Figure 5.10: The loading and collision terms in Equation 5.23 for a low-gradient trap of: $I = 0.5$ A, $B_{bias} = 1$ G, and $\nabla_y \mathbf{B} = 10$ G/cm.

$$\mu(x) = e^{\int_0^x p(x) dx}. \quad (5.32)$$

Neglecting the $R(t)$ term in Equation 5.23 and using the above solution to the Bernoulli equation, we find that

$$N(t) = \frac{N_0 e^{-t/\tau_{back}}}{N_0 \beta \int_0^t \frac{e^{-t/\tau_{back}}}{V(t)+V_0} dt + 1}, \quad (5.33)$$

with N_0 being the initial atom number. Is it possible to neglect the loading term? We see from Figure 5.10 that in a low gradient trap $R(t)$ is only significant in the beginning of compression, and Figure 5.11 shows that $R(t)$ is insignificant in a high gradient trap. Qualitatively, these figures demonstrate that it is safe to disregard the loading term. Moreover, using the population equation without the loading term and with the errors determined by Equation 5.30, we still find reduced χ^2 's very close to unity for all of the data sets (see Figure 5.12).

We use the Levenberg-Marguardt method, supplying our own Jacobian, to fit the non-linear model to our U-MOT decay data sets. The derivatives of Equation 5.33 needed to form the Jacobian matrix are

$$\frac{dN}{d\beta} = \frac{-N_0^2 e^{-t/\tau_{back}} \int_0^t \frac{e^{-t/\tau_{back}}}{V(t)+V_0} dt}{\left(N_0 \beta \int_0^t \frac{e^{-t/\tau_{back}}}{V(t)+V_0} dt + 1\right)^2}, \quad (5.34)$$

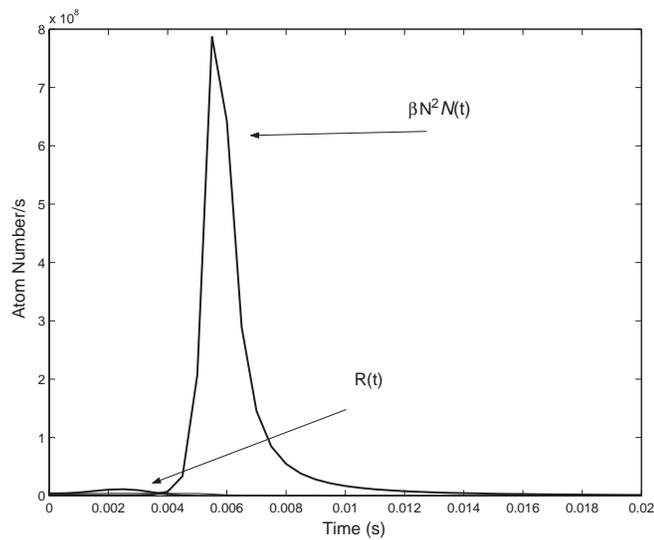


Figure 5.11: The loading and collision terms in Equation 5.23 for a high-gradient trap of: $I = 1.5$ A, $B_{bias} = 9$ G, and $\nabla_y \mathbf{B} = 270$ G/cm.

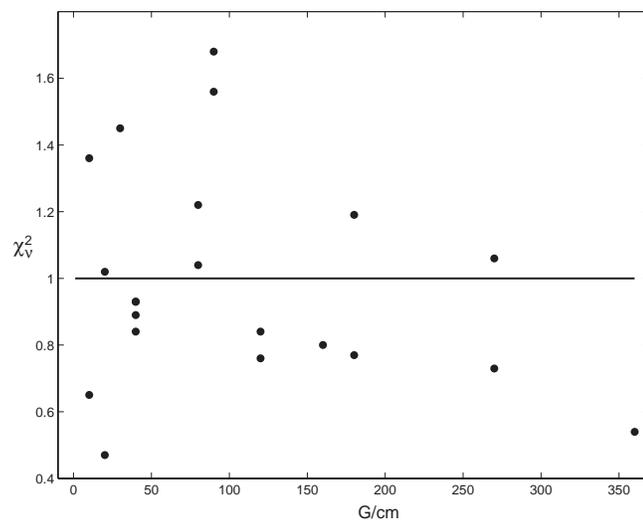


Figure 5.12: The reduced χ^2 without the $R(t)$ loading term in Equation 5.23.

$$\frac{dN}{dV_0} = \frac{-N_0^2 \beta e^{-t/\tau_{back}} \int_0^t \frac{e^{-t'/\tau_{back}}}{(V(t')+V_0)^2} dt'}{\left(N_0 \beta \int_0^t \frac{e^{-t'/\tau_{back}}}{V(t')+V_0} dt' + 1\right)^2}. \quad (5.35)$$

Qualitatively, the fits are remarkably successful throughout the large parameter range for U-MOT gradients and heights. Figure 5.13 shows typical U-MOT decay curves with the fits superimposed. We expect the value of β to remain constant for all the fits regardless of the current and bias field for the trap, and it should also be consistent with previously measured values of cesium excited-state losses (see Section 5.1.1). Indeed, we do not see much variation of β with respect to trap parameters (see Figure 5.14). There does appear to be some small increase of β with gradient, but this is less than a factor of ten over the large variation of $\nabla_y \mathbf{B}$. The fit value for V_0 does not seem to vary much for U-MOTs of different heights and final gradients. Figure 5.15 shows these values for V_0 . The error bars take into account the covariance of V_0 with β . The fractional error in each fit value varies between 0.2 and 0.8 depending on the degree of the covariance between β and V_0 . The errors in the fitted values originate from the errors in measured atom number as described in Equation 5.30, and because the overall level of uncertainty in atom number varies little between data sets, we expect the fractional errors in β and V_0 to not vary by more than a factor of ten between data sets. This relatively constant fractional error explains the correlation between small fit values and small error bars for β and V_0 . Moreover, there does not seem to be anything obvious about the raw data or curve fits that would suggest an alternative explanation for this correlation.

The covariance matrices for each fit can be combined to produce an error estimate on the mean of β and V_0 . The expressions for the weighed mean and combined covariance matrix are

$$\mu = \frac{\sum \mu_i / \sigma_i^2}{\sum 1 / \sigma_i^2}, \quad (5.36)$$

$$\bar{\sigma}_\mu^2 = \frac{1}{\sum 1 / \sigma_i^2}, \quad (5.37)$$

where σ_i^2 is the error on each mean value and $\bar{\sigma}_i^2$ is the covariance matrix for each fit [130]. Figure 5.16 shows the confidence interval of 68.3%, and the weighted mean values of the fit parameters with $1\text{-}\sigma$ errors are $\beta = (4.1 \pm 0.3) \times 10^{-11} \text{ cm}^3\text{s}^{-1}$ and $V_0 = (7.8 \pm 0.7) \times 10^{-8} \text{ cm}^3$. If we repeat the fitting procedure with fixing β at the above value, then we find that

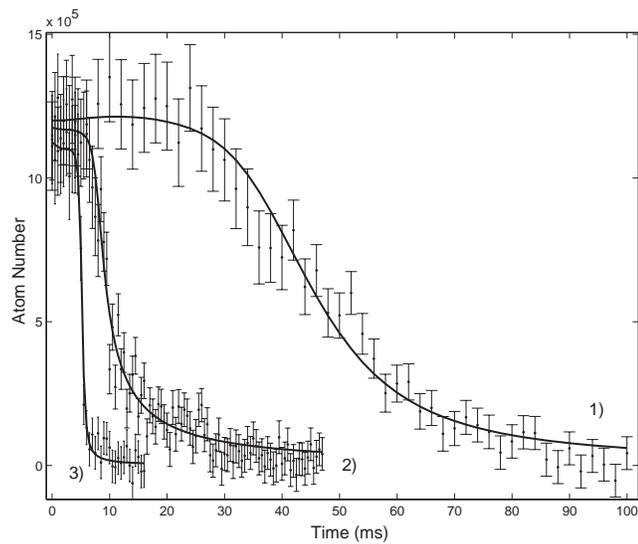


Figure 5.13: Typical fits to U-MOT decay curves. 1) $y_{final} = 1$ mm; $I = 0.5$ A; $B_{bias} = 1$ G. 2) $y_{final} = 0.5$ mm; $I = 0.5$ A; $B_{bias} = 2$ G. 3) $y_{final} = 0.33$ mm; $I = 2$ A; $B_{bias} = 12$ G.

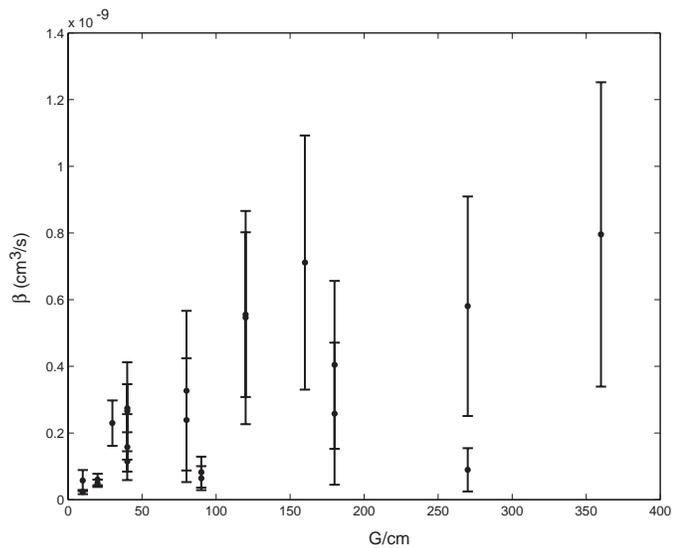
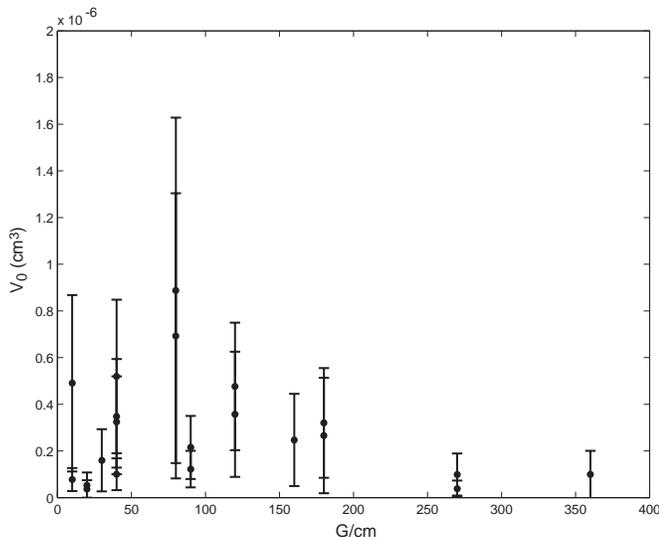


Figure 5.14: U-MOT loss rate, β , versus trap gradient.

Figure 5.15: V_0 versus trap gradient.

the weighted mean of V_0 is a factor of seven lower and the mean χ^2 increases to 2.3.

Although the above double parameter fit is optimal for minimizing the χ^2 for all the various trap parameters, physically we expect β to be a fixed value. We repeated the fitting procedure while setting β to a fixed value. We compiled V_0 fits for several values of β to find the fixed β that minimizes the the root mean square of $\chi^2 - 1$. As shown in Figure 5.19, a minimum exists for a β value approximately equal to $12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$: the optimal V_0 is equal to $4.8 \pm 0.2 \times 10^{-8} \text{ cm}^3$. Figure 5.17 shows the values of V_0 as a function of trap gradient when $\beta = 12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$. The weighted mean and error of V_0 is $4.8 \pm 0.2 \times 10^{-8} \text{ cm}^3$. Figure 5.18 shows the reduced χ^2 for each fit. The degree to which the mean of these χ^2 's deviates from unity provides some measure of the quality of fit for this specific fixed β .

Our values for the excited-state loss rate are consistent with that of the previous measurements described in Section 5.1.1. Moreover, the values for V_0 are consistent with what we measure from trap images. We feel that our model, Equations 5.23 and 5.29, does a reasonable job of explaining U-MOT loss and dynamics during capture and compression from a mirror MOT.

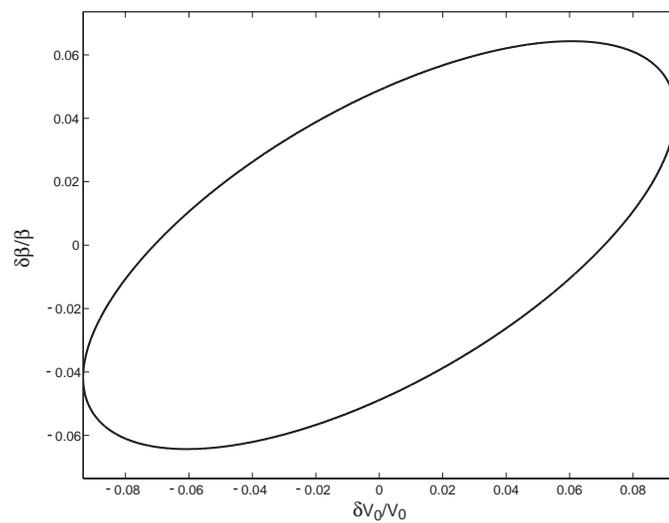


Figure 5.16: Covariance of β and V_0 at a confidence interval of 68.3%. $\beta = 4.1 \pm 0.3 \times 10^{-11} \text{ cm}^3/\text{s}$; $V_0 = 7.8 \pm 0.7 \times 10^{-8} \text{ cm}^3$.

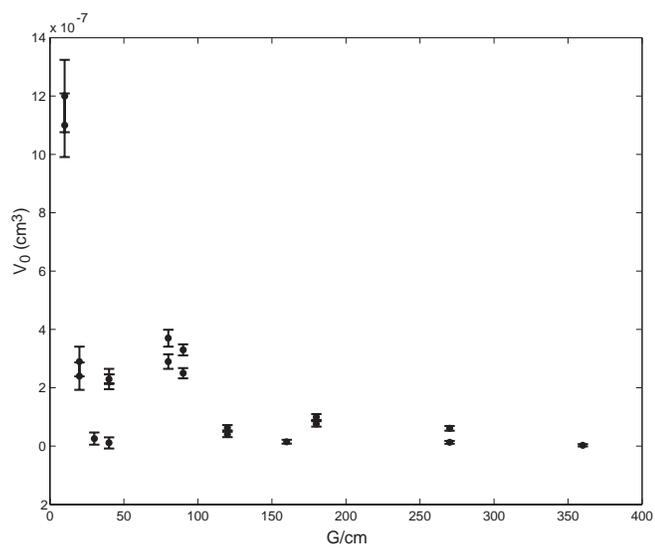


Figure 5.17: V_0 with β fixed at $12.5 \times 10^{-11} \text{ cm}^3 \text{ s}^{-1}$. The V_0 weighted mean is $4.8 \pm 0.2 \times 10^{-8} \text{ cm}^3$.

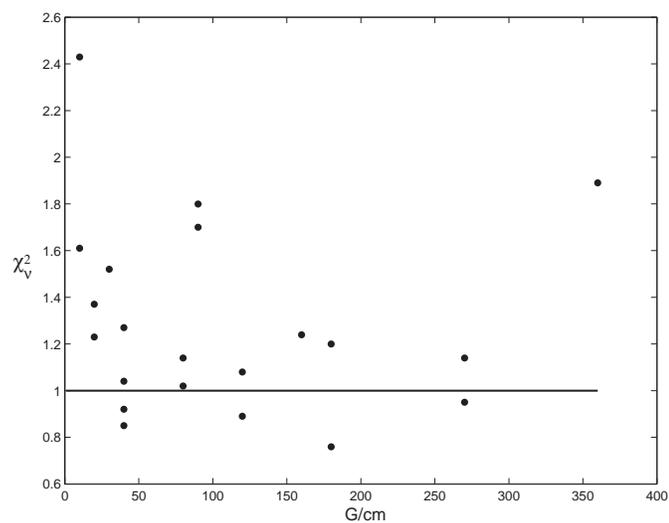


Figure 5.18: The reduced χ^2 for the V_0 fits in which $\beta = 12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$.

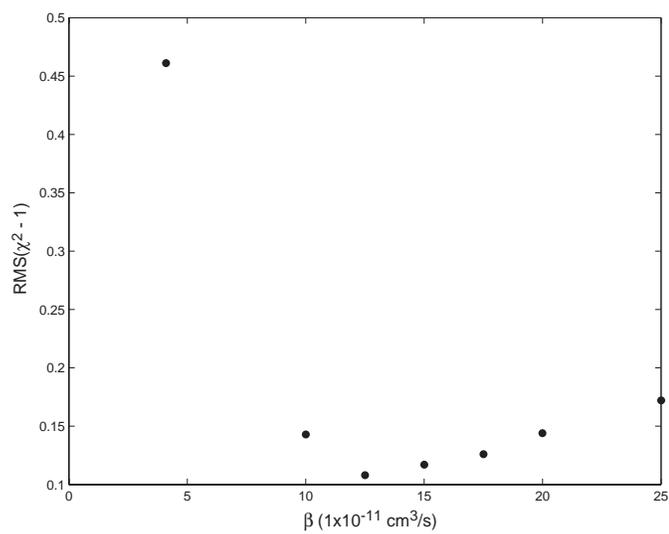


Figure 5.19: RMS of $\chi^2 - 1$ for various V_0 fits in which β is fixed. The optimal fixed β is $12.5 \times 10^{-11} \text{ cm}^3\text{s}^{-1}$.

5.2.2 U-trap losses

To load the U-trap, sub-doppler cooling and optical pumping is performed subsequent to the U-MOT loading and positioning. Once captured inside the U-trap, the atoms need to be held long enough to enable transfer into an Ioffe trap which is usually in the form of a Z-trap. This section examines the loss mechanisms that limit the lifetime of a cesium U-trap.

5.2.2.1 Background collisions

The vacuum chamber used for these experiments typically has a pressure between 1×10^{-9} to 1×10^{-8} Torr. We are limited to this pressure due to the quantity and sub-optimal vacuum quality of the atom chip materials in the chamber. The use of a cesium oven without a quick shut-off valve precludes the extinction of cesium vapor pressure during the post-MOT phase of the experiment. We believe the background gases—dominated by the lighter elements such as helium—are in sufficient concentrations to put the vacuum quality of the chamber in the $P_{He}/P_{Cs} \approx 0.3$ regime (see Figure 5.2). The lifetime due to background gas collisions is on the order of 0.5 s for a pressure of 1×10^{-8} Torr and increases to around 5 s for pressures lower than 1×10^{-9} Torr. Techniques to increase the trap lifetime while maintaining large trapped atom numbers are discussed in Chapter 10. In these experiments—performed in collaboration with the group in Munich—a desorption-based dispenser of rubidium is used instead of an oven as the alkali source and a UV light for chamber wall alkali desorption is used to transiently increase the loading rate of the MOT. The dispenser and UV light can be shut-off quickly during the experiment, allowing an initial capture of a large number of atoms while providing a low background pressure later in the experiment to obtain long magnetic trap lifetimes.

5.2.2.2 Ground-state collisions and Majorana spin-flips

From Equations 5.16 and 5.17, we calculate the ground-state collisional rates in the U- and Z-trap for various trap volumes and initial atom numbers. Table 5.1 lists these lifetimes for $N = 1 \times 10^6$ and 3×10^4 atoms, for atom temperatures equal to $T = 120$ and $1 \mu\text{K}$, and for trap dimensions consistent with measurements from U-trap images. The $T = 1 \mu\text{K}$ [$T = 120 \mu\text{K}$] entries are for the case in which the U-trap is loaded with [without] sub-

		T = 120 μ K		T = 1 μ K		
		$\tau_{4,4}, \tau_{3,-3}$		$\tau_{4,4}, \tau_{3,-3}$		
r	l	τ_{Major}	$n = 1 \times 10^6$	$n = 3 \times 10^4$	$n = 1 \times 10^6$	$n = 3 \times 10^4$
22 μ m	102 μ m	0.91 s	0.41 s, 5.7 s	14 s, 191 s	0.02 s, 0.14 s	0.67 s, 4.57 s
44 μ m	204 μ m	3.6 s	3.3 s, 46 s	110 s, 1,531 s	0.16 s, 0.14 s	5.4 s, 37 s
60 μ m	280 μ m	6.9 s	8.6 s, 121 s	287 s, 4,017 s	0.42 s, 2.88 s	14.1 s, 96 s

Table 5.1: Lifetimes for Majorana (τ_{Major}) and ground-state losses ($\tau_{4,4}$ and $\tau_{3,-3}$) in a U-trap. The ground-state loss rates are applicable to the Z-trap as well.

Doppler cooled atoms. (Actually, sub-Doppler cooling only provides ~ 3 μ K cesium atoms at best).

One can see that colder atoms exacerbate the collisional rate due to their higher density in the trap. Of course, lower temperature provides a more efficient loading into subsequent small phase-space traps. It is also apparent that trapping in the F=3, $m_f=-3$ state is preferable to the F=4, $m_f=4$ state due to the absence of hyperfine changing collisions. We performed an experiment in which the atoms were optically pumped into the F=4, $m_f=4$ state (as opposed to the F=3, $m_f=-3$ state) and noticed a marked decrease in the trap lifetime. It seems possible to have a ground-state collision-limited Ioffe trap of more than a few seconds with $\sim 1 \times 10^5$ atoms at sub-Doppler temperatures.

Our U-trap has high magnetic field gradients of more than 300 G/cm which may increase the Majorana spin-flip rate beyond those of background or ground-state collisions. The non-zero field minimum of the Z-trap prevents these spin-flips, but this trap is typically loaded with an initial U-trap whose Majorana spin-flip rate must be considered to maximize atom loading number and transfer efficiency. A U-trap is highly elongated and this anisotropy slightly enhances the loss rate with respect to a spherical trap of equal volume. We use the method of Section 5.1.5 to estimate the spin-flip loss rate in a U-trap, now accounting for the trap's anisotropy.

We begin calculating the loss rate by first defining two lengths associated with the nonadiabatic crossover region:

$$r_r = \sqrt{\frac{v\hbar}{g\mu_0\nabla_r\mathbf{B}}}, \quad \text{and} \quad r_x = \sqrt{\frac{v\hbar}{g\mu_0\nabla_x\mathbf{B}}}. \quad (5.38)$$

The coordinate notation follows that of Figure 5.4 with \hat{y} the strongly-confined trap axis and \hat{x} the weakly-confined trap axis. Images of the trap confirm that the confined cloud of

atoms is approximately cylindrically symmetric about \hat{x} , and we make the approximation that the spatial extent and gradients of the trap in \hat{y} can be used for those in \hat{r} , where $r^2 = y^2 + z^2$. The trap gradients, $\nabla_r \mathbf{B} = \nabla_y \mathbf{B}$ and $\nabla_x \mathbf{B}$, are listed in Appendix A.

The virial theorem enables us to write

$$m(v_r^2 + v_x^2) = g\mu_0(r\nabla_r \mathbf{B} + l\nabla_x \mathbf{B}), \quad (5.39)$$

where r and l are the radii of the trapped cloud along the \hat{r} and \hat{x} , respectively. The decay rate is

$$1/\tau_{Maj} = \frac{1}{\pi r^2 2l} \cdot \sqrt{v_r^2 + v_x^2} \cdot 2\pi r_r r_x, \quad (5.40)$$

$$1/\tau_{Maj} = \frac{\hbar}{rm} \left(\frac{1}{l} \sqrt{\frac{\nabla_r \mathbf{B}}{\nabla_x \mathbf{B}}} + \frac{1}{r} \sqrt{\frac{\nabla_x \mathbf{B}}{\nabla_r \mathbf{B}}} \right). \quad (5.41)$$

Comparing the lifetimes of a spherical trap and a U-trap with the same volumes, we find

$$\frac{\tau_{sph}}{\tau_{Utrap}} = \left(\frac{\sqrt{rl}}{2\sqrt{3}} \right)^{2/3} \left(\frac{1}{l} \sqrt{\frac{\nabla_r \mathbf{B}}{\nabla_x \mathbf{B}}} + \frac{1}{r} \sqrt{\frac{\nabla_x \mathbf{B}}{\nabla_r \mathbf{B}}} \right). \quad (5.42)$$

For a U-trap 0.33 mm above the substrate,

$$\sqrt{\frac{\nabla_r \mathbf{B}}{\nabla_x \mathbf{B}}} \approx 3.9, \quad (5.43)$$

and

$$\frac{\tau_{sph}}{\tau_{Utrap}} = 1.33. \quad (5.44)$$

Majorana spin flip losses are slightly worse in a U-trap than in a spherical trap of equal volume.

From trap images we find that typically $60 \mu\text{m} \geq r \geq 22 \mu\text{m}$ and $280 \mu\text{m} \geq l \geq 102 \mu\text{m}$. With these upper and lower bounds, we conclude that $6.9 \text{ s} \geq \tau_{Maj} \geq 0.91 \text{ s}$. The U-trap's Majorana spin-flip loss rate is of the same order of magnitude as the loss rate from ground-state collisions, and the U-trap should be converted to a Z-trap as soon as possible after loading.

5.2.3 Elastic collisions

Evaporative cooling is an essential ingredient in achieving BEC. In its simplest form, the trap height is lowered to allow the highest energy atoms to escape, and the remaining atoms rethermalize to a lower temperature through elastic collisions. (More commonly, an RF “knife” is employed [46].) In the case of a U-trap, the atoms are compressed by increasing the trap gradients. This leads to heating and the trap height must be simultaneously increased to prevent the atoms from “spilling over” the trap edge [68].

For a two-dimensional linear trap the temperature scales as

$$T = \alpha^{4/7} \quad (5.45)$$

when the potential is compressed by α [68, 131]. In a U-trap, the trap depth increases proportionally to the bias field, and the trap gradient increases with the square of the bias field: Increasing the bias field from B_{bias}^i to $B_{bias}^f = \beta B_{bias}^i$ causes the gradients at the center of the U-trap to increase by

$$\frac{\nabla B_{bias}^f}{\nabla B_{bias}^i} = \beta^2, \text{ and} \\ \alpha \sim \beta^2. \quad (5.46)$$

The final temperature is $T \propto \beta^{8/7}$, and this increases only slightly faster than the trap depth, which is proportional to $\mu_b \beta B_i$. For the $F=4, m_f=4$ and the $F=3, m_f=-3$ states, the bias field only needs to be 1.8 G and 2.4 G, respectively, to have a trap depth larger than the typical initial temperature of 120 μK . Moreover, increasing the trap gradients by a factor of 10 only increases the temperature by a factor of 1.4. Sub-doppler cooling the atoms to much less than 120 μK , or starting with a trap bias field larger than 3 G ensures no spillover atom loss.

5.2.4 Concluding remarks

As the above analyses suggest—and the experiments discussed in other chapters confirm—it is possible to load cesium atoms into magnetic microtraps for a time long enough to perform cavity QED experiments in spite of this atom’s unusually high collisional loss rates. However, we do need to be careful in how we load these traps. Using a macroscopic

wire located below the substrate to capture the atoms at the mirror MOT position with strong gradients allows us to be able to immediately shut-off the trapping lasers, thereby minimizing excited-state collisions. After a cooling and optical pumping stage, quickly exchanging a U-trap with a Z-trap suppresses Majorana spin flips. If the atoms are optically pumped to the $F=3$, $m_f = -3$ level with a temperature in the tens of μK , then we can trap several hundred thousand cesium atoms for more than a second.

Magnetic microtraps provide a new experimental setting to measure collisional losses as demonstrated in our measurement of the excited-state collisional loss rate of cesium. Further studies of collisional properties may be performed with the use of atoms confined in carefully engineered magnetic microtraps. For instance, tight confinement of two atoms in a harmonic microtrap potential could be employed to determine their scattering length. A spectroscopic measurement of the trap's level shifts would reveal the strength of their collisional interaction [132].

Chapter 6

The Atom-Cavity Chip: Combining microwire traps with photonic bandgap cavities and microdisks

This chapter is loosely based on paper [44]. The analyses have been updated and extended. The use of microdisk cavities for our experiments is discussed and a section comparing various cavity systems is included.

6.1 Introduction

The development of techniques necessary to manipulate single atoms and photons and to control their interactions is an important addition to the toolboxes of nanotechnology and quantum control. An important advance would be the development of a compact and integrable device to serve as a single atom detector [54, 23] or a repeater and processing node in a quantum network [1]. The system comprised of a strongly interacting atom and photon—cavity quantum electrodynamics (QED) [133, 134, 33]—provides the basis for realizing such devices. Single atom detectors could play as important a role in the burgeoning field of atom optics [135] as single photon detectors do in conventional optics. The advent of Bose-Einstein condensates (BECs) of neutral atoms and the production of degenerate fermionic condensates [111] further highlights the importance of developing single atom read-out devices.

To achieve these goals in cavity QED, a neutral atom must be inside the mode of a high-quality cavity with small mode volume: the atom-cavity system must be in the strong

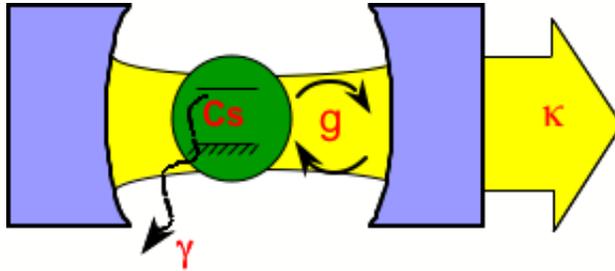


Figure 6.1: Model system for neutral atom cavity QED. A cesium atom—often approximated by a two-level system—is localized inside a high finesse resonator shown here as a Fabry-Perot cavity.

coupling regime. Strong coupling requires the atom-cavity coupling, g_0 , to be much larger than both the atomic dipole decay rate, $\gamma_{\perp} = \gamma/2$, and the decay rate of the cavity field, κ (see Figure 6.1). Specifically, the saturation photon number, $m_0 = \gamma_{\perp}^2/2g_0^2$, and the critical atom number, $N_0 = 2\gamma_{\perp}\kappa/g_0^2$, must both be much less than unity. Achieving strong coupling in the lab becomes the challenge to design, fabricate, and ultimately load atoms into cavities that simultaneously minimize mode volume ($g_0 \propto 1/\sqrt{V_m}$) and maximize the Q ($\kappa \equiv \pi c/(\lambda Q)$).

State-of-the-art neutral atom cavity QED experiments have achieved strong-coupling parameters as small as $[m_0, N_0] \approx [10^{-4}, 10^{-3}]$ by either dropping [37], or vertically tossing [38] a cold neutral atom between the mirrors of a high-finesse, low-mode volume Fabry-Perot cavity. Intracavity atom trapping for durations up to 3 s has been demonstrated by coupling a secondary optical beam into the Fabry-Perot cavity to form a Far Off Resonance Trap (FORT) [40]. Recently, cavity cooling effects and a side-coupled optical dipole trap have increased this trapping time to ~ 15 s [136, 137].

The intent of this chapter is to introduce a cavity QED system based on magnetostatic delivery of atoms to a photonic bandgap (PBG) or microdisk cavity, and to discuss the ability of this system to detect single atoms. Figures 6.2 and 6.3 show scanning electron microscope (SEM) images of a PBG and microdisk cavity, respectively. This experimental system—magnetostatic confinement of atoms inside the field mode of these microcavities—raises the possibility of achieving an experimentally robust, integrated, and scalable cavity QED apparatus. Mastering the integration of a single atom and photons—quintessentially quantum components—presents an entirely new prospect for technology: quantum compu-

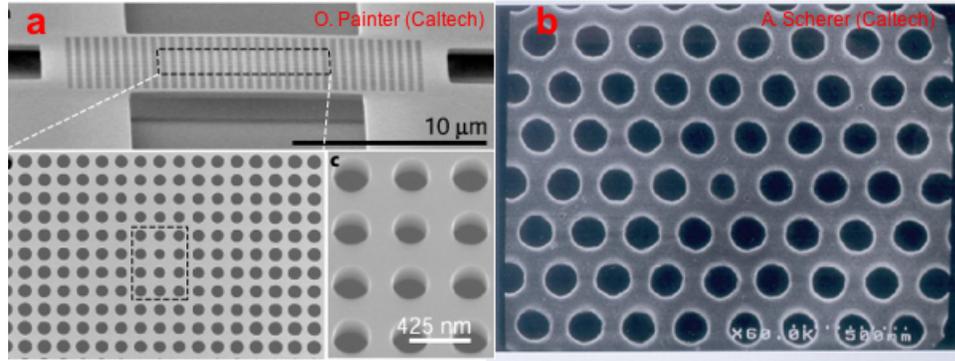


Figure 6.2: Images courtesy of O. Painter and A. Scherer’s groups at Caltech.

tation and communication. Cavity QED provides a rich experimental setting for quantum information processing (QIP), both in the implementation of quantum logic gates and in the development of quantum networks [34, 1]. While not necessary for single atom detection, confining the atom in the Lamb-Dicke regime inside the cavity for long periods of time is an important step towards accomplishing QIP using cavity QED. Moreover, quantum feedback experiments with this system will provide an excellent setting to explore the real-time actuation and measurement of an open quantum system.

Patterns of micron-sized wires can create magnetic field gradients and curvatures sufficiently large to accurately guide and trap atoms above the surface of the substrate [3]. These magnetic microtrap devices—commonly known as atom chips [51, 50]—can be fabricated using standard photolithography techniques [49, 8] and have been successfully used not only to trap and waveguide neutral atoms, but also to create and manipulate Bose-Einstein condensates [60, 61]. The proximity of the atoms to the chip’s surface naturally facilitates the integration of magnetically trapped atoms with on-chip cavities such as microdisks or photonic crystals. Chapters 3, 2, and 4 discuss atom chip design, fabrication, and operation in more detail, but we include a short introduction in the following.

Atom chips exploit the interaction potential, $V = -\vec{\mu} \cdot \vec{B}$, between an atom’s magnetic moment, $\vec{\mu}$, and a wire’s magnetic field, \vec{B} , to trap or guide weak-field seeking states of a neutral atom. The simplest example of a magnetic microtrap involves the combination of the field from a U-shaped wire with a homogenous bias field, B_{bias} [2]. The bias field, parallel to the wire substrate and perpendicular to the base of the U-wire, serves to cancel the curling field of the wire to form a two-dimensional quadrupole trap for the weak-field

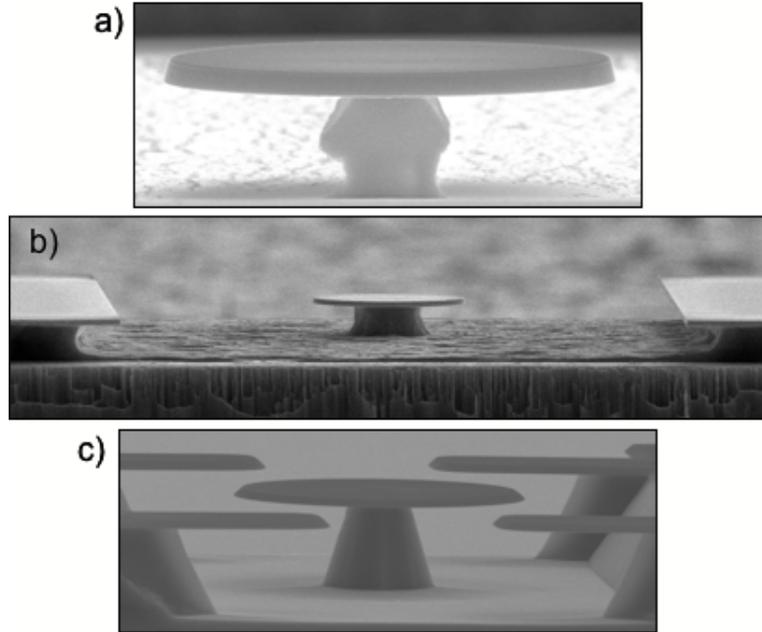


Figure 6.3: a) and b) SEM of AlGaAs microdisks of $\sim 9 \mu\text{m}$ diameter. c) SEM of SiN_x microdisk of diameter $8.4 \mu\text{m}$. Image courtesy of O. Painter’s group at Caltech. The diving board-like structures are supports for holding the fiber taper in place.

seeking atoms. The atoms are confined in the third dimension by the fields from the side wires of the U-trap, forming a cigar-shaped trap above the wire surface. The position of the trap minimum above the wire surface, r , and the gradient of the trap are completely determined by the magnitude of B_{bias} and the current, I , in the U-wire,

$$r = \frac{\mu_0 I}{2\pi B_{bias}}, \quad \nabla B = \frac{2\pi B_{bias}^2}{\mu_0 I}. \quad (6.1)$$

For example, with a wire current of 1 A and a bias field of 10 G, the atoms are trapped $200 \mu\text{m}$ above the surface in a field gradient—perpendicular to the base of the U-wire—of 500 G/cm.

An atom is trapped in the Lamb-Dicke regime when its recoil energy is less than the trap’s vibrational level spacing, $\eta_{LD} = (E_{recoil}/E_{vib})^{1/2} < 1$, and this regime has been achieved inside a Fabry-Perot resonator by using a FORT [40]. Ioffe traps—which are non-susceptible to trap losses due to Majorana spin flips—may be formed either by a similar Z-trap [2] or by using wires forming patterns of nested arcs [3]. Although this latter Ioffe trap is more complicated, it does allow the possibility of magnetically trapping atoms three-

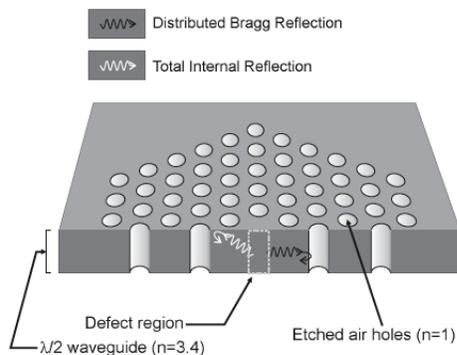


Figure 6.4: Schematic of a photonic bandgap cavity. An array of holes of periodicity commensurate with λ induces distributed Bragg reflection of in-plane light. Light is confined perpendicular to the plane in an optically thin transparent membrane of high index with respect to vacuum. A defect hole—typically of smaller diameter—supports a localized cavity mode. Figure from Reference [1].

dimensionally in the Lamb-Dicke regime inside a photonic bandgap cavity coplanar with the wires (see Figure 6.5 (b)) [1].

Magnetic microwire traps and PBG cavities form a scalable architecture for a quantum information processor using neutral atoms [1]. A chip would contain many PBG cavities, each with its own independently controlled magnetic microtrap. The atoms can be introduced to each cavity at will, and on-chip photonic waveguides network the light between the cavities, built-in detectors, and laser sources. Figure 6.4 sketches the operating principle of a PBG cavity, and Figure 6.5 depicts the scheme for integrating a microwire Ioffe trap with this resonator system.

Simple waveguides for the atoms can be formed from the Z-trap by extending the base of the Z-wire, allowing the atoms to ballistically expand along the field minimum above the elongated wire. Wires perpendicularly intersecting the elongated base of the Z-wire can serve as independently controllable gates to aid in the local confinement of the atoms in the waveguide. Chapter 2 discusses this waveguiding technique in detail. Beam splitters and conveyor belts have been demonstrated using similar techniques [51, 50].

6.2 Cavity QED with microcavities

This section discusses the use of PBG and microdisk cavities for neutral atom cavity QED with atom chips, and compares their properties to other resonator systems used for cavity

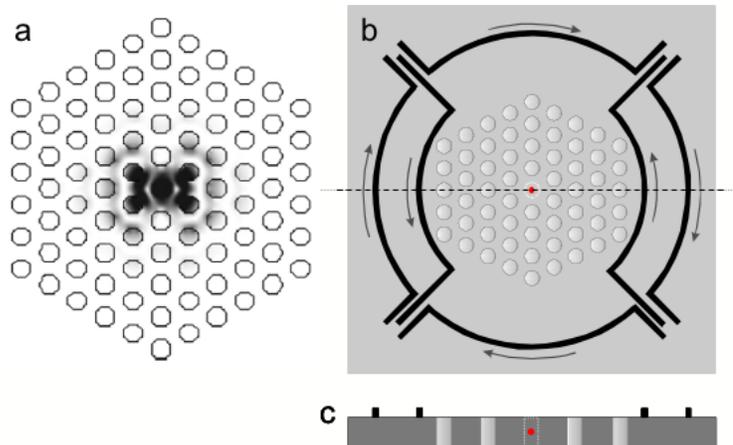


Figure 6.5: a) Simulation of the cavity mode in a hexagonal lattice photonic crystal. The cavity mode is centered and well-localized about the defect hole. b) Sketch of a PBG cavity with an integrated Ioffe microwire trap. The red dot shows the location of the trapped atom within the defect hole. Given current microwire technology, the wire pattern radius and wire thickness would need to be scaled-up by a factor of 10. The PBG hole size is ~ 100 nm, and the wire diameter and width would be $\sim 10 \mu\text{m}$ and $\sim 1 \mu\text{m}$, respectively. c) Cross-section of PBG cavity and Ioffe trap through a line intersecting the defect hole. The trapped atom—depicted as a red dot in the center of panels (b) and (c)—would be located inside the cavity field maximum at the center of the defect hole. Figures adapted from References [24] and [1].

QED. Section 6.3 details our scheme for integrating these microcavities with atom chip-based magnetic waveguides. The following Chapter 7 discusses a similar experimental system that utilizes a fiber-gap Fabry-Perot cavity.

6.2.1 Photonic bandgap cavities

Two-dimensional photonic bandgap (PBG) cavities—perforated semiconductor structures that confine light through the dual action of distributed Bragg reflection and internal reflection—are in many respects ideal for cavity QED [24]. Their small mode volume and modest quality factors open the near-term possibility of achieving extremely small strong coupling parameters: $[m_0, N_0] = [10^{-8}, 10^{-4}]$. As inherently stable, flat, monolithic structures, PBG cavities do not need the support structure for active stabilization that Fabry-Perot cavities require. Moreover, their compactness and compatibility with fiber optics-based input and output couplers allow one to envision an array of PBG cavities, atom microtraps, input/output couplers, and other processing devices all on the same integrated

chip for the formation of a node of a quantum network. Reference [138] presents simulations of the fiber taper-to-cavity photonic crystal coupling and references [139, 140, 141] describe experimental demonstrations. This work is done in close collaboration with Oskar Painter’s group at Caltech, which has developed the PBG cavity design and fabrication as well as the system for fiber taper coupling.

We plan to use PBG cavities of the graded defect design discussed in reference [142] and demonstrated in references [143, 144]. These consist of a rectangular lattice of air holes in an optically thin, high refractive index slab waveguide. The holes gradually decrease in diameter towards the center, forming the cavity. Experimental measurements of such cavities fabricated in silicon membranes (see Figure 6.6 (a)) and operating at $\lambda \sim 1.6 \mu\text{m}$ possess Q’s as high as 40,000 with modal volumes of $V_m \sim 0.9$ cubic wavelengths $(\lambda/n)^3$ [144]. Scaling the system down to an operating wavelength suitable for a cesium transition (852 nm), the central hole diameter is ~ 100 nm and the membrane thickness is ~ 170 nm. For the Q and V_m values mentioned above and taking $\lambda = 852$ nm and $n = 3.4$, the atom-cavity coupling at the center of the cavity can be as high as $g_0 = 2\pi \cdot 16$ GHz while the decoherence rates are $[\kappa, \gamma_{\perp}]/2\pi = [4.4 \text{ GHz}, 2.6 \text{ MHz}]$. This gives strong coupling parameters of $[m_0, N_0] = [1.3 \times 10^{-8}, 8.8 \times 10^{-5}]$, which are much smaller than those achieved in recent experiments using Fabry-Perot cavities, $[m_0, N_0] = [2.8 \times 10^{-4}, 6.1 \times 10^{-3}]$ [37].

For experiments with single neutral atoms, the PBG cavities need to be fabricated out of membranes transparent in the near infrared (≤ 900 nm). This precludes the use of silicon or InP that have been used previously. The wavelength of the D2 transition of cesium—our atom of choice for trapping experiments—is 852 nm, and we are investigating the suitability of both AlGaAs and SiN_x as PBG substrates at this wavelength. The Al_xGaAs_{1-x}, with $x \approx 0.3$, is more desirable due to its higher index of refraction ($n = 3.4$) compared to SiN_x’s $2 \leq n \leq 2.3$. AlGaAs microdisks tested at the 1.4 μm band have yielded Q’s as high as 3.6×10^5 [145], but unfortunately, high Q operation at 852 nm has proven more difficult. The material we have used so far has been too absorptive at 852 nm, greatly suppressing what should have been an equally high cavity Q. We have not been able to fine-tune the Al percentage and quality of the AlGaAs substrates to eliminate this absorption problem. Part of the difficulty is due to the fact that this material is not grown by us, several months to procure each batch, and is quite expensive. Obtaining and verifying a target Al percentage is therefore difficult.

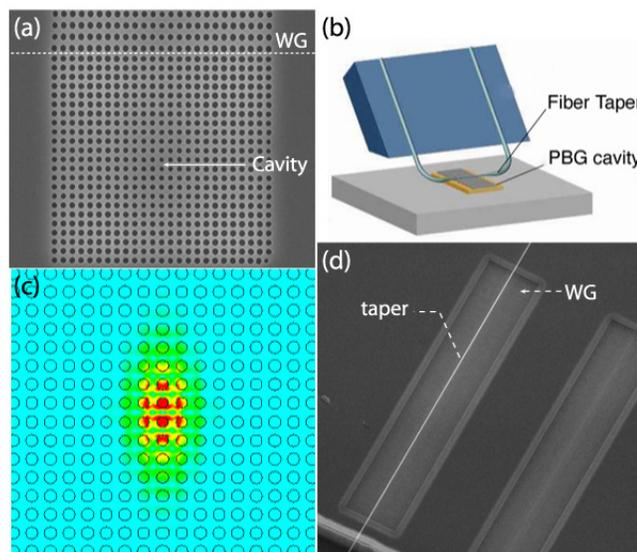


Figure 6.6: (a) SEM image of a photonic bandgap cavity and waveguide (WG) fabricated in silicon. (b) Schematic of the fiber taper coupler. (c) Finite-difference time-domain calculated electric field amplitude of the cavity mode taken in the center of the membrane. (d) SEM image of an optical fiber taper aligned above a photonic crystal waveguide. Figure from Reference [44].

SiN_x , though possessing a smaller n , is perhaps a much more promising substrate material for the PBG cavities. Substrates of usable quality can be grown in-house, and higher quality samples are inexpensive to purchase and can be obtained with short lead-times. Moreover, PBG and microdisk fabrication in SiN_x is simpler because SiN_x can be etched directly with a mask formed from e-beam resist whereas AlGaAs needs a harder mask. This requires an additional step to grow and etch a secondary, sacrificial mask between the e-beam resist and the AlGaAs. Experiments in Oskar Painter's group have demonstrated Q's as high as 3×10^6 at 852 nm in SiN_x microdisks, which leads us to believe that high Q's might be obtainable in SiN_x PBG cavities as well.

The cavity is coupled to a photonic crystal waveguide, which in turn is evanescently coupled to an optical fiber taper. By positioning the fiber taper—whose minimum diameter is on the order of a micron—along the axis and in the near field of the photonic crystal waveguide, highly efficient (greater than 97%) fiber coupling into and out of the photonic crystal waveguide can be achieved [139] (see Figures 6.6 [b] and [d]). Light coupled into the photonic crystal waveguide is reflected by the PBG cavity and recollects in the backward propagating fiber taper mode. Coupling and recollection efficiencies of 44% have been

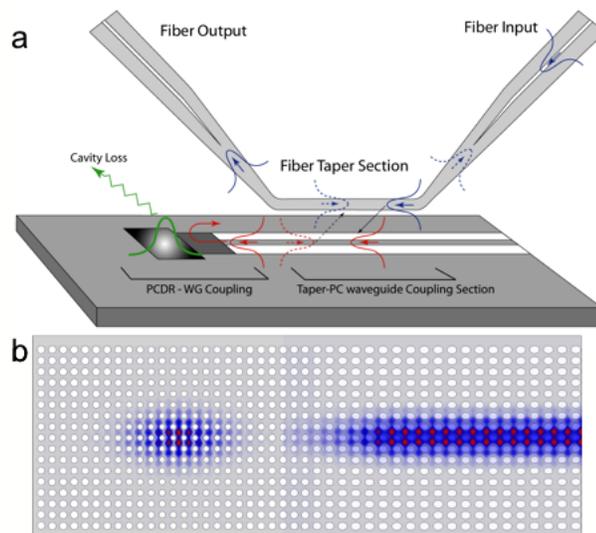


Figure 6.7: (a) Schematic of the fiber taper coupling system. (b) Simulation of the photonic crystal waveguide to cavity interface. Note that the field maximum in this simulation is in the material—not the holes. A mode with the maximum in the hole will be used for our neutral atom experiments. Figure courtesy of O. Painter’s group at Caltech.

achieved (see Figure 6.7) [141].

Figure 6.6 (a) shows the boundary between the waveguide and the cavity: the top four rows of holes are the end of the waveguide, which is formed in a similar fashion to the cavity, except that the holes are graded in only the lateral dimension. This boundary is also depicted in Figure 6.7 (b). This design maximizes the mode matching between the waveguide and the cavity modes [138]. The waveguide may be bent to allow atom insertion access into the cavity unencumbered by the fiber.

6.2.2 Microdisks

The microdisk cavity—akin to microsphere resonators, but two-dimensional—support whispering gallery modes around its rim [146, 145]. The cavity mode has an evanescent tail that extends a fraction of a wavelength into the vacuum. Single atoms positioned outside the disk’s rim can couple to this evanescent tail. Microdisks are easier to fabricate than PBG cavities and possess much higher Q ’s at the expense of a factor of ~ 10 increase in mode volume. We have made microdisks out of materials suitable for operation at 852 nm (AlGaAs and SiN_x), and used them for fabrication and quality factor diagnostics. Figure 6.3 is an SEM image of typical microdisks fabricated by O. Painter’s group [145]. Recently, $8.4 \mu\text{m}$

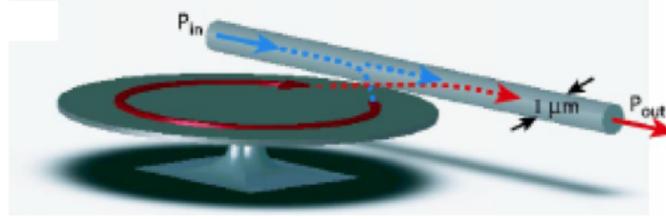


Figure 6.8: Schematic of the fiber taper-to-microdisk coupling scheme. Figure drawn approximately to scale. Courtesy of O. Painter’s group.

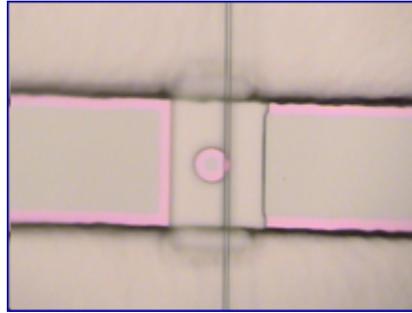


Figure 6.9: Fiber taper coupled to a $9 \mu\text{m}$ diameter AlGaAs microdisk cavity. The pink regions reveal the undercut material, and the black shadows demarcate the plateau regions.

diameter microdisks in SiN_x have been fabricated. At $\lambda = 852 \text{ nm}$, these microdisks have a $Q = 3 \times 10^6$ and a mode volume of $V_m = 12(\lambda/n)^3$, with $n = 2.0$. Since the maximum of the cavity field is in the dielectric, the resulting g_0 is reduced by a factor of ~ 3 at the position of an atom assumed to be at the edge of the disk. Nevertheless, strong coupling can be achieved with $[g_0, \kappa, m_0, N_0] = [1.5 \text{ GHz}, 180 \text{ MHz}, 1.5 \times 10^{-6}, 1.3 \times 10^{-4}]$.

As these microdisks are also suitable for our cavity QED experiments with magnetically guided cesium atoms, we will be using them in our first atom-cavity chip experiments. We use a fiber taper to couple light into and out of the microdisk resonator’s whispering gallery mode. Figure 6.8 depicts the coupling scheme. Cold atoms will be magnetically guided to the edge of the disk opposite the fiber taper. An image of a fiber taper coupled to a $9 \mu\text{m}$ diameter microdisk is shown in Figure 6.9. The microdisks are fabricated on a plateau to prevent the fiber taper from touching—and coupling to—the bulk material. Coupling efficiencies of $\geq 97\%$ have been demonstrated in both microdisks and microtoroids with this technique.

6.2.3 Comparison to other cavity systems

The table in this section compiles cavity QED parameters for various cavity QED resonator systems.¹ The first five entries are resonator systems employed by other researchers, and the last three are resonators that are featured in this thesis. We have included microwave and circuit QED to provide an interesting comparison. Detection of vacuum Rabi splitting in quantum dot-based systems has recently been reported, and we refer the reader to References [147, 148, 149] for more information.

For all but the microwave, circuit QED, and fiber-gap Fabry-Perot systems, the g_0 and Q is quoted assuming the operating wavelength of $\lambda = 852$ nm and cesium's D2 transition dipole moment ($\mu = 2.686 \times 10^{-29}$ C m). In this table, g_0 is defined as half the single-photon Rabi frequency and accounts for the fact that the atom might not be located in the maximum of the cavity's electric field:

$$\hbar g_0 = \vec{\mu} \cdot \vec{E} = \xi \mu \sqrt{\frac{\hbar \omega}{2e_0 V_m}} = \xi \gamma_{\perp} \sqrt{\frac{3c\lambda^2}{4\pi\gamma_{\perp} V_m}}, \quad (6.2)$$

where $\xi \leq 1$ accounts for the atom's position. The mode volume is

$$V_m \equiv \frac{\int \int \int \epsilon(\vec{r}) |\vec{E}(\vec{r})|^2 d^3 \vec{r}}{\max[\epsilon(\vec{r}) |\vec{E}(\vec{r})|^2]}, \quad (6.3)$$

where $\epsilon(\vec{r})$ is the cavity's spatially-dependant dielectric function. The fiber-gap Fabry-Perot entries are based-on the D2 line of ^{87}Rb ($\lambda = 780$ nm), which is the atom used in our experiment reported in Chapter 7. For this transition in ^{87}Rb , the dipole moment is only 5.7% smaller than cesium's, and the g_0 's listed in Table 6.1 have been adjusted accordingly. The microwave and circuit QED systems use Rydberg atoms [150] and Cooper pair boxes [151, 152], respectively. The atomic decay rate for the microwave cavity system is $\gamma/2\pi = 5.3$ Hz. Experimentally, the Cooper pair box decay rate is $\gamma/2\pi = 0.7$ MHz [152], but is projected to be less than $\gamma/2\pi = 0.08$ MHz [151].

The table lists the Q 's for each system, though for the two Fabry-Perot systems—one using macroscopic mirrors and the other with glued films on the tips of fiber optics—cavity

¹Cavity QED is an exciting, and therefore, rapidly advancing field. In the formation of Table 6.1, the author has attempted to compile a fair and up-to-date snapshot of this field. Please excuse any omissions and inaccuracies—none were intentional.

finesse, \mathcal{F} , is quoted as this is a more common measure. For conversions:

$$\kappa = \frac{\pi c}{\lambda Q}, \quad \mathcal{F} = \frac{\pi c}{2L\kappa}, \quad \mathcal{F} = \frac{\lambda Q}{2L}. \quad (6.4)$$

The last two columns of the table list the coupling to dissipation ratio, $g_0/\max[\kappa, \gamma_\perp]$, and the “rate of optical information per atom,” $I = g_0^2/\kappa$ [134, 153, 154]. The former is a measure of the number of Rabi flops the system can undergo, and the latter is an indicator of the amount of information that can be gleaned about some aspect of the atom-cavity dynamics. Note that the coupling to dissipation ratio for the microwave case is limited not by κ or γ_\perp , but rather by the limited transit time of the Rydberg atom through the superconducting microwave cavity ($t_{tr} \sim 100 \mu\text{s}$). In all other cases, atoms can be trapped inside the resonator mode longer than $1/\max[\kappa, \gamma_\perp]$

An attempt has been made to list the current experimental state-of-the-art along with projected limits of each system. This is done explicitly to remove any confusion regarding the fundamental capabilities of each cavity design. A similar table found in Spillane *et al.* [57] denoted the PBG cavity experimental result [44, 144] as a theoretical upper bound for this resonator system’s potential for cavity QED experiments. Comparisons between the PBG and the microtoroid are best made on an equal basis as is attempted in Table 6.1. Contrary to the conclusions in Spillane *et al.*, current theoretical projections indicate that PBG cavities can attain smaller critical photon and atom numbers, a higher coherent-to-dissipation ratio, and a greater information rate. With respect to experiments in the near term, PBG cavities in silicon have slightly better atom-cavity coupling characteristics. Of course, this assumes that the fabrication of PBG cavities in 852 nm transparent materials will not degrade their performance, and this is under current investigation.

With regard to neutral atom cavity QED, PBG cavities have the advantage over other microcavities (i.e. microdisks [145, 56], microtoroids [57], microspheres [20, 155, 156, 55]) in that the field maximum of the cavity mode can be located in the vacuum rather than inside the dielectric material ($\xi = 1$). Moreover, they are much more two-dimensional which aids in large-scale integration. Cavity QED experiments using these microcavities will all have to contend with the Casimir-Polder potential since the atoms will be placed within a fraction of an optical wavelength from the dielectric. In comparison to these microcavities and Fabry-Perot resonators [156, 157], the PBG cavity’s mode volume is significantly smaller—less

than a cubic wavelength—which results in a much larger g_0 . Although the PBG cavity’s Q is not currently as large as the other optical resonators, experiments in silicon have produced PBG cavity Q’s of 4×10^4 [144] and 6×10^5 [158] which allow g_0 to be larger than κ . Theoretical predictions indicate that Q’s as high as 2.2×10^7 might be achievable by using a double-heterostructure of photonic crystals, leading to unmatched atom-cavity coupling parameters.

Note that the g_0 quoted in reference [144] and listed as *Painter Group* in Table 6.1 is lower than one would expect from the quoted mode volume, V_m . This is due to the fact that this particular cavity was designed for a field maximum at the edge of the central hole, not in the center, and this results in an $\xi = 0.4$. PBG cavities used for our cavity QED experiments will have the mode maximum in the central hole.

The experimental state-of-the art for SiN_x microdisks—which operate at 852 nm with $n = 2.0$ —achieve critical atom and photon numbers of $[m_0, N_0] = [1.6 \times 10^{-6}, 1.3 \times 10^{-4}]$ with $n = 2.0$ and a diameter of 8.4 μm . These cavities are projected to reach $[m_0, N_0] = [1.0 \times 10^{-7}, 2.6 \times 10^{-7}]$ in the future by refining the fabrication process, decreasing the diameter, and using $n = 2.3$ material. These values account for $\xi \sim 0.33$, corresponding to an atom positioned at the edge of the disk ². Since the normalized field maximum outside the disk scales roughly as $1/r$ and the mode volume as $r^{-4/3}$, $g_0 \propto r^{-5/3}$ is the scaling of the coupling as a function disk radius. Radiation losses do not pose as stringent a limit on the diameter—and hence mode-volume—of the microdisk as for the microtoroid. For instance in SiN_x , radiation-limited loss affects the Q only below 6 μm at 852 nm. Consequently, the mode volume of the microdisk can be roughly a factor of 10 smaller than that of the microtoroids. This largely compensates for the lower Q of the microdisk: the projected m_0 and coupling to dissipation ratio are comparable to those of the microtoroid, though the N_0 and I are a factor of 10 worse.

6.3 Experimental proposal

As a first generation experiment, we would like to bring a trapped cloud of cold neutral atoms—cesium in our case—into contact with a PBG or microdisk cavity, simultaneously demonstrating the integration of a microfabricated cavity with an atom chip and the strong

²The atom should not actually be placed on the edge of the disk, but this forms a reasonable reference point that is consistent with previous microsphere treatments.

Table 6.1: Comparison of cavity designs. See text for details and citations.

Cavity design	$[g_0, \kappa]/2\pi$ (MHz)	Q	m_0	N_0	$\frac{g_0}{\max[\kappa, \gamma_\perp]}$	$I \equiv g_0^2/\kappa$ (Mbits/s)
Fabry-Perot						
experimental	[110, 14.2]	$\mathcal{F}=4.8 \times 10^5$	2.8×10^{-4}	6.1×10^{-3}	7.8	5.4×10^3
projected	[770, 21.7]	$\mathcal{F}=7.9 \times 10^6$	5.7×10^{-6}	1.9×10^{-4}	36	1.7×10^5
Microsphere						
experimental	[24, 3.3]	5.3×10^7	5.3×10^{-3}	3.0×10^{-2}	7.2	1.1×10^3
projected:						
V minimized	[750, 7900]	1.4×10^4	6.1×10^{-6}	7.3×10^{-1}	0.01	4.5×10^1
Q maximized	[280, 0.048]	2.4×10^{10}	4.3×10^{-5}	3.1×10^{-6}	107	1.1×10^7
Microtoroid						
experimental	[86, 1.4]	1.2×10^8	4.6×10^{-4}	1.0×10^{-3}	33	3.3×10^4
projected:						
near term	[450, 1.8]	1.0×10^8	1.7×10^{-5}	4.5×10^{-5}	173	7.3×10^5
V minimized	[700, 18.8]	9.3×10^6	6.0×10^{-6}	2.0×10^{-4}	37	1.6×10^5
Q maximized	[430, 0.007]	2.5×10^{10}	2.0×10^{-5}	2.0×10^{-7}	165	1.6×10^8
Microwave						
experimental	[0.024, 0.00017]	3×10^8	3.3×10^{-6}	2.5×10^{-8}	3	2×10^7
Circuit QED						
experimental	[5.8, 0.8]	$\sim 10^4$	7.3×10^{-3}	3.3×10^{-2}	7.7	2.6×10^2
projected	[50, 0.6]	$\sim 10^4$	6×10^{-5}	1×10^{-6}	83	3×10^4
Cavity designs featured in this thesis						
Fiber-Gap						
Fabry-Perot						
experimental	[186, 2650]	$\mathcal{F}=1.1 \times 10^3$	1.0×10^{-4}	4.2×10^{-1}	0.07	7.8×10^1
projected	[423, 500]	$\mathcal{F}=1 \times 10^4$	1.9×10^{-5}	1.5×10^{-2}	0.8	2.2×10^3
Microdisk						
experimental:						
semiconductor	[8600, 1200]	1.5×10^5	4.5×10^{-8}	8.1×10^{-5}	7.3	3.9×10^5
SiN _x	[1500, 58.7]	3×10^6	1.5×10^{-6}	1.3×10^{-4}	26	2.5×10^5
projected:						
SiN _x	[2000, 1.8]	1×10^8	1.0×10^{-6}	2.6×10^{-6}	110	1.2×10^7
Photonic Bandgap						
experimental:						
Painter Group	[16000, 4400]	4.0×10^4	1.3×10^{-8}	8.8×10^{-5}	3.6	3.7×10^5
Noda Group	[33400, 293]	6.0×10^5	3.0×10^{-9}	1.4×10^{-6}	114	2.4×10^7
projected	[33400, 8]	2.2×10^7	3.0×10^{-9}	3.7×10^{-8}	4175	8.8×10^8

coupling of a neutral atom to a microcavity. Standard laser cooling and trapping techniques [46] are used to load cold atoms into the magnetic microtraps and waveguides. Typically, atoms are collected in a variant of the magneto-optical trap (MOT) that uses the atom chip surface as a mirror to form four of the six required laser cooling beams [2]. This mirror MOT and subsequent sub-doppler cooling allows the collection of 10^6 , ~ 10 μK atoms a few millimeters above the chip's surface. Conveniently, the quadrupole field from the U-trap is in the same orientation as the magnetic field required to form a mirror MOT. The atoms can be transferred to the U-trap by replacing the mirror MOT's quadrupole field with that of the U-trap while maintaining the cooling lasers in the same configuration. This creates a MOT using the microwire magnetic field—a U-MOT—which automatically spatially mode-matches the atoms to the U-trap once the cooling lasers are extinguished. Sub-doppler cooling and optical pumping stages are added as well. The mirror MOT employed in our lab is actually a macro U-MOT that can trap atoms directly from vapor. It uses a large copper U-shaped block carrying 20 to 30 A located underneath the atom chip [64] and eliminates the need for bulky coils external to the chamber, while ensuring a smooth, spatially mode-matched atom transfer to the smaller, magnetostatic U-traps on the atom chip surface. Chapters 2 and 3 contain more technical details regarding our atom chip trapping system.

Figure 6.10 is a rough schematic of the atom-cavity chip experiment. The chip is divided into two regions, one for laser trapping and cooling of the atoms in a U-MOT and U-traps, and the other for the microcavity and its fiber taper coupler. The two regions are connected by a microwire waveguide to transport the atoms from the laser cooling region to the cavity. These regions must be separated by at least 1 cm in order for the bulk of the cavity to not obstruct the 45° 1 cm^2 U-MOT beams. The thickness of both the microcavity's substrate and the fiber taper mount is small enough that it does not obstruct the horizontal U-MOT beam, allowing a simple, straight Z-trap waveguide as depicted in Figure 6.10 and 6.11. For experiments using taller, less compact cavities—such as the Fiber-Gap cavity discussed in Chapter 7—a waveguide with a 90° bend is required to convey the atoms out of the horizontal beam before being guided into the cavity mode. This can be accomplished either by using a two-wire guide [74] or by rotating the atoms in a P-trap—similar to a U-trap but with the base wire bent allowing a rotating bias field to change the orientation of the atoms [159]—before transferring the atoms into a Z-trap waveguide aligned perpendicular to

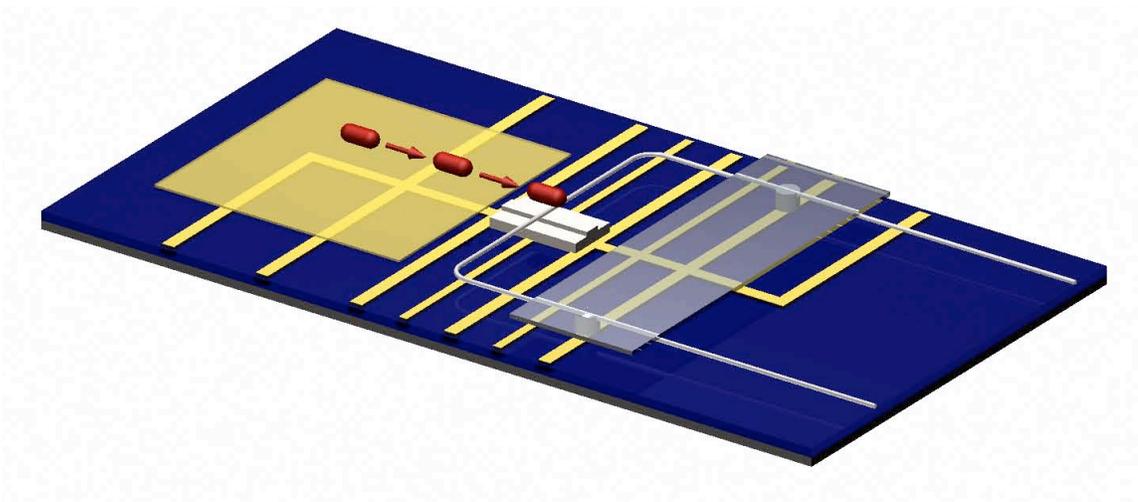


Figure 6.10: Schematic of the atom-cavity chip experiment. The microwire U-traps and atomic waveguides are shown as yellow wires. The light gold area centered about the initial U-trap is the mirror for the macro U-MOT, and its size roughly represents the footprint of the reflected trapping laser beams. The atoms are the red cylinders, pictured as they are transported towards the PBG or microdisk cavity which is shown as the white chip glued to the substrate's surface. The grey line is the fiber taper which is glued to the thin glass cover slip shown as the transparent white rectangle.

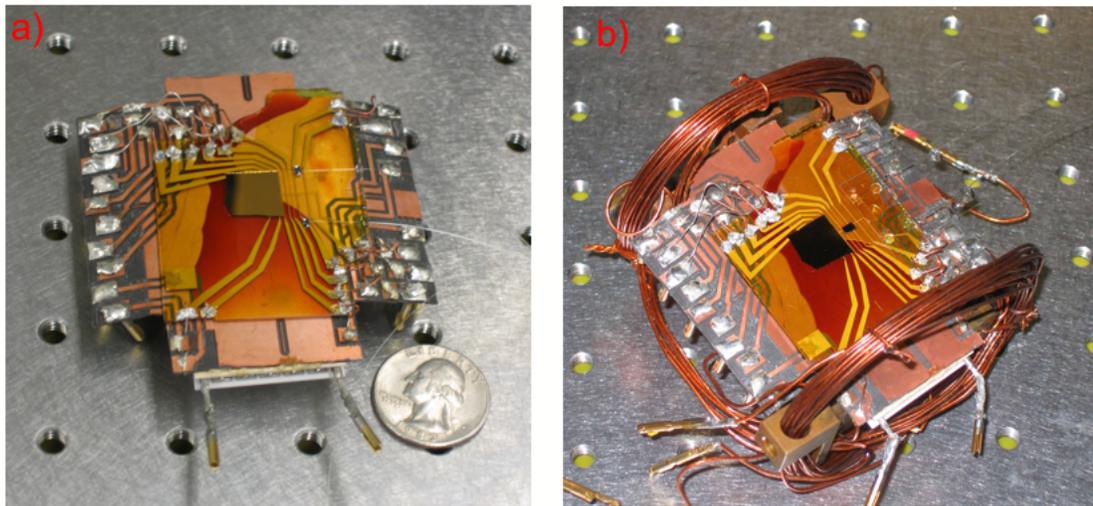


Figure 6.11: a) The atom-cavity chip to be used in the experiment. The fiber taper is attached. b) The atom chip with microdisk cavity mounted to the chip assembly block. The fiber taper is coupled to a microdisk cavity on the SiN_x substrate.

the initial U-trap. The Z-trap waveguide design has the advantage that the straightforward addition of a few coplanar wires can serve as gates and loosely confine the atoms once they reach the cavity. We have demonstrated such a P-trap waveguide scheme in our lab at Caltech and in the experiments in Munich. Chapter 2 describes the operation and relative merits of these waveguides in detail.

In the cavity region, the atoms are suspended 400 to 500 μm above the surface of the waveguide's microwires. This allows enough room for the ~ 300 μm thin microcavity substrate and fiber taper mount to be placed in the gap between the atoms and the microwires. Once the atoms are transported to a position above the PBG or microdisk cavity, the current and bias field of the guide are adjusted to lower the cold atom cloud into the mode of the cavity. A thermoelectric cooler (TEC) is located underneath the atom chip to counteract heating due to the microwire waveguide and aids in maintaining a specific cavity detuning from the frequency of the driving laser and atomic resonance. We have not yet measured how well the microcavity is thermally isolated from the microwire heating, but a significant time lag in thermalization could aid in controlling the detuning during each experimental shot. We estimate a cavity tunability of 20 GHz/ $^{\circ}\text{C}$ for the semiconductor substrates and 5 GHz/ $^{\circ}\text{C}$ for the SiN_x . With TEC control of 10^{-2} $^{\circ}\text{C}$, we should be able to achieve a 200 MHz tuning resolution with the semiconductor substrates and 50 MHz with the SiN_x . This resolution is sufficient for the PBG cavity, since its linewidth is on the order of several GHz. The ~ 180 MHz loaded linewidth of the microdisk poses more of a difficulty with regards to detuning control, and operation in the dispersive regime might be necessary for initial experiments. For the PBG [microdisk] cavity, the vacuum Rabi splitting spans ± 16 [1.5] GHz, which is well within the temperature tuning capability of the TEC. For both the PBG and microdisk cavities an error signal could be derived by probing on a greatly detuned cavity resonance and separating this beam from the detection beam with a dichroic filter. However, this probe power needs to remain under ~ 1 μW to avoid melting the fiber taper. In the future, a tuning mechanism other than temperature would be desirable. For instance, a piezoelectric material incorporated with the substrate could tune through stress, or an electro-optical material imbedded in the cavity mode could tune the resonance by changing the index of refraction via DC electric fields. Chapter 3 contains more details regarding the laser probe detuning, detection, and cavity detuning control and locking.

To create a compact structure, the fiber tapers are glued to the substrate in the exact position necessary for coupling to the microcavity. The taper is formed by pulling a fiber by micro-steppers in a hydrogen torch. It is crucial that the fiber be introduced to the edge of the blue flame to prevent heated air currents from sucking the taper into the center of the flame. To produce straight tapers, the fiber should be held by the micro-steppers as close to the flame as possible. Transmission is monitored while pulling the taper which provides a measure of the taper's thickness. As it is thinned, the fiber changes from single-mode to multimode and back to single mode. In this last transition, the transmission begins to drop as the diameter becomes $\leq \lambda$, and the pulling should be halted at this point. The taper is quite fragile: The fiber taper is bent in a U-shape which applies natural strain on the taper, keeping it straight and rigid. Instead of a U-shaped fiber taper, we tried to glue straight-tensioned fiber tapers to the atom chip but found they broke even under the slightest chip agitation. We now glue the U-shaped fiber taper to the top of a thin ($\sim 120 \mu\text{m}$) glass coverslip and then glue this cover slip to the atom chip. The fiber taper is cantilevered off the coverslip with glue joints located a few millimeters from the tapered section (see Figure 6.10). This minimizes optical power loss at the joints: if placed too close to the tapered region the glue joints act as fiber position-to-optical power transducers. We use a low-shrink UV curable glue (Dymax, 50 W UV lamp, OP-66-LS or OP-4-20663) for securing the fibers. The OP-4-20663 glue is more transparent and easy to manipulate than the opaque OP-66-LS, though this latter glue seems to shrink less. The cure occurs within 3 s at room temperature. Although the UV glue has relatively little shrinkage, we found that the position of the fiber taper can move up to a micron in a random direction during the curing. This is not as important for the PBG waveguide coupling, but this movement does ruin the taper-to-microdisk coupling. By gluing the U-shaped taper to a coverslip before gluing the cover slip to the atom chip, we found that the taper can be accurately secured due to the added mass and rigidity the coverslip provides. In a vacuum chamber, the taper is susceptible to breaking during the initial roughing. This should be done as adiabatically as possible to avoid a quick pressure change inducing "wind." Both our group and that of Kimble/Vahala have noticed that for pressures ≤ 1 Torr, heat can no longer be carried away from the taper and mW's of optical power melts the fiber taper. We have found that the tapers can withstand at least $1 \mu\text{W}$ of power in the vacuum chamber and perhaps slightly more. As discussed in the following section, optical powers no greater than a few 100 nW

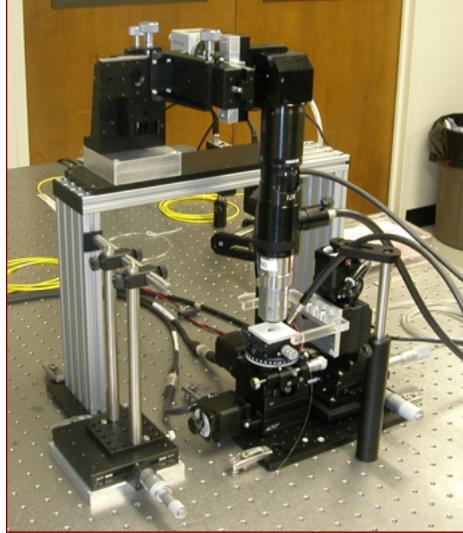


Figure 6.12: The fiber taper-to-atom chip alignment set-up. A 3D translation stage with a two-axis goniometer stage assists in cavity-taper coupling alignment. The transmission is probed with a scanning New Focus Velocity laser, and the alignment is monitored optically through a high-power objective.

are required for single atom detection.

The delivery scheme described above provides a non-deterministic source of weakly trapped atoms to the cavity mode. For the PBG cavity, the field of the cavity mode is concentrated in the central ~ 10 holes (see Figure 6.6 (c)). For the specific PBG design shown in Figure 6.6, the field has two maxima, one in each of the two central holes and offset by 45 nm from each hole's axis. We expect to transport a few times 10^5 atoms in a cigar-shaped cloud of density $10^{11}/\text{cm}^3$. The cross-sectional area of this cloud parallel to the chip is larger than the $0.4 \mu\text{m}^2$ area of the PBG cavity that is occupied by the field, and we estimate there is a $\sim 10\%$ probability of an atom encountering one of the central 10 holes per cloud interaction. With an experimental repetition once every ~ 5 seconds—limited by the U-MOT replenishing time—we foresee the accumulation of a significant number of events in a reasonable amount of time. As discussed in Section 6.4.1 below, we expect to detect strong signals during single atom transits through the PBG cavity's central holes. If we assume a cesium cloud temperature of $10 \mu\text{K}$, then a cesium atom whose velocity is parallel to the axis, \hat{z} , of a central hole will interact with the mode for a time duration of $\sim 10 \mu\text{s}$. The microdisk cavity has a mode volume ~ 10 times larger than the PBG cavity, and because the cavity diameter—typically $< 10 \mu\text{m}$ —is smaller than the cloud extent, single

atom transits should be detected more frequently than in the PBG cavity experiment.

6.4 Single atom detectability using PBG cavities

To investigate the PBG cavity's response to a strongly coupled atom falling through a central hole, we solve the semi-classical optical bistability equation for a qualitative understanding of the interaction and—using a two-level atom—the quantum master equation to obtain a more quantitative description. Although neither of these treatments fully encompasses the complexity of the system, we presume that they are sufficient for demonstrating the feasibility of the device for single atom detection. For tractability, these calculations ignore the fact that g_0 and the detunings are of the same order or much larger than both the hyperfine ground-state and excited-state splittings, which for cesium are 9.2 GHz and 151 to 251 MHz, respectively. In other words, the atom-photon coupling is much stronger than the coupling between the electron and nuclear spins. This is an unusual situation and requires a full quantum calculation of the atom-PBG cavity interaction that includes the full cesium D2 manifold of states. Kevin Birnbaum and Scott Parkins recently wrote a Matlab code (using the Quantum Optics Toolbox by Sze Tan [160, 161]) to account for the full manifold of states and the optical pumping from a linearly polarized probe beam. Details and results of this calculation are presented in Kevin's thesis [83]. The simulation was run at weak driving—intracavity photon number much less than unity. Qualitatively, the vacuum Rabi peaks persist, but appear at different detunings and with different widths than expected from the two-level atom treatment. In addition, at least two sharp transmission peaks emerge. These might be caused by quantum interference phenomena from the field coupling to the multitude of atomic states [162]. The code used for the simpler simulations presented in this chapter and Chapter 7 is included in Appendix B, Sections B.1 and B.1.

The optical bistability equation is a semi-classical description of the transmission of a cavity containing atoms [163],

$$y = \frac{x}{\left[\left(1 + \frac{2}{N_0(1+(\Delta/\gamma_\perp)^2+y^2)} \right)^2 + i \left(\frac{\theta}{\kappa} - \frac{2\Delta}{\gamma_\perp N_0(1+(\Delta/\gamma_\perp)^2+y^2)} \right)^2 \right]^{\frac{1}{2}}}. \quad (6.5)$$

In the above equation, x is the input field, $E/\sqrt{m_0}$, where E is the amplitude of the driving field; y is the output field, $\alpha/\sqrt{m_0}$, where α is the intracavity coherent state amplitude; Δ

is the atom-laser detuning; and θ is the cavity-laser detuning³. This semi-classical equation is derived under the assumption that operators can be replaced by their expectations (i.e. $\hat{a} \rightarrow \langle \hat{a} \rangle = \alpha$, where α is a c-number) and joint operator moments may be factored (i.e. $\langle \hat{a}\hat{\sigma}^\dagger \rangle \rightarrow \langle \hat{a} \rangle \langle \hat{\sigma}^\dagger \rangle$) [25]. These approximations break down for the PBG and microdisk cavity experiments since quantum fluctuations can no longer be ignored when $[m_0, N_0] \ll 1$ for a single intracavity atom.

The solutions to the unconditional master equation paint a more accurate picture of the atom-cavity system. Under the two-level atom, electric dipole, and rotating-wave approximations, the equation for the density matrix, ρ , of the joint state of the atom and cavity is as follows:

$$\begin{aligned} \dot{\rho} = & \frac{-i}{\hbar} [\hat{H}_0, \rho] + \gamma_{\perp} (2\hat{\sigma}\rho\hat{\sigma}^\dagger - \hat{\sigma}^\dagger\hat{\sigma}\rho - \rho\hat{\sigma}^\dagger\hat{\sigma}) \\ & + \kappa(2\hat{a}\rho\hat{a}^\dagger - \hat{a}^\dagger\hat{a}\rho - \rho\hat{a}^\dagger\hat{a}), \end{aligned} \quad (6.6)$$

$$\hat{H}_0 = \hbar\Delta\hat{\sigma}^\dagger\hat{\sigma} + \hbar\theta\hat{a}^\dagger\hat{a} + i\hbar E(\hat{a}^\dagger - \hat{a}) + \hat{H}_{int}, \quad (6.7)$$

$$\hat{H}_{int} = i\hbar g_0 \psi(\hat{r}) [\hat{a}^\dagger\hat{\sigma} - \hat{\sigma}^\dagger\hat{a}]. \quad (6.8)$$

In this equation, $\hat{\sigma}$ is the atomic lowering operator and \hat{a} is the cavity field annihilation operator. Along the axis of the central cavity hole, the mode function, $\psi(z)$, closely approximates a Gaussian of width ~ 225 nm, centered about the midpoint of the ~ 170 nm thick cavity membrane (the hole diameter is ~ 100 nm). In $\psi\rho$, the field is roughly azimuthally symmetric and decreases in magnitude towards the walls of the defect hole in an approximately Gaussian fashion. The steady-state density operator, ρ_{ss} , as a function of various drive strengths, coupling strengths, and detunings is found by solving equation 6.6 with $\dot{\rho}_{ss} = 0$. Operator expectations are $\langle \hat{O} \rangle = \text{Tr}[\rho_{ss}\hat{O}]$.

The coupled (dressed) eigenstates of the atom-cavity are characterized by the Jaynes-Cummings ladder as depicted in Figure 6.13. The dressed states at each rung of the ladder are split by $\sqrt{n}\hbar g_0$, where $n \equiv \langle \hat{a}^\dagger\hat{a} \rangle$ is the number of photons in the intracavity field. As n becomes large, $\sqrt{n+1}/\sqrt{n} \rightarrow 1$, and the ladder rungs become equally spaced by $\hbar\Omega$, where Ω is the Rabi frequency. In this case, the system exhibits phase bistability [164] and the system resembles that of a strongly-driven free-space atom: a Mollow triplet is exhibited in fluorescence, and power broadening in absorption. When $n \ll 1$ the atom-cavity system

³ $\Delta = \omega_a - \omega_l$ and $\theta = \omega_c - \omega_l$.

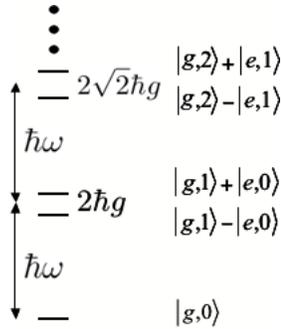


Figure 6.13: Jaynes-Cummings ladder of atom-cavity eigenstates for $[\Delta, \theta] = 0$.

is in the linear, weak driving regime and only the lowest rung is driven. This latter case is depicted for the PBG cavity in Figure 6.14 (b) and (c), black \diamond curve (though E should be much less than one rather than 0.1 to truly represent this regime). The vacuum Rabi peaks are separated by $2g_0$ and of width set by the mean of κ and γ_{\perp} .

6.4.1 Single atom detection signal-to-noise

The remainder of this section is devoted to analyzing the master equation for various drive powers and detunings. Our main purpose is to judge the feasibility of detecting single atoms using PBG cavities, and a thorough study of the phenomena associated with the Jaynes-Cummings Hamiltonian is beyond the scope of this chapter. For more information on atom-cavity structure and dynamics, see References [67, 134] and the various theses from H. J. Kimble's group [165, 166, 153, 21, 22, 154]. Kevin Birnbaum's thesis contains information specific to a high- g_0 cavity QED system [83].

The expected cavity output in photons per detector integration time, Δt , and photon detection efficiency, η , is⁴

$$N = 2\kappa\Delta t\eta\langle\hat{a}^{\dagger}\hat{a}\rangle, \quad (6.9)$$

with noise fluctuations of variance

$$(\Delta N)^2 = 2\kappa\Delta t\eta(\langle\hat{a}^{\dagger}\hat{a}\hat{a}^{\dagger}\hat{a}\rangle - \langle\hat{a}^{\dagger}\hat{a}\rangle^2), \quad (6.10)$$

For all the parameters considered here in the strong coupling regime, the intracavity field is in not quite a coherent state, and the transmitted field varies from sub- to super-Poissonian

⁴The energy decay rate of the cavity is 2κ . All mentions of g_0 , κ , and γ_{\perp} include the factor of 2π .

counting statistics as the intracavity field sweeps across a mean intracavity photon number of ~ 1 . However, the field statistics are never far from Poissonian, and for computational simplicity we assume a coherent field when calculating shot noise: $\Delta N \approx N_{\text{shot}} = \sqrt{N}$.

We define the signal corresponding to an atom transit as

$$S = 2\kappa\eta\Delta t(\text{abs}[\langle\hat{a}^\dagger\hat{a}\rangle_1 - \langle\hat{a}^\dagger\hat{a}\rangle_0]), \quad (6.11)$$

where $\langle\hat{a}^\dagger\hat{a}\rangle_i$ is the intracavity photon number in a (possibly detuned) cavity with an intracavity atom ($i = 1$) and without an intracavity atom ($i = 0$). The noise obscuring S is caused by shot noise on both of these detected cavity transmissions:

$$S_{\text{noise}} = \sqrt{2\kappa\eta\Delta t[\langle\hat{a}^\dagger\hat{a}\rangle_1 + \langle\hat{a}^\dagger\hat{a}\rangle_0]}. \quad (6.12)$$

The signal-to-noise, \mathcal{S}/\mathcal{N} , for single atom detection is $\mathcal{S}/\mathcal{N} \equiv S/S_{\text{noise}}$.⁵ The number of spontaneously emitted photons during an atom transit detection of time Δt is $N_{\text{Spont}} = \gamma\Delta t\langle\sigma^\dagger\sigma\rangle$.⁶ Note that instead of photon counting, heterodyne detection may be used, in which case expectations of \hat{a} rather than $\hat{a}^\dagger\hat{a}$ are the relevant quantities. The results of the simulations presented here are qualitatively similar for either case, though the calculated \mathcal{S}/\mathcal{N} might be lower for heterodyne detection [153].

Figure 6.14 shows master equation simulations of the PBG cavity transmission as a function of (a) atom detuning [$\theta = 0$], (b) laser detuning [$\Delta = \theta$], and (c) cavity detuning [$\Delta = 0$]. We assume cavity QED parameters of $[g_0, \kappa]/2\pi = [16, 4.4]$ GHz for the PBG cavity (see Table 6.1). As the intracavity power is increased from weak to strong driving (black \diamond to blue \square to red \circ), the vacuum Rabi peaks—initially separated by $2g_0$ —fill-in towards zero detuning as higher rungs on the Jaynes-Cummings ladder are excited. Figures 6.15 through 6.19 are slices of Figure 6.14 for various detunings chosen to highlight different atom-cavity response regimes. Panels (a) plot the transmission versus drive power which is measured in intracavity photon numbers for a resonant and empty cavity. The red curve is the empty cavity transmission at the given θ , the solid black curve is the solution to the

⁵This expression is a lower bound on the \mathcal{S}/\mathcal{N} since for post-processed detection of the atom transits we could average away the shot-noise from the empty cavity transmission and exclude $\langle\hat{a}^\dagger\hat{a}\rangle_0$ from Equation 6.12. However, for real-time measurements—in quantum-limited feedback experiments, for instance— $\langle\hat{a}^\dagger\hat{a}\rangle_0$ should be included.

⁶ $\gamma = 4\pi\gamma_\perp$

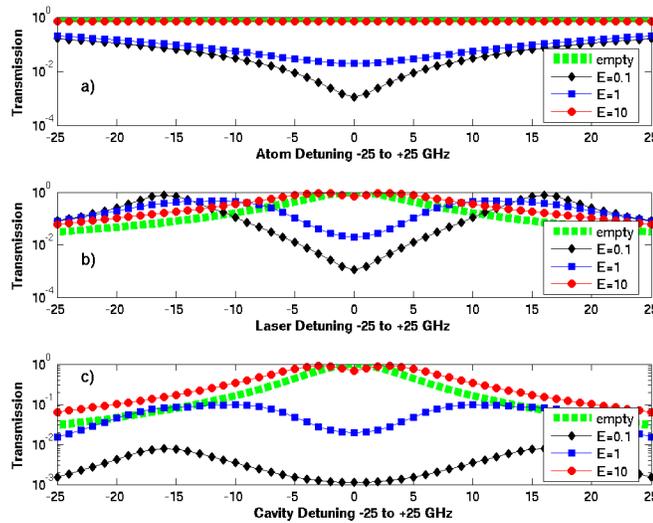


Figure 6.14: PBG cavity transmission as a function of a) atom detuning [$\theta = 0$], b) laser detuning [$\Delta = \theta$], and c) cavity detuning [$\Delta = 0$] for one intracavity atom and $[g_0, \kappa]/2\pi = [16, 4.4]$ GHz. The drive, E (measured as the empty cavity intracavity photon number on resonance) is $E = 0.1$ in the black \diamond curve, $E = 1$ in the blue \square curve, and $E = 10$ in the red \circ curve. The green solid curve is the empty cavity transmission. All curves are normalized to the empty cavity transmission at each curve's drive strength.

semi-classical optical bistability equation 6.5, and the black dots are the solutions to the master equation 6.6.

The semi-classical solution is added as both a guide to the eye and as an indicator of the transition from linear to non-linear system response. The saturation photon number designates this onset, $\langle \hat{a}^\dagger \hat{a} \rangle_1 = m_0$, in the semi-classical regime. The most striking aspect of the two solutions pictured in panels (a) is the “washing out” of the bistability in the quantum treatment for this strongly-coupled, on-resonant single-atom cavity system. In the single atom and low intracavity photon number regime, the quantum fluctuations are non-negligible $n \approx \sqrt{n}$, invalidating the assumptions of the semi-classical description. In loose terms, the quantum fluctuations are of the same order of magnitude as the hysteresis loop, and in the steady-state behavior of the system, bistability is averaged away. One should note that not all bistability is washed-out in steady-state functions of the strongly-coupled QED system: one can find detuned parameter regimes where amplitude bistability persists [32]; and moreover, the Q-function of the strongly-coupled system at high drive powers ($E \gg 1$) exhibits a bifurcation in phase [164].

With the cavity-laser detuning set to zero, Figures 6.15, 6.16, and 6.18—in which $\theta = 0$ but $\Delta/2\pi = [0, 9.2, 16]$ GHz, respectively—show that a deficit of photons transmitted through the cavity—a “down-transit”—can be detected for a drive of a few intracavity photons. The $\Delta/2\pi = 16$ GHz detuning is chosen to probe the atom-cavity eigenstate, and the $\theta = 0$, $\Delta/2\pi = [0, 9.2]$ GHz detunings are chosen for their ease in experimental implementation: the cesium hyperfine splitting is 9.2 GHz and the laser can easily be locked to the lower hyperfine transitions, naturally providing this detuning. Because we do not have good control over the cavity resonance, it is much simpler to match the laser resonance with the cavity [$\theta = 0$] and detune the laser from the atom. Chapter 3 discusses these issues in more detail. Figures 6.17 and 6.19 plot transmission of the system with the atom and cavity on-resonance with each other: laser detunings of 9 GHz and 16 GHz, respectively. In these detuning regimes, an excess of transmitted photons—“up-transits”—are observed during the presence of an intracavity atom.

Panels (b) and (c) of Figures 6.15 through 6.19 show the signal, S , and signal-to-noise, S/\mathcal{N} , respectively, that we expect from probing the PBG atom-cavity system at these detunings and drive powers, and with a detection efficiency of $\eta = 0.44$ and integration time of $\Delta t = 1 \mu\text{s}$. Table 6.2 summarizes the achievable S/\mathcal{N} for each detuning set. The “Maximum S/\mathcal{N} ” set of columns lists the maximum S/\mathcal{N} obtained by optimizing over drive power. The highest S/\mathcal{N} ’s are attained for $\theta = 0$ and drive powers of ~ 40 nW. For the $\theta \neq 0$ case, the S/\mathcal{N} ’s—though lower—saturate at nearly a factor of ten higher drive powers. This could be quite useful for ameliorating the typically stringent detector requirements for implementing shot noise-limited detection (see Chapter 3). While these S/\mathcal{N} ’s of ~ 150 are encouraging, the number of spontaneously emitted photons per $1 \mu\text{s}$ integration, $N_{\text{SponE}}/\mu\text{s}$, at these drive powers is not insignificant (though constitutes only $\sim 0.01\%$ of S). This is not terribly detrimental to our initial single-atom detection experiments since we mainly want to know if the atom is coupled and the driving light at these large detunings doubles as a repumping beam: brief signals are acceptable. However, for future experiments we would like to have the atom magnetically trapped in the Lamb-Dicke regime for long periods of time using an Ioffe trap. Both heating and optical pumping to untrapped states due to the spontaneous emissions will limit the lifetime of the magnetic trap, thereby reducing the coupling time of the atom-cavity system. Calculations of the sort performed by K. Birnbaum and S. Parkins [83, 162] should be investigated to understand the severity of

Table 6.2: Comparison of \mathcal{S}/\mathcal{N} for various cavity detunings. For each cavity system and detuning, the table’s columns are divided into two sections: the first lists the cavity drive that maximizes \mathcal{S}/\mathcal{N} regardless of N_{SponE} during a $\Delta t = 1 \mu\text{s}$ detection, and the second lists the \mathcal{S}/\mathcal{N} for a cavity drive that limits the $N_{SponE}/\mu\text{s}$ to equal unity. The drive is listed here as both the resonant, empty intracavity photon number $\langle \hat{a}^\dagger \hat{a} \rangle_{r0}$ and the cavity input drive power $P = 2\kappa\hbar\omega\langle \hat{a}^\dagger \hat{a} \rangle_{r0}$ to facilitate comparisons between panel (a) and panels (b) and (c) in Figures 6.15 through 6.19 and Figures 6.24 through 6.26. The “ \geq ” symbols indicate that the maximum \mathcal{S}/\mathcal{N} is possibly outside the range of drive powers accessible to this Fock basis-limited simulation.

Cavity	$[\Delta, \theta]/2\pi$ (GHz)	Maximum \mathcal{S}/\mathcal{N}			$N_{SponE}/\mu\text{s} = 1$	
		\mathcal{S}/\mathcal{N}	$[\langle \hat{a}^\dagger \hat{a} \rangle_{r0}, \text{P(nW)}]$	$N_{SponE}/\mu\text{s}$	\mathcal{S}/\mathcal{N}	$[\langle \hat{a}^\dagger \hat{a} \rangle_{r0}, \text{P(nW)}]$
PBG	[0, 0]	208	[3, 39]	11	98	[0.4, 5.2]
	[9.2, 0]	198	[3, 39]	10	94	[0.4, 5.2]
	[9.2, 9.2]	≥ 141	$\geq [20, 258]$	15	9.7	[0.2, 2.6]
	[16, 0]	181	[3, 39]	9	87	[0.4, 5.2]
	[16, 16]	≥ 97	$\geq [40, 516]$	13	26	[0.04, 0.5]
Microdisk	[0, 0]	154	[15, 7.7]	12	65	[2, 1]
	[1.5, 0]	149	[15, 7.7]	11	64	[2, 1]
	[1.5, 1.5]	≥ 30	$\geq [50, 25]$	12	8	[0.04, 0.5]

the optical pumping, even for these large probe detunings. The heating problem is more tractable. Each spontaneously emitted photon has a probability—suppressed by the Lamb-Dicke parameter—to excite the atom out of the trap’s ground state and into ever higher vibrational levels. One needs to calculate the average N_{SponE} —and average time—it takes for the atom’s wavepacket to be lost either by being excited to an energy larger than the trap’s depth or by being extended to such a width that the atom has a chance to encounter the wall of the PBG cavity. The magnetic trap’s depth, oscillation frequencies, and Lamb-Dicke parameter can be accurately known, ensuring the accuracy of this calculation. The second set of columns, under the “ $N_{SponE}/\mu\text{s} = 1$ ” heading, lists the \mathcal{S}/\mathcal{N} for drive powers that limit N_{SponE} to be equal to unity during a $\Delta t = 1 \mu\text{s}$ integration time. This demarcates the drive power below which spontaneous emission becomes negligible. For drives of a few nW—roughly ten times lower than for the maximum \mathcal{S}/\mathcal{N} case—the signal-to-noise is still significant for most detuning regimes. It may be possible to find other detuning regimes that exhibit better \mathcal{S}/\mathcal{N} while minimizing N_{SponE} .

6.4.2 Simulated atom transits and cavity induced force

Simulated photon counts during atom transits are shown in Figures 6.20 and 6.21. We assume the atom moves with constant velocity, $v = 2.5 \text{ cm/s}$, through the axis of the cavity mode $\psi(z)$, making a full transit of the Gaussian waist in $10 \mu\text{s}$. In both plots the

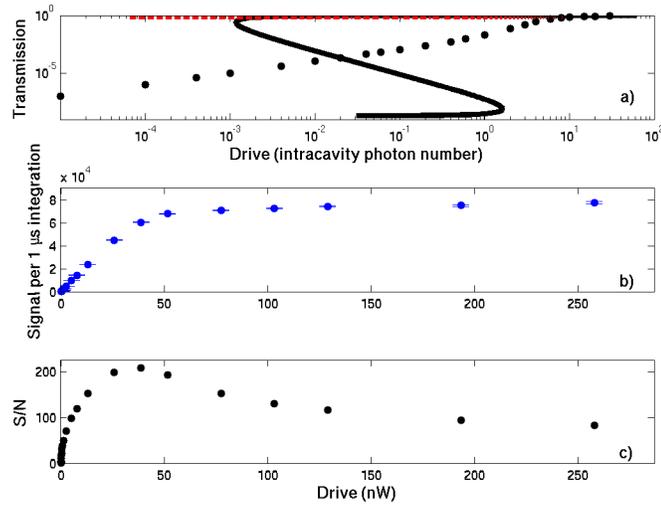


Figure 6.15: (a) The transmission of the PBG cavity as a function of drive strength—measured in intracavity photon numbers for a resonant and empty cavity—calculated from Equations 6.5 (black line) and 6.6 (black dots). The empty cavity transmission is shown as a dashed red line. (b) The expected signal S as a function of drive power—quoted as $\hbar\omega$ times the resonant, empty intracavity photon number—during an atom transit with a $\Delta t = 1 \mu\text{s}$ integration time. Error bars correspond to S_{noise} . (c) The expected signal-to-noise \mathcal{S}/\mathcal{N} in a $\Delta t = 1 \mu\text{s}$ integration time. For this plot, $[g_0, \kappa]/2\pi = [16, 4.4]$ GHz and $[\Delta, \theta] = [0, 0]$.

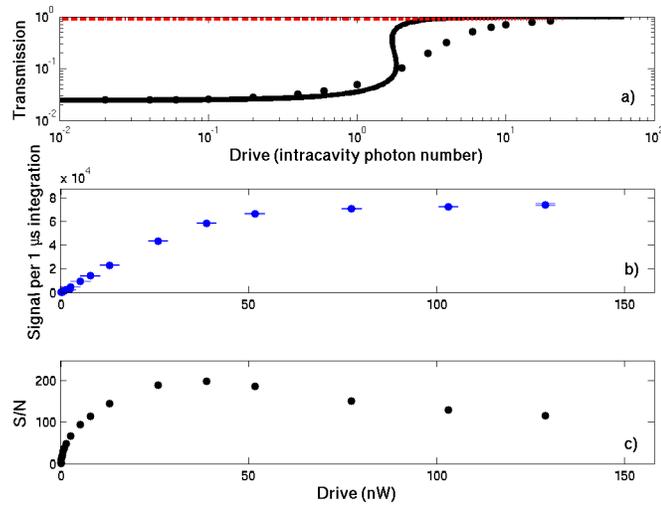


Figure 6.16: Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [9.2, 0]$ GHz.

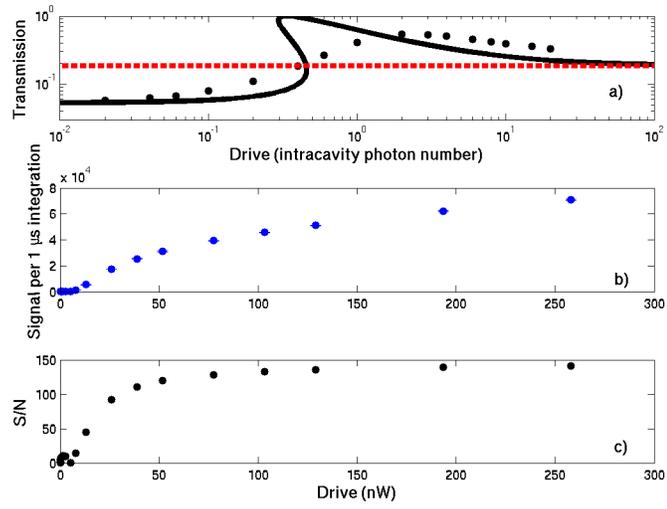


Figure 6.17: Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [9.2, 9.2]$ GHz.

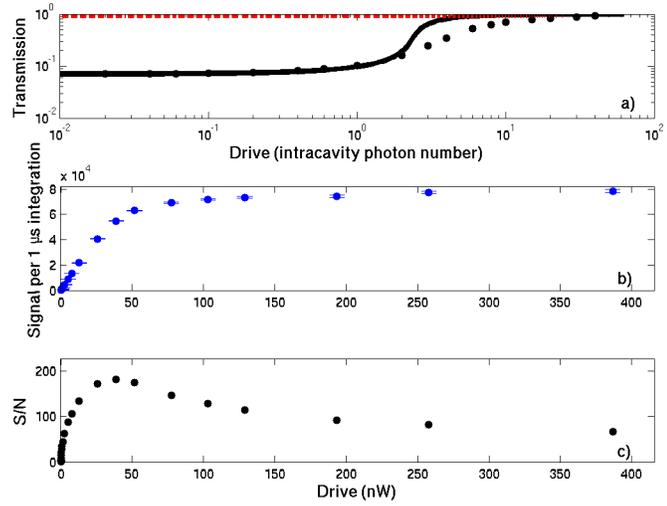


Figure 6.18: Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [16, 0]$ GHz.

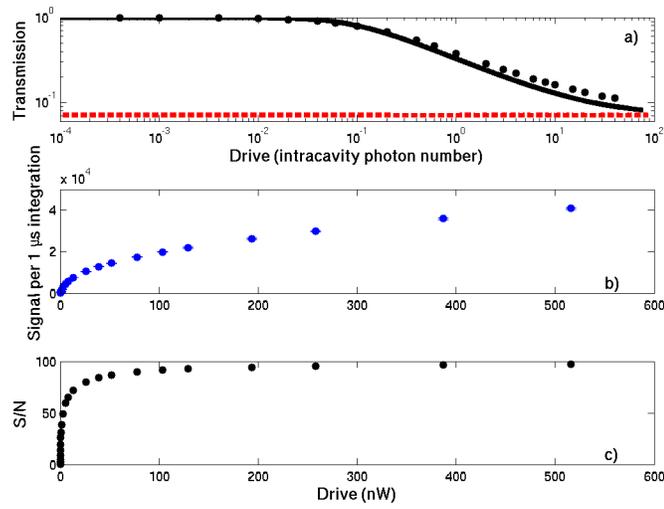


Figure 6.19: Same as Figure 6.15 except $[\Delta, \theta]/2\pi = [16, 16]$ GHz.

drive strength is 2 intracavity photons. As the atom transverse the cavity, the coupling $g(t) = g_0\psi(vt)$ also varies as a Gaussian, which modulates the output photon flux. The mean photon count, N , and variance, $(\Delta N)^2$, are found by solving for ρ_{ss} for each $g(t)$ in time steps of $\Delta t = 1 \mu\text{s}$, chosen to simulate a finite bandwidth photodetector. Each point includes additional shot-noise selected randomly from a normal distribution of standard deviation ΔN . The figures show that even with shot-noise, both up-transits (a) and down-transits (b) of single atoms through the axis of the central PGB cavity hole are clearly detectable. Moreover, it seems possible to detect atom transits that only experience 20% to 30% of g_0 . During an experiment, we expect to detect a low background of signals from marginally coupled atoms—such as those grazing the field extending from the surface of the PBG membrane or slipping into holes away from the central region—punctuated by sharp pikes representing atoms fully coupled to the field inside the central holes. It should be noted that the mean photon numbers and noise in Figures 6.20 and 6.21 are not derived from a quantum trajectory calculated from the conditional master equation [25], but are simply calculated using ρ_{ss} from the unconditional equation 6.6. This is acceptable given the inherent limitations of the model as mentioned at the beginning of this section.

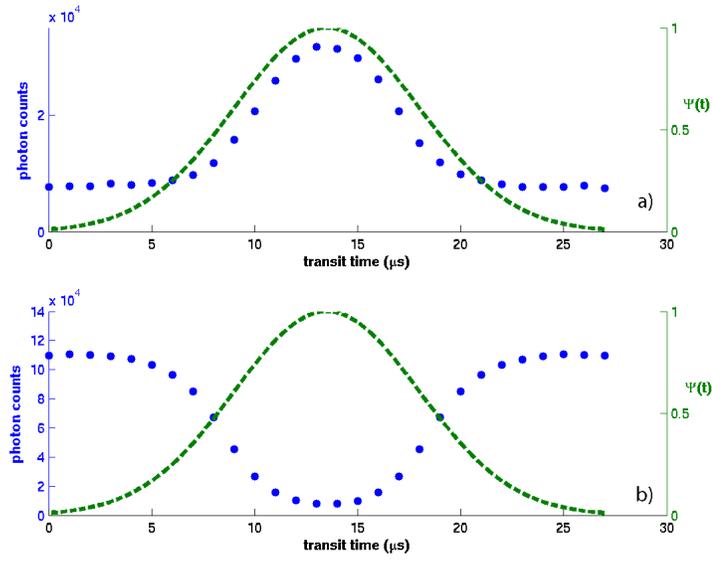


Figure 6.20: Simulated photon counts due to atom transits through the axis of the cavity's central hole. Blue dots (left axis) are the photon counts, and the green, dashed curve (right axis) is the Gaussian variation of $g(t)/g_0 = \psi(z(t))$ experienced by the atom during its transit. Calculations are for detunings of (a) $[\Delta, \theta]/2\pi = [16, 16]$ GHz and (b) $[\Delta, \theta]/2\pi = [0, 0]$ GHz. Drive power is set for optimal \mathcal{S}/\mathcal{N} .

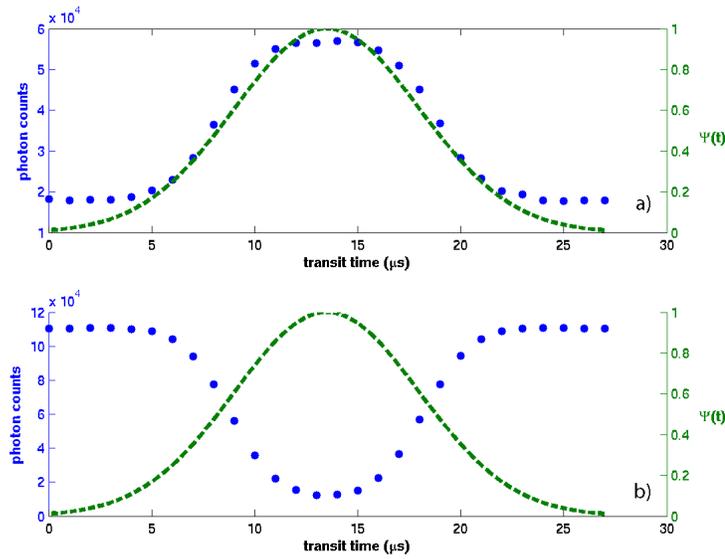


Figure 6.21: Same as Figure 6.20. Calculations are for detunings of (a) $[\Delta, \theta]/2\pi = [9.2, 9.2]$ GHz and (b) $[\Delta, \theta]/2\pi = [9.2, 0]$ GHz. Qualitatively, the atom transit signals are nearly identical to those in Figure 6.20.

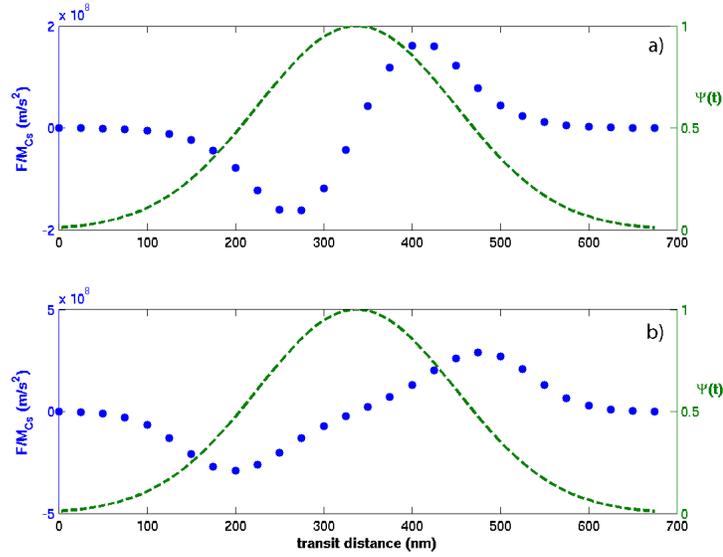


Figure 6.22: The force on an atom traversing the PBG cavity mode ($[g_0, \kappa]/2\pi = [16, 4.4]$ GHz). a) $[\Delta, \theta]/2\pi = [16, 16]$ GHz. b) $[\Delta, \theta]/2\pi = [16, 0]$ GHz. Drive power is set for optimal \mathcal{S}/\mathcal{N} .

The atom will experience a force,

$$\langle \vec{f} \rangle = -i\hbar \nabla g(\vec{r}) \langle \hat{a}^\dagger \hat{\sigma} - \hat{a} \hat{\sigma}^\dagger \rangle, \quad (6.13)$$

as it encounters the cavity mode. Figure 6.22 shows simulations of this force on a cesium atom as it traverses the mode of a central hole of the PBG cavity. The maximum acceleration on an atom dragged through the cavity mode at velocity 2.5 cm/s is $|\langle f_{max} \rangle|/M_{Cs} \approx 2 \times 10^8$ m/s², corresponding to a change in velocity of

$$\Delta v = \sqrt{\frac{|\langle f_{max} \rangle| \Delta z}{M_{Cs}}} \approx 5 \text{ m/s} \quad (6.14)$$

over half the length of the cavity mode, $\Delta z = 100$ nm. In the above equations, M_{Cs} is the mass of a cesium atom. This agrees with a simple estimate using

$$\hbar g_0 = 0.5 M_{Cs} (\Delta v)^2, \quad (6.15)$$

which yields $\Delta v = 10$ m/s. Fabry-Perot experiments have detected effects of the cavity interaction on the atomic motion [167]. The simple estimate using equation 6.15 gives a

smaller value of $\Delta v \approx 0.7$ m/s for the Fabry-Perot experiments, implying that the motion of the atom traversing the mode of the PBG cavity will also be significantly affected. A more detailed calculation [168, 169] of the force and momentum diffusion using a master equation beyond the two-level atom approximation is necessary to make predictions about the behavior of an atom in an attractive, red-detuned cavity mode or in a repulsive, blue-detuned mode. The close proximity of the atom to the sides of the PGB cavity's holes will surely affect the system's dynamics due to the Casimir-Polder potential [170], and this will need to be addressed in more detailed simulations. In future experiments, the force on the atom due to the Casimir-Polder potential may be counterbalanced by careful cavity light detuning and/or by using the magnetic microtrap. Regardless, any experimental information about this potential in this unique geometry will be quite interesting.

6.5 Single atom detectability using microdisk cavities

We repeat a similar analysis as in Section 6.4.1 for the feasibility of detecting single atoms with a microdisk cavity. For the master equation simulations, we assume a SiN_x microdisk of diameter ~ 8 μm and a (loaded) Q of 1×10^6 , which is a factor of 3 less than that quoted in Table 6.1 for the intrinsic Q attainable in current experiments. This results in cavity QED parameters of $[g_0, \kappa]/2\pi = [1.5, 0.18]$ GHz. Figure 6.23 plots the microdisk cavity transmission versus various detunings, and Figures 6.24, 6.25, and 6.26 are transmission versus drive power plots for $[\Delta, \theta]/2\pi = [0, 0]$, $[1.5, 0]$ GHz, and $[1.5, 1.5]$ GHz, respectively. As with the PBG cavity, the $[\Delta, \theta] = [0, 0]$ and $[\Delta = g_0, \theta]/2\pi = [1.5, 0]$ GHz cases show that an atom down-transit can be detected, while the $[\Delta = g_0, \theta = g_0]/2\pi = [1.5, 1.5]$ GHz detuning enables the detection of an up-transit. Adjustment of θ is not as experimentally feasible as Δ , and down-transit detection will be attempted initially.

Table 6.2 lists the projected \mathcal{S}/\mathcal{N} for the three sets of representative detunings. As discussed in the previous section, spontaneous emission is detrimental to the lifetime of the atom in a magnetic trap. We can achieve \mathcal{S}/\mathcal{N} 's that are slightly less than half that for optimal cavity driving even when limiting the drive power so that the number of spontaneously emitted photons is equal to 1 during the detection time $\Delta t = 1$ μs . Single atom detection with microdisks seems feasible with drive powers in the tens of nW. For the $[\Delta = g_0, \theta = g_0]$ case, \mathcal{S}/\mathcal{N} is maintained out to a drive power of 25 nW and beyond, which eases the de-

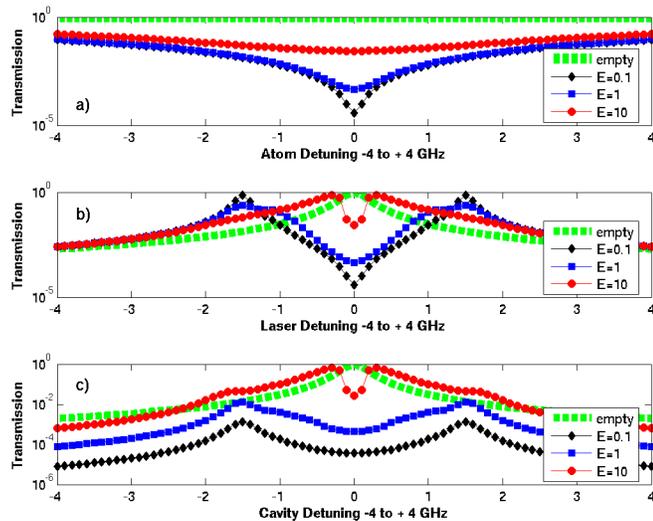


Figure 6.23: Microdisk cavity transmission as a function of a) atom detuning [$\theta = 0$], b) laser detuning [$\Delta = \theta$], and c) cavity detuning [$\Delta = 0$] for one intracavity atom and $[g_0, \kappa]/2\pi = [1.5, 0.18]$ GHz. The drive, E (measured as the empty cavity intracavity photon number on resonance) is $E = 0.1$ in the black \diamond curve, $E = 1$ in the blue \square curve, and $E = 10$ in the red \circ curve. The green solid curve is the empty cavity transmission. All curves are normalized to the empty cavity transmission at each curve's drive strength.

tector requirements for shot noise limited detection (see Chapter 3). The resonance of the microdisk exhibits a double peak due to the degeneracy lifting of two counterpropagating modes by disk defects [171]. These modes are out-of-phase standing-waves and are separated in frequency by a few hundred MHz. In future experiments, the two modes could be driven separately—one blue-detuned, one red-detuned—to form an optical trap for an atom near the rim. Effects due to the simultaneous coupling of these two driven modes to the atom should be observable.

6.6 Conclusion

The integration of atom trapping and cooling with photonic bandgap and microdisk cavities on a chip introduces a robust and scalable cavity QED system to the toolbox of nanotechnology. A device allowing cooled neutral atoms to be delivered via a magnetic microtrap and waveguide to the mode of a graded lattice PBG cavity or a microdisk cavity is feasible given present technology. Calculations using the semi-classical optical bistability equation

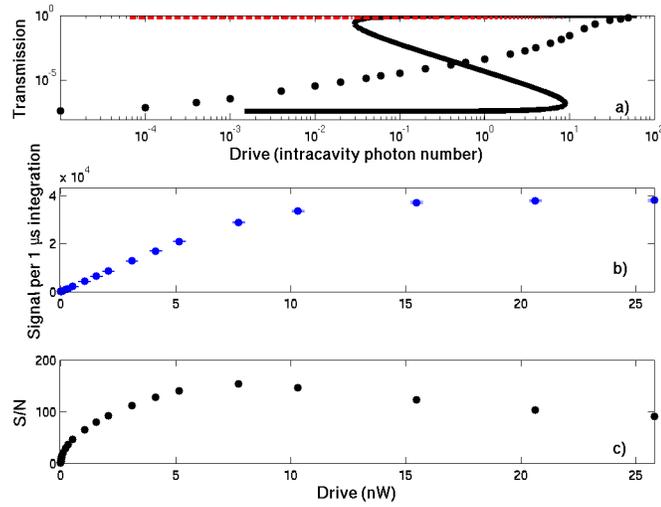


Figure 6.24: (a) The transmission of the microdisk cavity as a function of drive strength—measured in intracavity photon numbers for a resonant and empty cavity—calculated from equations 6.5 (black line) and 6.6 (black dots). The empty cavity transmission is shown as a dashed red line. (b) The expected signal S as a function of drive power (quoted as $\hbar\omega$ times the resonant, empty intracavity photon number) during an atom transit with a $\Delta t = 1 \mu\text{s}$ integration time. Error bars correspond to S_{noise} . The detection efficiency is assumed to be $\eta = 0.97$. (c) The expected signal-to-noise \mathcal{S}/\mathcal{N} in a $\Delta t = 1 \mu\text{s}$ integration time. For this plot, $[g_0, \kappa]/2\pi = [1.5, 0.18]$ GHz and $[\Delta, \theta] = [0, 0]$.

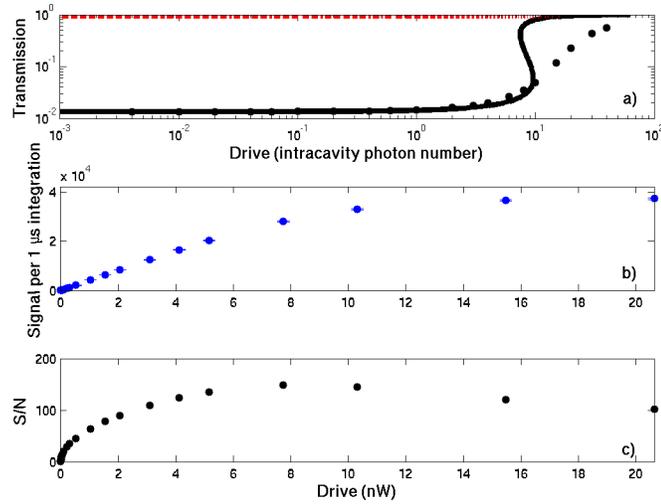


Figure 6.25: Same as Figure 6.24 except $[\Delta, \theta]/2\pi = [1.5, 0]$ GHz.

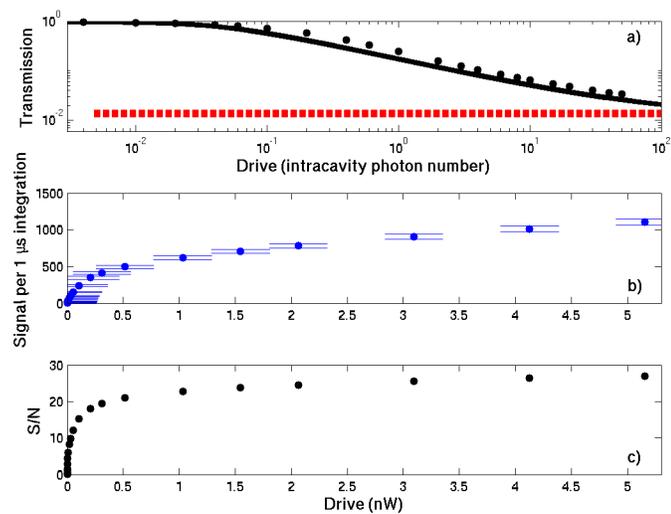


Figure 6.26: Same as Figure 6.24 except $[\Delta, \theta]/2\pi = [1.5, 1.5]$ GHz.

and the unconditional master equation indicate that it will be possible to detect single strongly-coupled atoms with this atom-cavity chip.

Chapter 7

Fiber-Gap Fabry-Perot Cavity: The first atom-cavity chip

This chapter describes an ongoing experiment performed in collaboration with Professors J. Reichel and T. W. Hänsch and their student and postdoc, Tilo Steinmetz and Dr. Yves Colombe, respectively. The experiment is located at the MPQ/LMU in the Schellingstraße labs. Several of the figures in Section 7.2 were made by D. Hunger and the members of the Munich group listed above.

7.1 Introduction

This experiment constitutes the first demonstration of an atom-cavity chip: atoms magnetically delivered to and trapped within the mode of an on-chip optical resonator. The presence of the atoms modulates the cavity transmission with a detection sensitivity approaching the level of a single atom. This is a step forward along the route to create miniaturized and robust single atom detectors, and is an important addition to the toolbox of nanotechnology. Moreover, it is a proving ground for studying the dynamics of a magnetically trapped atom coupled to an optical resonator. The atom chip presents the near-term possibility for creating an on-chip BEC [60] and magnetically transporting it into the mode of the fiber-gap cavity.

Several other atom chip and Fabry-Perot combinations have been proposed in recent years. At Caltech, we have investigated a system with an atom chip combined with countersunk supermirrors, though we tabled this experiment several years ago in favor of the photonic band gap (PBG) and microdisk experiments discussed in Chapter 6¹. Though the

¹We have recently learned that the group of V. Vuletić at MIT has built a similar device, but with cavity

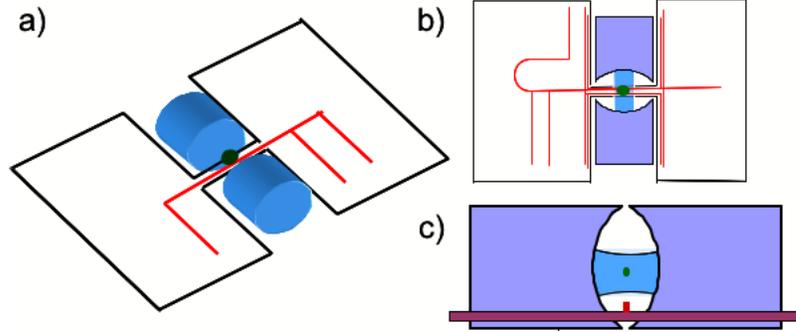


Figure 7.1: Fabry-Perot atom-cavity chip using commercial supermirrors. The minimum mirror spacing is roughly 1 mm and the radius of curvature is $R = 10$ cm, resulting in strong coupling parameters of $[m_0, N_0] = [7.2 \times 10^{-2}, 8.3 \times 10^{-3}]$ for $\mathcal{F} = 10^6$ and $[g_0, \kappa, \gamma_\perp]/2\pi = [6.8 \text{ MHz}, 75 \text{ kHz}, 2.6 \text{ MHz}]$.

mirrors in this proposal are large (7.5 mm diameter), they are high finesse ($\mathcal{F} = 10^5$ to 10^6) and easily obtained², which allows immediate incorporation into the experiment: no extra fabrication is necessary. The atom chip, made from sapphire, is cut in-house with a diamond blade post-microwire fabrication. Figure 7.1 is the schematic for this proposed experiment. The mode of the Fabry-Perot cavity is a distance four times the mode waist, $w_0 = 24 \mu\text{m}$, from the chip surface, which is far enough to not be clipped by the chip's surface. The atom chip bridge that supports the microwires must be no smaller than around 1 mm in width to maintain structural integrity. This forces the mirrors to be ~ 1 mm from each other. (With custom-cut mirrors like those used by Kimble's group, the mirrors could be spaced much closer.) This large cavity length severely constrains the minimum mode length, but with a cavity finesse of $\mathcal{F} = 10^6$, one could achieve strong coupling with $[g_0, \kappa, \gamma_\perp]/2\pi = [6.8 \text{ MHz}, 75 \text{ kHz}, 2.6 \text{ MHz}]$ and $[m_0, N_0] \approx [7.2 \times 10^{-2}, 8.3 \times 10^{-3}]$. This assumes a cavity of radius $R = 10$ cm. The microwire waveguides pictured in Figure 7.1 are simple one- and three-wire guides, but using double-layer atom chips would allow Z-traps or Libbrecht-style Ioffe traps to be fabricated on the bridge providing intracavity trapping in 3D and in the Lamb-Dicke regime (see Chapters 2 and 4).

Other groups have proposed using fiber pairs on atom chips to excite and detect the fluorescence of magnetically guided atoms [54]. The group of N. P. Bigelow has detected atoms with such a scheme, though not with single atom sensitivity [172]. E. A. Hinds' group and

mirrors separated by several centimeters.

²From REO, for instance: Research Electro-Optics, Inc., 1855 South 57th Court, Boulder, Colorado 80301.

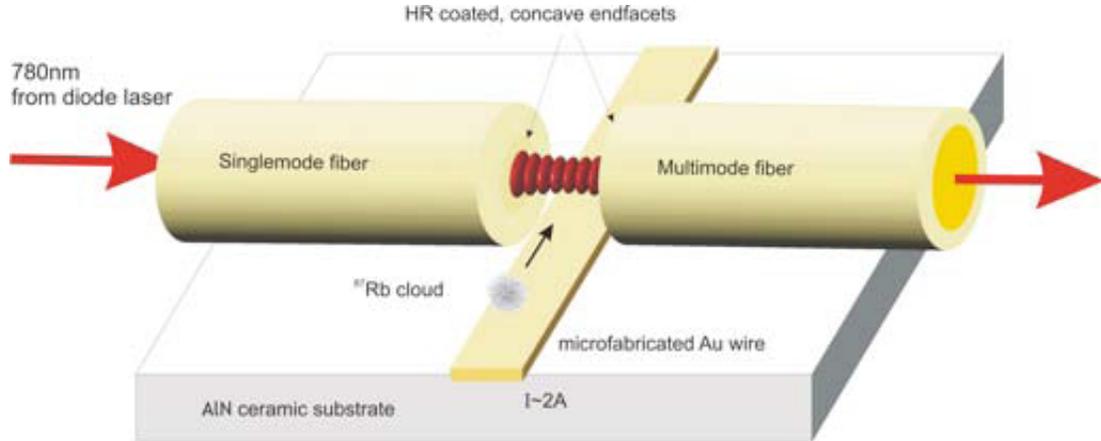


Figure 7.2: Schematic of the fiber-gap cavity atom chip experiment.

collaborators have recently built miniaturized Fabry-Perot cavities using microfabrication techniques [173].

The atom-cavity chip experiment discussed in this chapter was originally proposed in Reference [23]. A cavity is formed from opposing ends of optical fibers spaced $27 \mu\text{m}$ from one another. Onto each end of the optical fiber is glued a concave dielectric mirror of radius 1 mm. A finesse of $\mathcal{F} = 1050$ has been achieved, and the cavity QED parameters of such a device are $[g_0, \kappa, \gamma_{\perp}]/2\pi = [186, 2640, 3]$ MHz and $[m_0, N_0] = [1.0 \times 10^{-4}, 4.2 \times 10^{-1}]$ for ^{87}Rb atoms. This is in the “bad cavity” limit of cavity QED, $\kappa > g_0 > \gamma_{\perp}$ [163], but nevertheless single atom detection can still be performed as discussed in Section 7.3.

We have been able to magnetically guide and trap dilute clouds of cold atoms in the resonator mode and detect their presence via the modifications to cavity transmission due to atom-cavity coupling. We can detect the atom cloud undergoing a retroreflection from the magnetic waveguide terminus by observing—via cavity transmission—a double passage of the atoms through the cavity. Recently, we have obtained data suggesting progress towards the observance of optical bistability and single atom transits in this system.

7.2 Experimental details

This section is divided into two parts: operation of the atom chip waveguide; and the construction and operation of the fiber-gap cavity. Figure 7.2 provides an overview of the experiment.

7.2.1 Atom chip loading and waveguiding

The vacuum chamber—with base pressure $\sim 2 \times 10^{-10}$ Torr—is of the glued-cell design, and is discussed in detail in Chapter 10, Section 10.2 and in Reference [17]. The atom chip is first glued to a “base chip.” This base chip contains several isolated wire pads and a large H-wire—several hundred microns in wire width—that can carry several amps of current and form a U-MOT. The contact pads of the microwires on the atom chip are wire-bonded down to the base chip’s wire pads. The base chip is glued to the polished edges of a missing face of a cubic glass cell, thereby sealing this face with the chip. The atom chip is fully contained inside the glass cell, but the base chip’s wire pads extend outside the vacuum chamber defined by the edges of the glass cell. In this manner, a simple, compact electrical feedthrough is formed for all of the atom chip’s wires (see Figure 10.8 in the discussion in Section 10.2). A clip formed from a modified computer female PCI slot is secured onto the edge of the base chip, forming the connection between the wire pads and macroscopic wires. The layout of the base chip’s wire pads and the structure of the PCI clip are fashioned to prevent the obscuring of the MOT and detection laser beams. Figures 7.3 and 7.4 show the glass cell with the base chip/atom chip attached to one face of the vacuum cell. Both the atom and base chips are AlN substrates with microwires fabricated on the top surface using the electroplating method discussed in Chapter 4. The wires, made of gold, are $5 \mu\text{m}$ tall. The P-trap has a wire width of $100 \mu\text{m}$, and most of the waveguide has wire widths of $70 \mu\text{m}$ except the wires nearest to the resonator position which are $50 \mu\text{m}$ wide.

The cell walls can be antireflection-coated, but are not in the current experiment. A thermoelectric cooler (TEC), temperature sensor, and water-cooled copper block are attached to the top (air-side) of the base chip. This is for the regulation of the chip’s temperature—important for stabilizing the cavity resonance—and for allowing larger microwire currents by providing a good heat dissipation channel. The surface of the atom chip is coated with a dielectric that is transparent for much of the visible spectrum but is highly reflective at 780.2 nm , the wavelength used for trapping ^{87}Rb atoms in a mirror MOT. This coating is discussed in more detail in the next section. The fiber optics leading to the fiber-gap cavity are threaded through a notch in the glass cell between the base chip and the glass cell. Ultrahigh vacuum safe glue—Epotek 353—is used to seal this hole and to glue the cell to the chip as well.

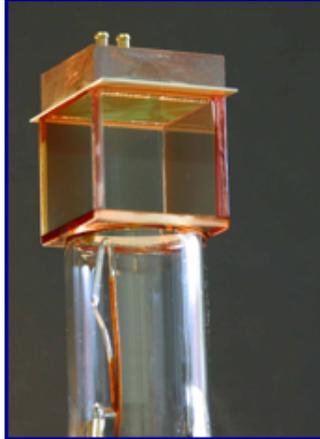


Figure 7.3: The vacuum cell assembly. The copper block at the top is a water-cooled heat sink. The glass cylinder is a glass-to-metal converter and is glued to the cubic glass cell. The Rb dispenser is located inside this cylinder. The base chip forms the top face of the cell and its wire pads extend outside of the vacuum to serve as the electrical feedthroughs. This picture is of the cell assembly used in the experiment discussed in Chapter 10, but it is similar that used in the fiber-gap cavity experiment discussed here.

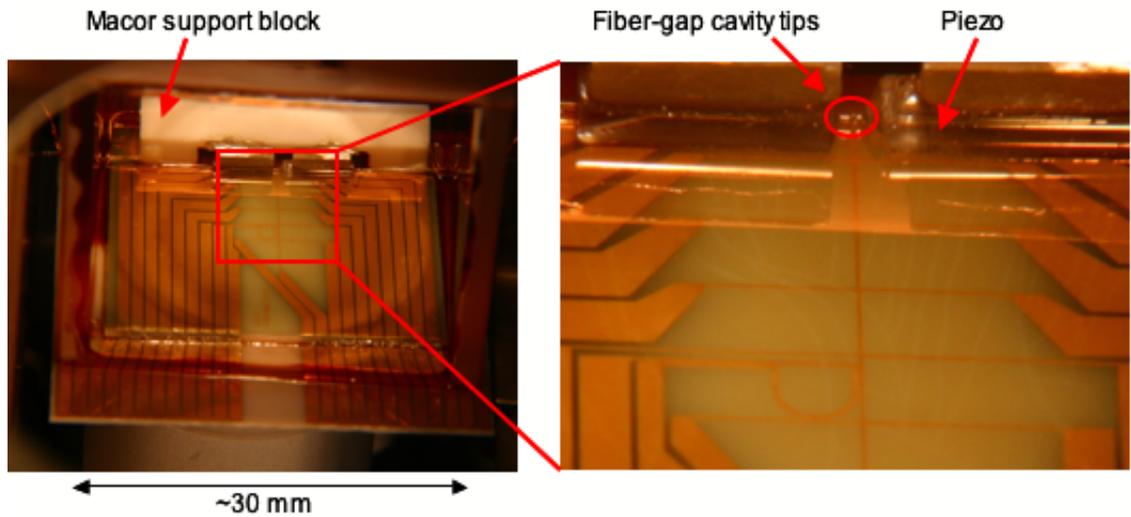


Figure 7.4: Fiber-gap Fabry-Perot cavity and atom chip assembly. Though transparent in visible band, the atom chip has a mirror dielectric coating on its surface for creating a mirror MOT at 780.2 nm.

The atom trapping, cooling, and microtrap loading proceeds in a similar manner as discussed in Chapter 2 for the Caltech-based experiments. Important differences are the use of Rb instead of Cs, which permits the relatively straightforward production of BECs; the use of anti-Helmholtz coils external to the vacuum chamber to form the initial MOT instead of a macro U-MOT; and the use of an alkali dispenser in place of an oven. The experiment begins with a mirror MOT loaded from a vapor of ^{87}Rb atoms produced by the combination of a dispenser run at low-current and the pulsed operation of a UV light source (see Chapter 10 for more information on this technique). The mirror MOT is formed by ~ 1 cm² beams and centered on the P-trap shown in Figure 7.5. The vacuum cell is surrounded by water-cooled coils, two pairs of coils per cardinal axis. One pair acts as both a permanent ambient magnetic field null and as a generator of an adjustable bias field. The second pair is in an anti-Helmholtz configuration for the production of a gradient field. These latter coils are seldom required in the experiment. The MOT's anti-Helmoltz coils are angled at 45° with respect to the chip surface to enable the formation of a mirror MOT. The external coils, despite the fact that they obscure optical access, are used instead of a macro U-MOT to maximize the number of atoms initially collected. Once collected in the mirror MOT for ~ 8 s, the atoms are transferred to the U-MOT formed using the base-chip's U-trap. The atoms are then loaded into the P-trap in a manner similar to that described in Chapter 2. Instead of a straight waveguide, the P-trap is used to deliver atoms to a 90° -rotated waveguide to enable the placement of the fiber-gap cavity assembly—which is several millimeters tall—away from the MOT trapping beams, thereby leaving them unobscured. Typically ~ 8 million ^{87}Rb atoms are initially loaded into the mirror MOT, and around 40% of these atoms are transferred all the way to the initial Z-trap of the waveguide in 200 ms. The lifetime of the Z-trap is ~ 8 s. This efficiency is dependent on both the transfer efficiencies of each step and the initial atom number, the latter due to the limited volume of the P-trap which cannot accommodate arbitrarily large numbers of atoms. Z-trap atom number optimization becomes a balancing game between obtaining a large initial trap population culled from a high vapor pressure of ^{87}Rb in the cell and attaining a long magnetic trap lifetime which is achieved by operating at low vapor pressure. The optimization difficulty lies in the fact that the optimal transfer parameters for each step—trap position, confinement, and hand-off time—are slightly dependent on the atom number at each step.

The atoms are ready to be guided to the fiber-gap cavity once they are confined in

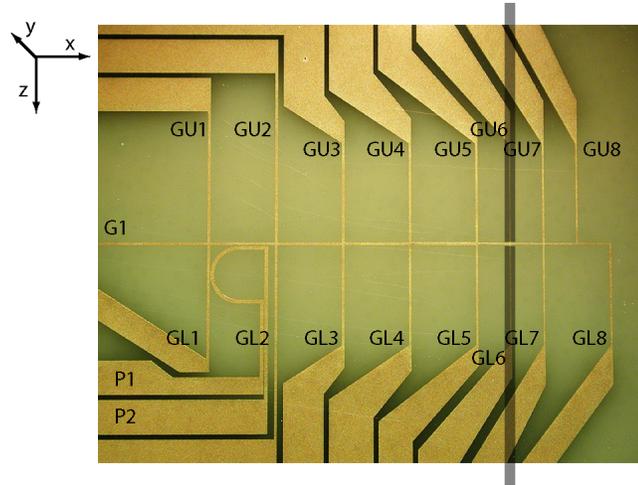


Figure 7.5: Close-up view of the atom chip. The wires are identified by “guide” (i.e. G1), “gate upper” (i.e. GU3), or “gate lower” (i.e. GL2). The P-trap wire is [P1,P2], where the brackets denote the beginning and end of a wire segment, respectively. The transparent grey vertical lines indicate the position of the fiber-gap cavity.

the initial Z-trap. Figure 7.5 shows a close-up of the microwires on the atom chip used in this experiment. The wire terminus labels (GU1, GL3, etc.) and the coordinate axes relate to the labels used in the Matlab script in Appendix B, Section B.3 that simulates the magnetic field, gradient, and curvature of the waveguide and the various traps. After the P-trap [+P1,-P2]³ transfer, the Z-trap may be formed by flowing current in wires [+GU1,-GL2] or [-GU2,+GL1] with a bias field in $-\hat{z}$. Typically, a wire current of ~ 2 A is used with a bias field of $B_{\hat{z}} \approx -16$ G resulting in a trap height of ~ 250 μm . The current and bias field are fine-tuned to match the trap height with the height of the resonator mode, 225-230 μm . For the case of [-GU2,+GL1], a weak bias field in $+\hat{x}$ should be added to enhance the trap curvature in \hat{x} : For this orientation of wire currents, a $+\hat{x}$ bias field adds to the field of the Z-trap’s side wires, while for fields $B_{\hat{x}} \leq 0$ (i.e. the field in $-\hat{x}$), two generally undesirable dips in the potential form at the intersection of the side wires and base wire. A field of $B_{\hat{x}} \approx 2$ G is sufficient for eliminating these dips.

In practice, an alternative Ioffe trap is used instead of the standard Z-trap: current is flowed through [+G1,-GU2] to form the base wire and one of the side wires, and current in [-GU1,+GL1] forms the other side wire. This allows easy atom injection into the waveguide by ramping-down the current in [+G1,-GU2] while simultaneously ramping-up the current

³The “+” in front of the wire terminus label designates the side with positive potential.

in [+G1,-GU8]. Notice that the same wire terminal is used, G1. This can be done by using two independently-controlled floating power supplies sharing the G1 connection but sinking current into different final wires. The speed of this wire exchange affects the length and velocity of the cloud of atoms propagating in the waveguide.

Before launching into the waveguide, the ^{87}Rb atoms are evaporatively cooled with an RF knife for 3 s [46]. To enhance cooling efficiency, the atoms are compressed in the initial Ioffe trap by increasing the bias field to 45 G while maintaining the [+G1,-GU2] current at 2 A. This increases the phase-space density of the cloud, resulting in a tighter packet of atoms propagating in the waveguide. Without this step, the atomic cloud is quite diffuse and only gradual changes in cavity transmission are observed when the cloud traverses the resonator mode. Greater than 10^4 atoms typically remain in the trap after the cooling stage. The RF-cooled atoms are decompressed into the original trap and the [+G1,-GU2] current is replaced with [+G1,-GU8]. The current in [-GU1,+GL1] may be increased to give the atom cloud an extra push. Depending on the exact manner of the trap release, the atoms travel down the waveguide at about 6 cm/s (8 mm in ~ 140 ms). Once the atoms have reached the resonator position—located above the [GU6,GL6] wire—the atoms may be confined in a H-trap using the wires [-GU5,+GL5] and [-GU7,+GL7] in addition to [+G1,-GU8] and the bias fields in $-\hat{z}$ and $+\hat{x}$. In addition to—or instead of—this H-trap, a dimple trap may be formed at the resonator. This is accomplished by running a current of approximately 1 A in [GU6,-GL6] and increasing $B_{\hat{x}}$ to 5 G. This bias field cancels the field from the wire, forming a dip in the potential directly above the wire at the position of the resonator mode. Sufficiently RF-cooled atoms will collect into this dimple trap. The atomic cloud’s position is adjusted in the resonator mode to maximize the atom-cavity coupling signal. This is done by making slight modifications to $B_{\hat{z}}$, which changes the trap’s height in \hat{y} , and by applying a non-zero $B_{\hat{y}}$, which moves the trap minimum in \hat{z} . Figure 7.6, generated by the Matlab code in Section B.3 of Appendix B, plots the magnetic field of the waveguide plus dimple.

7.2.2 The fiber-gap cavity

The cavity is formed by gluing concave dielectric films to the tips of fibers and mounting an opposing pair of fiber tips a few tens of microns from one another. Each fiber tip is glued to a shear-mode piezo with the tip slightly cantilevered from the piezo end (see Figure 7.7

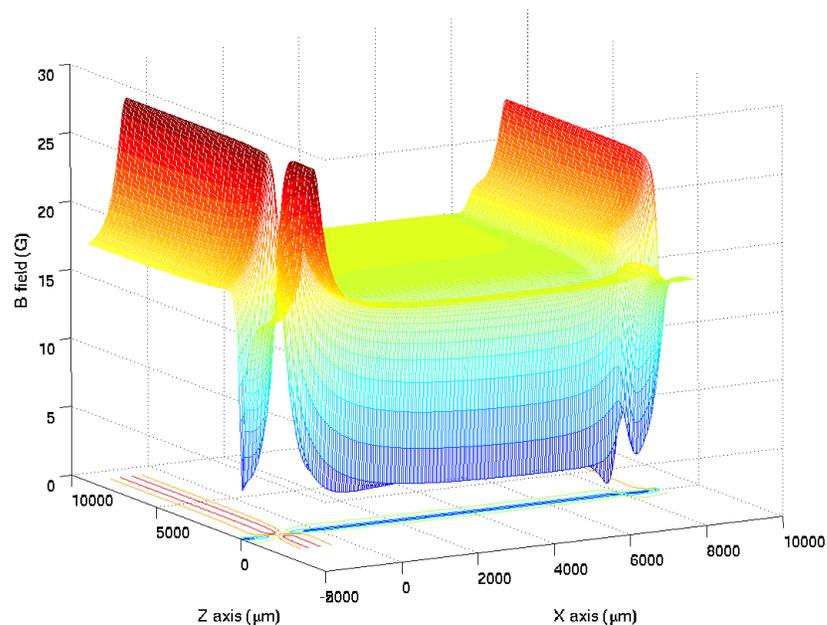


Figure 7.6: Magnetic field plot of a waveguide plus dimple (located at $x \approx 7750 \mu\text{m}$) generated by the Matlab script in Section B.3 of Appendix B. The wire currents and bias fields are as follows: [+G1,-GL8] has 2 A, [+GU1,-GL1] has 2.5 A, [-GU6,+GL6] has 1 A, $B_{\hat{x}} = -5 \text{ G}$, $B_{\hat{y}} = -1.8 \text{ G}$, $B_{\hat{z}} = -16.3 \text{ G}$, and $y_{\text{min}} = 245 \mu\text{m}$. Note, there are some differences between this field and current configuration and the one described in the text: the signs of some of the currents and fields are flipped and the guide wire extends to GL8 instead of GU8. The trapping principle remains the same.

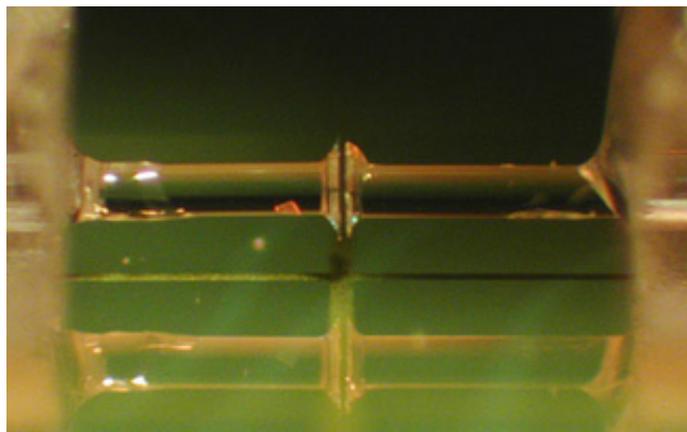


Figure 7.7: Image of the fiber-gap Fabry-Perot cavity. The mirror spacing is $27 \mu\text{m}$.

for a image of an actual cavity formed in this manner). The fibers are affixed to the piezos using UV curable glue (purchased from Dymax) which can be set with low shrinkage in less than 3 s. The two piezos are glued to a small Macor block that is itself glued to the atom chip. Cavity alignment is performed in the following manner: one fiber tip is initially glued to a piezo, while the other is held in a 5-axis translation and rotation stage. The stage is adjusted until the TEM₀₀ cavity resonance is optimized, and at this moment the second fiber tip is UV glued and the stage detached. The transmission is monitored during the cure to ensure that proper alignment is maintained throughout. Figure 7.8 illustrates the procedure for attaching the cavity mirrors to the fiber tips. The concave mirror is fabricated by coating a 1 mm radius hemispherical glass ball lens with a detachable dielectric coating: the layer between the mirror coatings and the glass is sacrificial which allows the mirror to “peel-away” from the glass substrate⁴. The fiber tip is glued to the apex of the coated ball lens with a drop of NIR transparent UV-curable glue as shown in panel (b) of Figure 7.8. After curing, the fiber is tugged upwards, separating the coating from the glass. Finally, the excess coating is broken-away leaving a small mirror coating glued to the end of the fiber tip as shown in panel (c) of Figure 7.8 and in Figure 7.7. The finesse of the cavity formed with these mirror coatings is $\mathcal{F} = 1050$. The cavity length is 27 μm and the radius of curvature is approximately the same as the ball lens, 1 mm.

Figure 7.9 is a schematic of the fiber-gap cavity laser coupling, detection, and locking set-up. A beam from a laser diode—locked to the 780.2 nm resonance of the D2 line in ⁸⁷Rb—is coupled into the single-mode fiber that forms the input to the fiber-gap cavity. Before the coupling, this free-space beam passes through $\lambda/2$ and $\lambda/4$ waveplates to control the intracavity polarization. Light reflected by the cavity is picked-off by a non-polarizing beamsplitter and directed into a Thorlabs photodetector. It is very important that angle-cleaved fibers are used. This eliminates oscillations in signal intensity that are caused by the formation of an etalon between the cavity input mirror and the fiber input facet. The fiber used for the output of the cavity is multimode to maximize the collection efficiency of the transmitted intracavity light. The transmitted light is out-coupled from the fiber to allow the interjection of a mechanical shutter before re-coupling the light into the fiber input of the APD detector. This mechanical shutter blocks the high-power cavity lock light that

⁴This coating is produced by the German company OIB (Optical Interference Components) <http://www.oib-jena.de/firmenpreng.html>

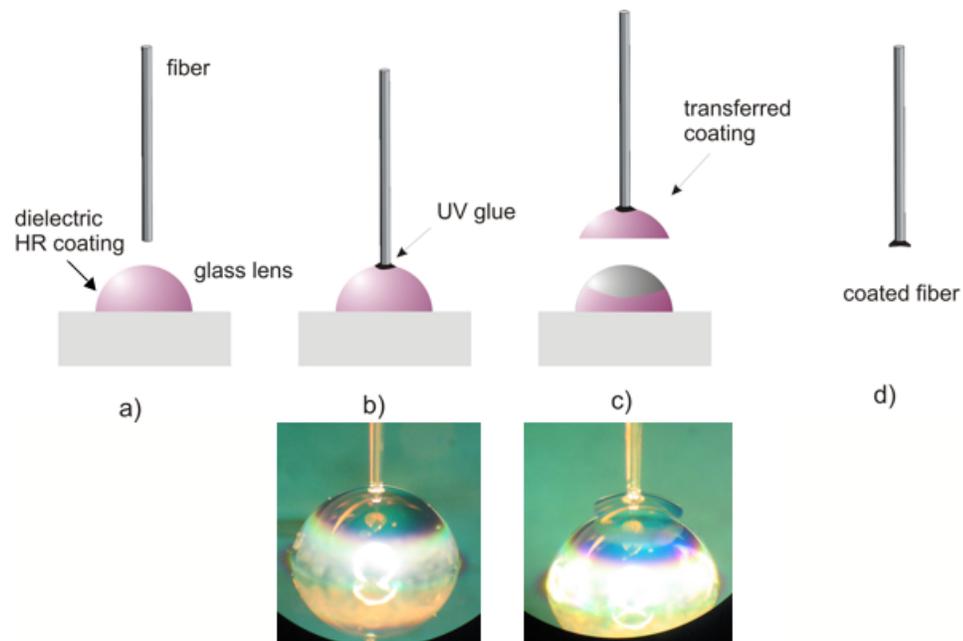


Figure 7.8: a) The fiber is carefully lowered to the top of the coated ball lens. b) UV-curable glue attaches the fiber tip to the dielectric mirror coating. c) After curing, the fiber is tugged upwards, detaching the coating from the glass lens. d) Excess dielectric film is flaked-off, leaving behind the small mirror coating attached to the fiber tip.

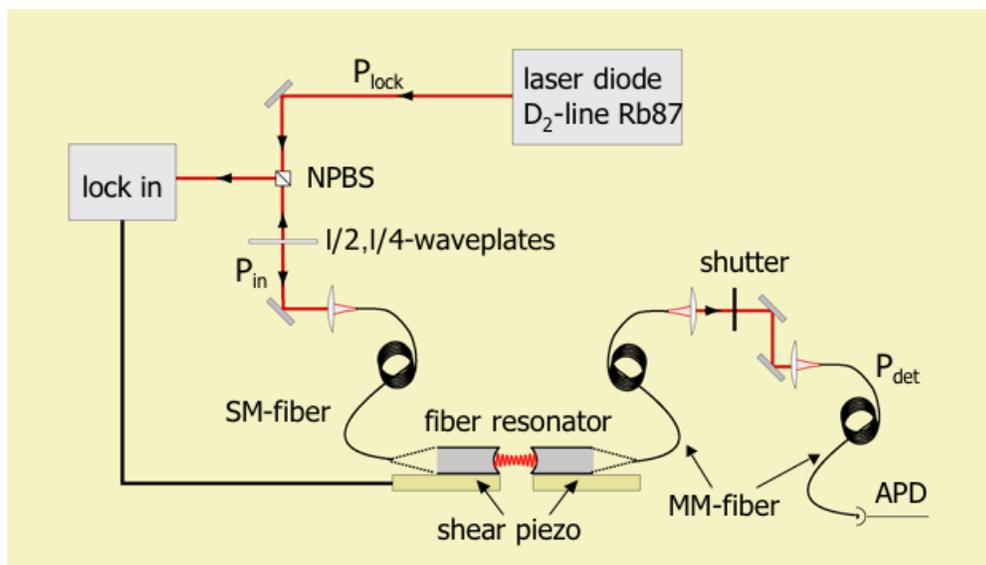


Figure 7.9: Schematic of the fiber-gap cavity set-up.

would damage the APD, but opens for the low-power probe light during the atom-cavity coupling phase of the experiment. The Si APD is from PerkinElmer and has a minimum dark count rate of 157/s and saturates at a photon flux of approximately 10^7 /s.

The cavity is tuned by applying a DC voltage to one of the two cavity piezos. The other piezo is driven with an AC voltage of frequency equal to ~ 11 kHz, which is limited by the piezo resonance. This adds a frequency dither to the reflected light that is used to produce an error signal with the aid of a lock-in detector. The error signal, suitably amplified, is fed back to the second piezo to maintain cavity lock. The lock bandwidth is a few hundred Hertz. The large linewidth of the cavity (> 2.6 GHz) and the use of piezos force this lock to be in the less-desirable slow modulation regime as discussed by E. Black in his tutorial on the Pound-Drever-Hall technique [174] (see also Reference [175]). Error signal detection in reflection is sub-optimal, but eliminates the need for reducing atom detection signal-to-noise by syphoning-off photons from the transmission signal. The reflected signal rides atop a large DC offset due to the superposition of rejected light from the free-space to fiber and fiber tip to cavity interfaces. We do not have separate laser sources for the cavity lock and probe beams. Consequently, the cavity lock must be turned-off when the beam is attenuated during the atom transit detection sequences. The optimal driving power for maximizing signal-to-noise is around 20 pW, and the max cavity lock power is approximately $10 \mu\text{W}$. The AOM that controls the probe laser detuning has an attenuation range of > -80 dB and is used to shift the drive power from probe mode to lock mode. Despite the lock non-optimality, the cavity is quite stable and the lock can be maintained for long periods of time. This is greatly aided by the containment of the cavity in a vacuum chamber and the large linewidth of the cavity which makes it less susceptible to mechanical vibrations.

Near-resonant stray light from the high lock power in the resonator quenches the lifetime of the atoms in the magnetic trap. To avoid this, we shut-off the cavity lock and attenuate the cavity probe to ≤ 0.2 pW before the RF-cooling stage. This occurs roughly 3 s before the atom transits, and the cavity length is free to wander off resonance in this time period. The TEC temperature stabilizes the chip to within 0.01 °C, aiding the passive stability of the cavity length. However this feedback loop is very slow, and cannot compensate for the transient heating caused by the microwire power dissipation. Within a few tens of milliseconds after the turn-on of the microwire current, we can observe a shift in the cavity resonance, and during a typical atom chip wire sequence the resonance can shift

by many linewidths. This temperature shift is mitigated by only gluing one end of the Macor block to the chip which reduces its thermal coupling. Nevertheless, a feedforward procedure is implemented to ensure that the cavity is on resonance as the atoms transit the resonator mode. The extent to which the cavity is detuned for a specific microwire sequence is measured and subsequently negated by applying an appropriate feedforward voltage to the piezo. This procedure is of course cumbersome since it needs to be re-calibrated each time the microwire sequence changes to account for the differing power dissipation. Future experiments should use a secondary laser resonant with a far-detuned cavity resonance to lock the cavity. This beam shouldn't affect the atom motion if suitably far-detuned and of low power. A dichroic mirror can pick-off this light from the probe beam.

7.3 Signal-to-noise and spontaneous emission analysis using the master equation

In a manner similar to that of Section 6.4.1 of Chapter 6, we analyze the fiber-gap Fabry-Perot cavity's capability for detecting single atoms. Refer to that section for definitions of signal, noise, detunings, etc. As mentioned earlier, the present cavity has $\mathcal{F} = 1050$, $L = 27 \mu\text{m}$, and $R = 1 \text{ mm}$. The cavity waist is $w_0 = 5.3 \mu\text{m}$. A previous incarnation of this experiment used a cavity of $\mathcal{F} = 600$, $L = 40 \mu\text{m}$, and $R = 1 \text{ mm}$, but in this section we will focus on the current, higher finesse system.

Table 6.1 includes the cavity QED parameters for the fiber-gap cavity as well as comparisons to other resonator systems used for cavity QED. While worse than most current resonator systems, the fiber-gap cavity is relatively compact, easy to build, and integrates well with magnetic microtraps. The cavity QED parameters for this system are $[g_0, \kappa, \gamma_\perp]/2\pi = [186, 2640, 3] \text{ MHz}$ and $[m_0, N_0] = [1.0 \times 10^{-4}, 4.2 \times 10^{-1}]$. This current implementation of the fiber-gap cavity is in the so-called "bad cavity" limit of cavity QED, $\kappa > g_0 > \gamma_\perp$ [163], but as we show in the following, single atom detection is still possible. A major difficulty, however, lies in the inability of this cavity to suppress the spontaneous emission of photons during a probe period long enough to detect a single atom. This hampers our ability to maintain the atom in a magnetic trap post-detection.

As a major goal of this experiment is to detect single atoms, the analysis presented here is for the case of a single intracavity atom. The Matlab script used in these simulations—

presented in Appendix B, Sections B.1 and B.2—can be modified to simulate the effects of more than one intracavity atom. Simulations in the multi-atom regime are important to enable the exact calibration of the intracavity atom number. From atom cloud density measurements at the position of the resonator, we can determine that there are ≤ 10 atoms in the resonator for most of the data presented in Section 7.4. However, at present we cannot determine the atom number more precisely. We will return to this problem in Section 7.4.

Cavity QED with multiple atoms and a single atom are similar to one another in the regime of weak driving. The first rung on the Jaynes-Cummings ladder (Figure 6.13) is unmodified for N atoms⁵ if one replaces the single atom-cavity coupling rate, g_0 , with an effective coupling rate $g_{eff} = \sqrt{N}g_0$. Here, g_0 is the atom-cavity coupling rate for each individual atom and may not necessarily be of equal magnitude to the g_0 in the single atom case. For higher rungs on the Jaynes-Cummings ladder the energy eigenvalues differ. For instance, the second rung of excitations for the single atom case is divided into two levels spaced by $2\sqrt{2}\hbar g_0$, while for the N atom case there are three levels, with the upper and lower level split from the middle by $\pm\sqrt{4N-2}\hbar g_0$. For $N \gg 1$ these levels are split by $\pm 2\hbar g_{eff}$. As noted in Reference [153], the eigenvalues for the second excitation in the $N \gg 1$ case approximate that of classically coupled oscillators: it is important to note that the semi-classical equation 6.5 can accurately describe the system of $N \gg 1$ for any drive strength even when g_{eff} is much larger than κ and γ_{\perp} , but this equation fails when $N=1$ for a similarly strongly coupled system. This highlights the important role that quantum fluctuations play in this system: the single atom case produces qualitatively different dynamics that are only properly described by the quantum master equation 6.6. Note that the multi-atom g_0 may be smaller than the single atom g_0 , but the total multi-atom coupling g_{eff} may be equal to the single atom coupling for the $N = 1$ case. In other words, even though the $N \gg 1$ atoms may be each individually weakly coupled but collectively strongly coupled to the cavity mode, the $N \gg 1$ system dynamics are not equivalent to that of a strongly coupled single atom.

⁵Please excuse the redundant notation with Equation 6.9. The meaning of N will be clear from the context.

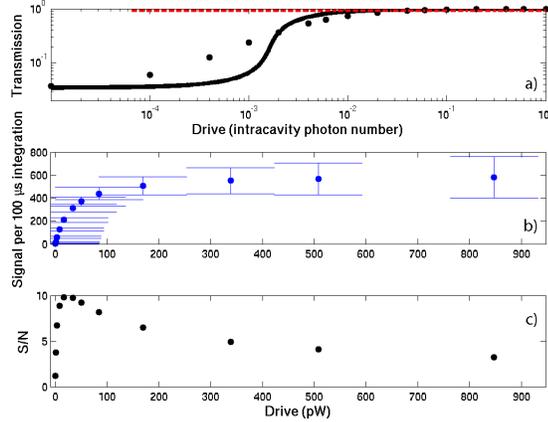


Figure 7.10: a) The transmission of the fiber-gap Fabry-Perot cavity as a function of drive strength—measured in intracavity photon numbers for a resonant and empty cavity—calculated from Equations 6.5 (black line) and 6.6 (black dots). The empty cavity transmission is shown as a dashed red line. b) The expected signal S as a function of drive power—quoted as $\hbar\omega$ times the resonant, empty intracavity photon number—during an atom transit with a $\Delta t = 100 \mu\text{s}$ integration time. Error bars correspond to S_{noise} . The photon detection efficiency is $\eta = 0.05$. c) The expected signal-to-noise, S/\mathcal{N} , in a $100 \mu\text{s}$ integration time. For this plot, $[g_0, \kappa]/2\pi = [186, 2640]$ MHz and $[\Delta, \theta] = [0, 0]$.

7.3.1 Signal-to-noise and spontaneous emission

We now turn to the determination of this fiber-gap cavity’s ability to detect single atoms. The atom-cavity system is most easily run on-resonance due to the lack of extra EOM and AOMs in the laser system. (Though the cavity could be detuned from probe laser resonance $\theta \neq 0$ by suitably offsetting the feedforward control.) Figure 7.10 shows the cavity transmission, signal S , and signal-to-noise S/\mathcal{N} versus drive for $[\Delta, \theta]/2\pi = [0, 0]$. This is plotted for a detection integration time $\Delta t = 100 \mu\text{s}$, which is typical for this experiment, and a detection efficiency of $\eta = 0.05$. This detection efficiency incorporates three mechanisms for photon loss: absorption and scattering of photons by the cavity mirrors, $\eta_1 = 0.1$; and photon loss during propagation from the output cavity mirror to the APD, η_2 , and the APD quantum efficiency itself, η_3 . In this experiment $\eta_2\eta_3 = 0.5$. The full detection efficiency, $\eta = \eta_1\eta_2\eta_3$, may be modeled as a beam splitter placed between a zero-loss cavity and a perfect detection channel. The S/\mathcal{N} scales as $\sqrt{\eta\Delta t}$.

The maximum signal-to-noise, $S/\mathcal{N} = 9.8$, is obtained at a drive power of 17 pW and an empty, resonant intracavity number of 2×10^{-3} . Single atom detection seems to be feasible with the current system. However, the number of spontaneously emitted photons

during the $\Delta t = 100 \mu$ detection is quite large—657—and poses a severe problem for non-destructively detecting the atom in a magnetic trap. As mentioned in Chapter 6, both the heating and the optical pumping to untrapped states due to the spontaneous emissions will limit the lifetime of the magnetic trap, thereby reducing the coupling time of the atom-cavity system. Each spontaneously emitted photon has a probability—suppressed by the Lamb-Dicke parameter if the atom is confined in this regime—to excite the atom out of the trap’s ground state and into ever higher vibrational levels. To estimate how severe this will be for our system, one needs to calculate the average N_{SponE} —and average time—it takes for the atom to be lost either by being excited to an energy larger than the trap’s depth or excited to such a spatial extent that the atom has a chance to encounter the fiber-gap cavity mirrors.

The fraction of the signal, \mathcal{S} , that is due to the spontaneous emission of photons out of the side of the cavity mode versus the suppression of transmission due to atom-cavity coupling is equal to 15.7% for these cavity QED parameters of $[g_0, \kappa]/2\pi = [186, 2640]$ MHz. While the spontaneous emission does not comprise the majority of the missing photons constituting the signal, it is still quite high. Even if we fix the \mathcal{S}/\mathcal{N} to equal a minimum of 3, 61 photons are spontaneously emitted in a time $\Delta t = 9.3 \mu s$. Heating and optical pumping due to spontaneous emission seems unavoidable with a fiber-gap cavity of this volume and $\mathcal{F} = 1050$. Setting the detuning parameters, $[\Delta, \theta]$, to be non-zero does not seem to help much. Even with $\eta = 1$, a reasonably thorough parameter search—which is difficult due to the need for optimizing over $\Delta, \theta, \Delta t$, and drive power—only revealed a narrow detuning set in which the \mathcal{S}/\mathcal{N} was ~ 2 while maintaining $N_{SponE} \leq 1$. The irrepressible N_{SponE} in this system is caused by the fact that $g_0 \ll \kappa$, resulting in a $N_0 = 4.2 \times 10^{-1}$ that is not quite small enough to ensure low intracavity fields at high \mathcal{S}/\mathcal{N} . In more realistic situations where $\eta < 1$, N_{SponE} is unaffected while \mathcal{S}/\mathcal{N} is suppressed—a combination which is detrimental to our goal of non-destructive single atom detection.

Figure 7.11 shows the transmission for various scanned detunings at different drive powers. Significant \mathcal{S}/\mathcal{N} can be achieved most simply with $[\Delta, \theta] = 0$, but detuning the cavity from the laser and atom also produces $\mathcal{S}/\mathcal{N} > 1$. Panels (a) through (c) show the most straightforward detuning schemes, but others are possible. In particular, detunings of $[\Delta, \theta]/2\pi \approx [\pm 0.8, 3]$ GHz reveal a sweet-spot in which the $\mathcal{S}/\mathcal{N} \approx 2$ while the $N_{SponE} \leq 1$. This is shown in Figure 7.12 for $\Delta t = 100 \mu s$ and drive of $E = 0.01$ intracavity photons

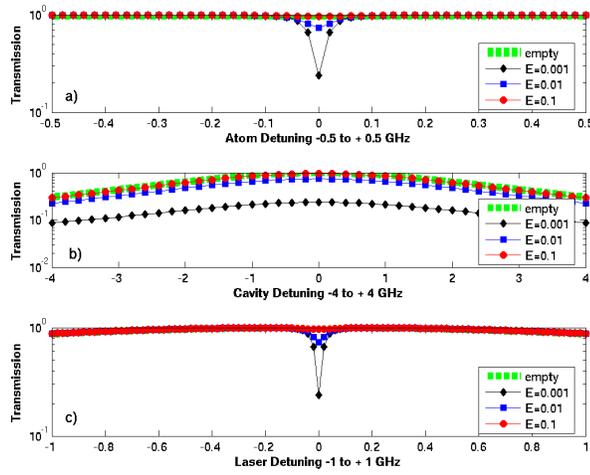


Figure 7.11: Fiber-gap cavity transmission as a function of a) atom detuning [$\theta = 0$], b) cavity detuning [$\Delta = 0$], and c) laser detuning [$\Delta = \theta$] for one intracavity atom and $[g_0, \kappa]/2\pi = [186, 2640]$ GHz. The drive (measured as the empty cavity intracavity photon number on resonance) is $E = 0.001$ in the black \diamond curve, $E = 0.01$ in the blue \square curve, and $E = 0.1$ in the red \circ curve. The green dashed curve is the empty cavity transmission. All curves are normalized to the empty cavity transmission at each curve's drive strength.

(referenced to an empty, on-resonant cavity). However, the use of $\eta = 1$ for this plot is unrealistically high. Adjusting the drive and Δt does not seem to help to increase the \mathcal{S}/\mathcal{N} where N_{SponE} crosses unity. For comparison to other choices of η and Δt , note that \mathcal{S}/\mathcal{N} scales as $\sqrt{\eta\Delta t}$ (see Equations 6.11 and 6.12) whereas N_{SponE} scales as Δt . Figure 7.13 repeats this \mathcal{S}/\mathcal{N} versus N_{SponE} analysis for the detunings used in Figure 7.11 and a drive of $E = 0.01$ intracavity photons. In each of the three cases, $\Delta t = 100 \mu\text{s}$ and $\eta = 1$, and unfortunately in none of these cases is it possible to adjust Δt to obtain a $\mathcal{S}/\mathcal{N} > 1$ while maintaining $N_{SponE} \leq 1$.

With the same cavity length and mirror radius of curvature ($[L, R] = [27 \mu\text{m}, 1 \text{mm}]$), one would need a cavity of finesse $\mathcal{F} = 20000$, given $\eta = 0.25$ and $\Delta t = 0.09 \mu\text{s}$, to have a $\mathcal{S}/\mathcal{N} = 3$ while spontaneously emitting less than one photon per Δt . The next generation of this experiment, to be built in Prof. J. Reichel's new ENS labs in Paris, will use a different technique to create the fiber-gap cavity mirrors that will greatly increase \mathcal{F} while also decreasing V_m . The new scheme uses a focused and pulsed CO_2 laser of carefully controlled intensity to melt the core of the fiber at its tip [176]. A dimple in the core is formed with radius of curvature equal to $200 \mu\text{m}$. Coating companies have estimated that

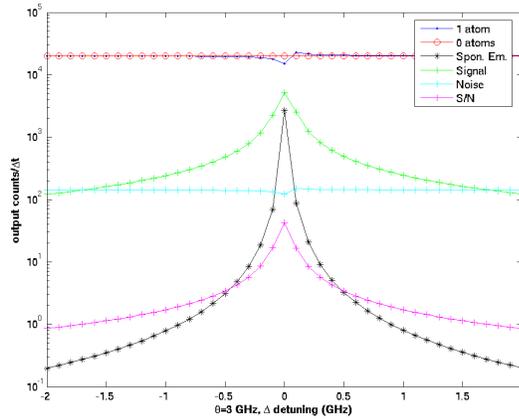


Figure 7.12: Fiber-gap cavity transmission with and without an intracavity atom along with plots of the signal, S ; noise, S_{noise} ; number of spontaneously emitted photons, N_{SponE} ; and S/\mathcal{N} . These are plotted for $\eta = 1$, $\Delta t = 100 \mu s$, and a drive $E = 0.01$. The laser-cavity detuning θ is fixed at 3 GHz, but the atom-cavity detuning varies between ± 2 GHz. The plots are for one intracavity atom and $[g_0, \kappa]/2\pi = [186, 2640]$ GHz. The y-axis for the panels lists the number of photon counts per Δt except for the magenta (Δ) curve for which the y-axis denotes the magnitude of S/\mathcal{N} . The $S/\mathcal{N} \approx 2$ while $N_{SponE} \approx 1$ at detunings of $[\Delta, \theta]/2\pi \approx [\pm 0.8, 3]$ GHz.

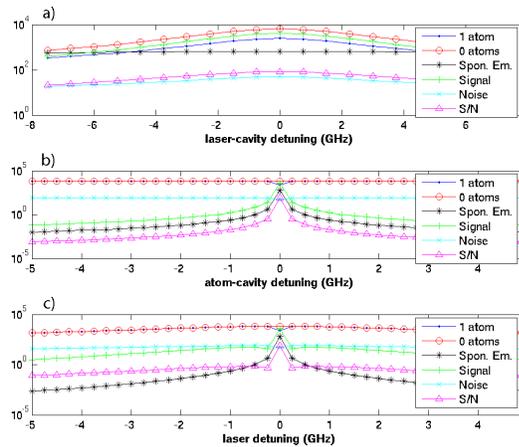


Figure 7.13: Fiber-gap cavity transmission with and without an intracavity atom, and plots of the signal, S ; noise, S_{noise} ; number of spontaneously emitted photons, N_{SponE} ; and S/\mathcal{N} . These are plotted for $\eta = 1$, $\Delta t = 100 \mu s$, and a drive $E = 0.01$. The three panels are plots as a function of a) θ , b) Δ , and c) laser detuning, $[\Delta = \theta]$, for one intracavity atom and $[g_0, \kappa]/2\pi = [186, 2640]$ GHz. The y-axis for the panels list the number of photon counts per Δt except for the magenta (Δ) curves in which the y-axis denotes the magnitude of S/\mathcal{N} .

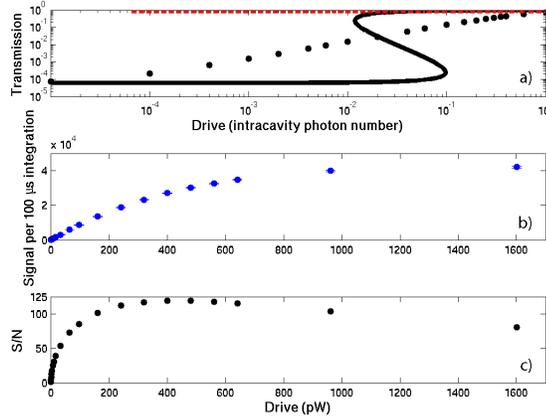


Figure 7.14: Same as for Figure 7.10 except $[g_0, \kappa]/2\pi = [423, 500]$ MHz and $[\mathcal{F}, L, R] = [10^4, 15 \mu\text{m}, 200 \mu\text{m}]$. $[\Delta, \theta] = [0, 0]$.

the smoothness of this dipole should readily allow a mirror finesse of $\mathcal{F} = 10^4$. Unlike the current system, the new mirror does not extend past the tip of the fiber, and this allows the fiber gap to be $15 \mu\text{m}$ or smaller. Taking cavity parameters of $[\mathcal{F}, L, R] = [10^4, 15 \mu\text{m}, 200 \mu\text{m}]$, the achievable cavity QED parameters are $[g_0, \kappa]/2\pi = [423, 500]$ MHz and $[m_0, N_0] = [1.9 \times 10^{-5}, 1.5 \times 10^{-2}]$. These are the values listed in Table 6.1 for the projected performance of the fiber-gap Fabry-Perot cavity system. While not quite in the “good cavity” limit of cavity QED, this improved cavity system allows single atom detection with much larger signal-to-noise, $\mathcal{S}/\mathcal{N} \approx 120$. Figure 7.14 shows the cavity transmission, S , and \mathcal{S}/\mathcal{N} versus drive power for zero detunings.

The maximum \mathcal{S}/\mathcal{N} occurs at a drive of $E = 0.25$ (400 pW). A detection time of $\Delta t = 100 \mu\text{s}$ is used. For these mirror coatings, an effort will be made to increase η_1 to 0.5, which combined with a pessimistic estimate of $\eta_2\eta_3 = 0.5$ gives $\eta = 0.25$. Spontaneous emission comprises only 0.87% of this cavity’s on-resonance S , which is 940 photons for $\Delta t = 100 \mu\text{s}$ and $E = 0.25$. Because the on-resonant \mathcal{S}/\mathcal{N} is so much larger for this cavity, we can decrease the drive power to $E = 2 \times 10^{-4}$ (0.3 pW) and have a $N_{SponE} \leq 1$ while still obtaining a $\mathcal{S}/\mathcal{N} = 5.6$ with $\Delta t = 100 \mu\text{s}$. Better performance might be achieved with non-zero detunings. We can also achieve $N_{SponE} \leq 1$ by driving at $E = 0.25$ (400 pW) but decreasing the detector integration time to $\Delta t = 0.06 \mu\text{s}$. With a cavity of $[L, R] = [15, 200] \mu\text{m}$ a finesse of only $\mathcal{F} = 6000$ is required to have $N_{SponE} \leq 1$ while maintaining a $\mathcal{S}/\mathcal{N} = 3$ ($\Delta t = 0.12 \mu\text{s}$). This sets an approximate lower bound on \mathcal{F} that we need to surpass in

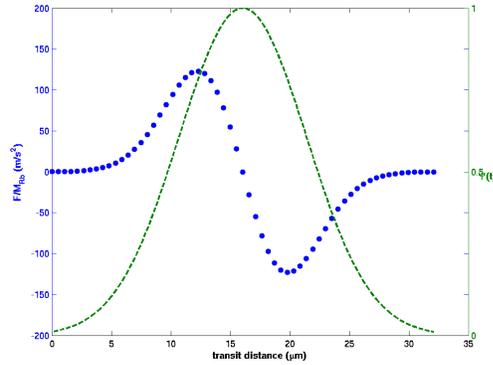


Figure 7.15: The force on an atom traversing the fiber-gap cavity mode. $[g_0, \kappa]/2\pi = [186, 2640]$ MHz, $[\Delta, \theta]/2\pi = [-800, 3000]$, and the drive is $E = 0.002$.

other to achieve non-destructive single atom detection with this length fiber-gap Fabry-Perot cavity.

7.3.2 Simulated atom transits and cavity induced force

Figure 7.15 shows the force on a ^{87}Rb atom dragged at $v = 6$ cm/s through the $\mathcal{F} = 1050$, $[L, R, w_0] = [27, 1000, 5.3]$ μm cavity (see Equation 6.13). The detuning is $[\Delta, \theta]/2\pi = [-800, 3000]$ and the drive is $E = 0.002$. The maximum force on the atom is $F/m_{\text{Rb}} = 123$ m/s^2 , which imparts a change in velocity of $\Delta v \approx 1$ to 2 cm/s on the atom. This is comparable to the atom's incoming velocity, and the cavity field at this detuning should noticeably affect the atom's trajectory.

Figure 7.16 presents a simulation of the photon flux transmitted by the cavity as an atom passes through the resonator mode in a trajectory perpendicular to the cavity axis and in the center of an anti-node. The atom's velocity is 6 cm/s. Panel (a) shows the signal with a detection integration time of $\Delta t = 100$ μs and drive $E = 0.002$ (17 pW), and Panel (b) is the same but with $\Delta = 10$ μs . As discussed in Section 6.4.2 of Chapter 6, the noise, ΔN , of this signal is found by solving for ρ_{ss} for each $g(t) = g_0\psi(t)$ in time steps of Δt , chosen to simulate a finite bandwidth photodetector. Each point includes additional shot-noise selected randomly from a normal distribution of standard deviation ΔN . Again, we make the caveat that simulation is not derived from a quantum trajectory calculated from the conditional master equation [25], but is simply calculated using ρ_{ss} from the unconditional equation 6.6.

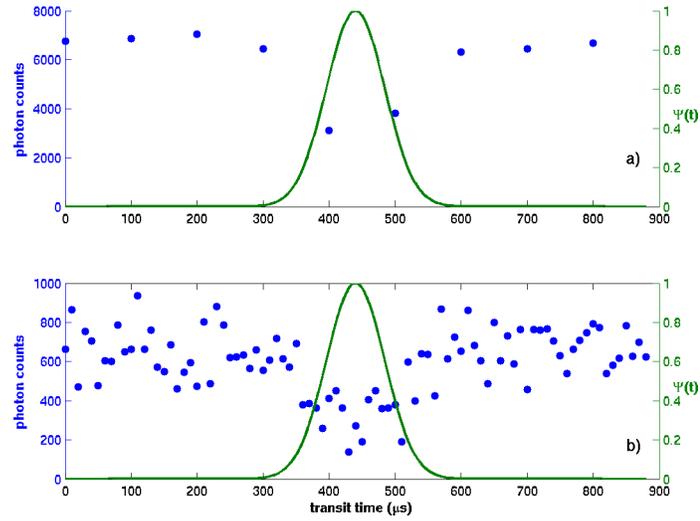


Figure 7.16: Simulated photon counts due to an atom transit through an anti-node and perpendicular to the fiber-gap cavity axis. For this cavity, $[g_0, \kappa]/2\pi = [186, 2640]$ MHz. Blue dots (left axis) are the photon counts, and the green, solid curve (right axis) is the Gaussian variation of $g(t)/g_0 = \psi(t)$ experienced by the atom during its transit. The calculation is for detunings of $[\Delta, \theta] = [0, 0]$, and a drive power $E = 0.002$ (17 pW). a) $\Delta t = 100 \mu\text{s}$ and the maximum $\mathcal{S}/\mathcal{N} = 44$. b) $\Delta t = 10 \mu\text{s}$ and the maximum $\mathcal{S}/\mathcal{N} = 14$. In both of these plots $\eta = 1$.

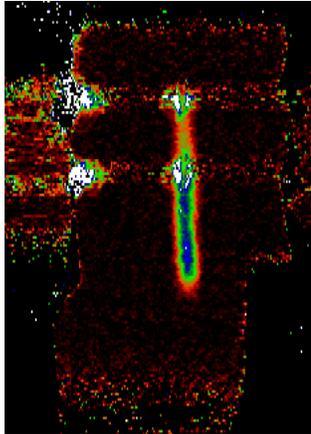


Figure 7.17: Absorption image of waveguided ^{87}Rb atoms as they approach the fiber-gap cavity mode. The image is doubled due to the absorption laser beam being reflected from the atom chip’s mirrored surface. The atoms are the bright, vertical cigar-shaped region in the middle; the cavity position is at the upper tip of the cigar; and the fiber tips are the faint (red) horizontal regions intersecting the atom cloud.

7.4 Experimental results

This section presents some of our preliminary experimental results. Figure 7.17 is an absorption image of the atoms—post-RF cooling—being magnetically guided toward the fiber-gap cavity mode. The first detection of atoms with a fiber-gap cavity was made with the $\mathcal{F} = 600$, $L = 40 \mu\text{m}$, $R = 1 \text{ mm}$ cavity. The cavity QED and critical coupling parameters for this cavity were $[g_0, \kappa]/2\pi = [139, 3125] \text{ MHz}$ and $[m_0, N_0] = [2.3 \times 10^{-4}, 9.7 \times 10^{-1}]$. Unfortunately, during the course of these initial experiments, the vacuum cell developed a leak from an unfixable crack. The atom chip and cavity were replaced, this time with the better cavity of $\mathcal{F} = 1050$, $L = 27 \mu\text{m}$, and $R = 1 \text{ mm}$ that was the focus of much of the analysis in Section 7.3. However, the first measurements were performed with the $\mathcal{F} = 600$ cavity and the first atom transit detection is shown in Figure 7.18. This signal had to be averaged over many trials (~ 20) since we had yet to RF-cool the atoms and fully optimize the position of the waveguided atoms as they pass through the resonator mode. Nevertheless, we were able to make transmission versus drive power curves as shown in Figure 7.19 (b). In this figure, the signal is averaged 10 times and the cavity and laser are on-resonance with the atom transition. Panel (a) shows the atom cloud transit signals for four different drive powers: The bottom black curve has the lowest drive power and the top blue curve has the largest. The variance in transmission is presumably due to a

varying intracavity atom number as the atom cloud of variable density passes through the resonator mode. By taking different time slices of these transmission curves, we can form plots such as those in panel (b). Each of the three curves is taken at time slices of relatively constant cavity transmission, and we therefore expect each curve to represent a different but approximately constant intracavity atom number, N . Since $g_{eff} = \sqrt{N}g_0$, the black (\square) curve should represent a system with stronger atom(s)-cavity coupling than the red (\circ) and green (\triangle) curves which are taken at later times and consequently lower atom densities. Provided there are less than 10 atoms in the cavity for the black (\square) curve, these transmission versus drive power plots qualitatively follow the shape, position, and relative depth that we would expect if the intracavity photon number is roughly 10 times more than what the “Reference level” (measured post-cavity) indicates⁶. (For rough comparisons, see Figure 7.10.) This factor of 10 roughly concurs with the estimation of $\eta_1 = 0.1$ for this batch of cavity mirror coatings, though some fraction of η_2 should be accounted for as well. This estimation is very loose and is not based on any stringent curve fitting. Such a fitting sequence should be performed, though it will be difficult since not only is the intracavity atom number uncertain but so is the calibration of the x-axis. Specifically, the calibration of the intracavity power is unknown due to the difficulty in measuring the photon loss from the interface between the coupling fiber and the microscopic mirror coating and the loss in the coatings themselves. Moreover, the rejection of light mode-mismatched into both the input and output fibers must be taken into account in such a measurement. While upper bounds can be placed, exact calibration would be difficult⁷. Although there is a considerable amount of light scattered by the coatings into 4π , the fraction of this light that interacts with the atoms—located in the center of the cavity—should be of the correct orientation to re-couple into the cavity mode and should not be treated differently. A major goal would be to perform a master equation fitting sequence that varies the relative intracavity atom number, N , and drive power, E , as well as the detunings Δ and θ . Although a major undertaking, the high sensitivity of the cavity transmission to these parameters provides enough constraints to pin-down the atom number and intracavity photon number calibrations.

⁶The high transmission spike at low drive powers remains unexplained. A simple cause would be the background and dark counts of the APD. Careful measurements of these count rates seem to rule this out, but a duplicate APD with perhaps better dark count rates should be used as a cross-check. Perhaps optical pumping at low drive powers in the magnetic trap is the culprit.

⁷However, Jun Ye recently communicated to Jakob Reichel that the techniques presented in Reference [157] might aid in determining this calibration.

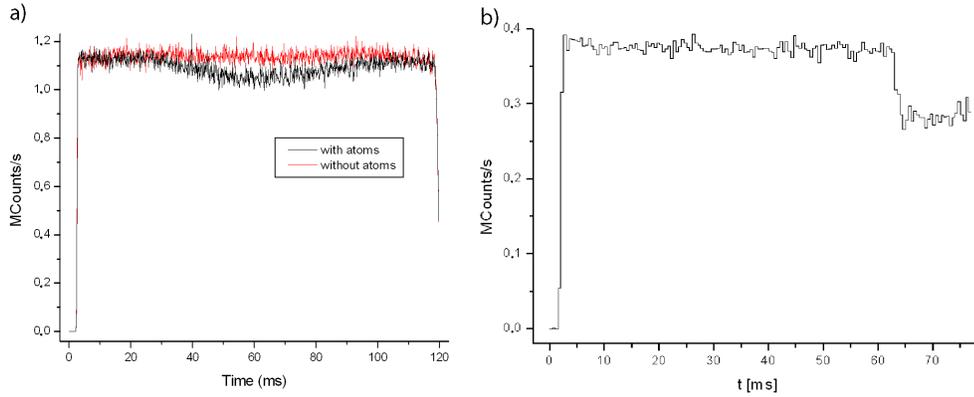


Figure 7.18: The first magnetically guided atom transits detected with a fiber-gap Fabry-Perot cavity integrated on an atom chip. For this cavity $\mathcal{F} = 600$, $L = 40 \mu\text{m}$, $R = 1 \text{ mm}$. The atoms generally take $\sim 140 \text{ ms}$ to arrive at the cavity mode. a) The first detection signal. b) After an optimization of the position of the atoms inside the cavity mode, the detection signal is enhanced. The time bins are $500 \mu\text{s}$ and the trace is averaged 20 times.

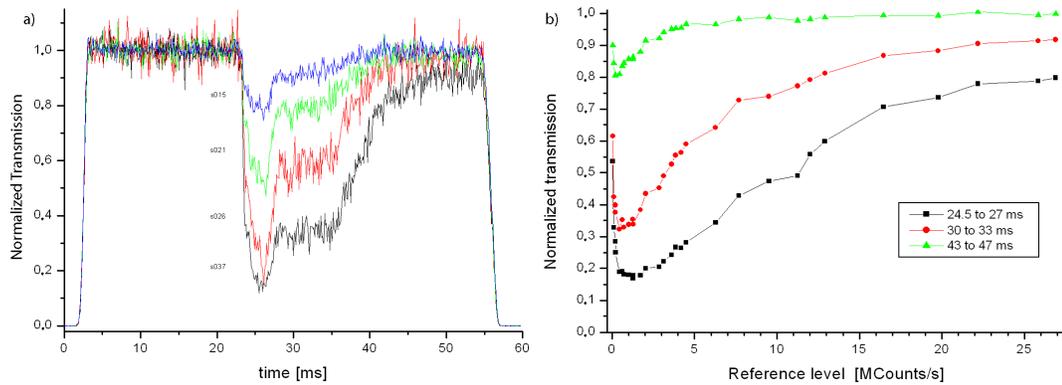


Figure 7.19: a) Atom transits through the fiber-gap cavity for different drive powers. Drive power increases from the black to the blue curves (bottom to top). For this cavity $\mathcal{F} = 600$, $L = 40 \mu\text{m}$, $R = 1 \text{ mm}$, $[g_0, \kappa]/2\pi = [139, 3125] \text{ MHz}$, and $[m_0, N_0] = [2.3 \times 10^{-4}, 9.7 \times 10^{-1}]$. Ten averages are taken for each curve. b) Transmission versus drive for three different intracavity atom numbers.

With the new, higher finesse fiber-gap cavity, we have been able to see atom cloud transits in a single-shot; double passage of the atom cloud through the resonator due to the reflection of the atoms at the terminus of the magnetic waveguide; a signal reminiscent of optical bistability; and a sharp dip in the transmission that may indicate the temporal resolution of a few or even a single intracavity atom. Panels (a) and (b) of Figure 7.20 show the on-resonance cavity transmission of a RF-cooled cloud of atoms guided through the cavity. In panel (a), zero averaging has been done, while 10 averages have been taken for the transmission signal in panel (b). Panel (c) shows a “double-bounce” of the atoms through the cavity caused by reflection at the terminus of the magnetic trap. The second dip is smaller than the first due to the loss of atoms from spontaneous emission, collisions with the cavity mirrors, and normal magnetic trap loss. As expected, by switching the waveguide current from [G1,GL8] to [G1,GU8] we observed a difference in the time delay between the dips in transmission. A possible future experiment that would highlight the unique abilities of this atom-cavity chip would involve the periodic introduction of magnetically trapped atoms into and out of the cavity mode. The atoms would be confined in a Z-trap centered at the resonator mode ([GU5,GL7], for instance), and the bias field and wire current adjusted to move the atoms vertically in and out of the resonator mode while maintaining a constant trap compression. For $[\Delta, \theta] = 0$, one should detect periodic dips in cavity transmission that decrease in amplitude with time. The attenuation of signal amplitude would be caused by the loss of atoms due to optical pumping, collisions with cavity mirrors, and the natural decay of atoms from the magnetic trap. These rates could be measured by fitting the decaying envelope of the periodic atom transit signals. By repeating the experiment with and without an intracavity light field, and noting the unperturbed lifetime of the magnetic trap, one could distinguish and measure the individual decay rates.

By detuning the cavity system, we have been able to enter a regime in which there are an excess of photons coming out of the cavity during the atom transits. These “up-transits” are shown in Figure 7.21. The detunings are $[\Delta, \theta] \approx [-3.5\gamma, -\kappa]$ in Panel (a) and $[\Delta, \theta] \approx [2.3\gamma, \kappa]$ in Panel (b). In Panel (c) the Δ detuning is the same as in Panel (b) but θ is slightly larger. The exact cavity-laser detuning is hard to calibrate in these data runs since the cavity is detuned by simply adjusting the feedforward signal and not by precisely offset-locking the cavity. For a single atom, Figure 7.12 shows the expected cavity transmissions, S , and S/\mathcal{N} for similar detunings used for the data runs in Figure 7.21.

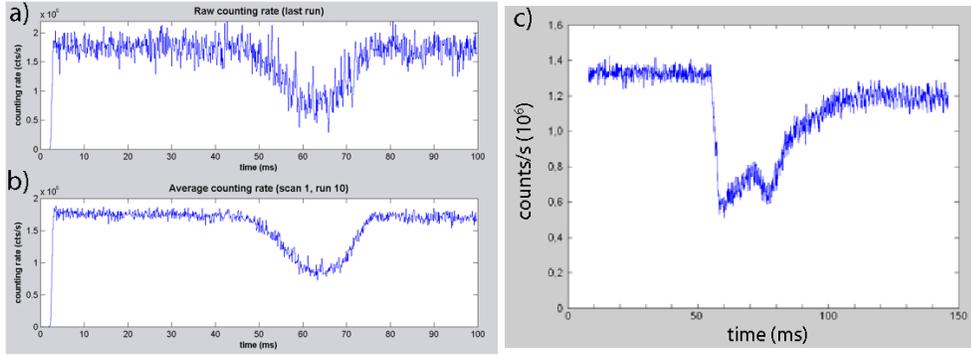


Figure 7.20: a) Single shot measurement of an atom cloud passing through the fiber-gap cavity. b) Same as (a) but averaged 10 times. c) Double passage of the atoms through the resonator mode due to reflection at the terminus of the magnetic waveguide. $\Delta t = 100 \mu\text{s}$, $[g_0, \kappa]/2\pi = [186, 2640] \text{ MHz}$.

The Figure 7.12 calculation is qualitatively consistent with the up-transit signal observed at early and late times in the atom cloud transit through the cavity. However, during the peak atom density at $t \approx 45 \text{ ms}$, the up-transit regime switches to a down-transit in Panels (a) and (b). This may be due to the increase in g_{eff} which shifts the eigenmodes of the atom-cavity system. This behavior is not seen for the detunings in Panel (c). Fitting these types of data curves for varying detunings and drive powers should allow us to calibrate the intracavity atom and photon numbers.

A hint of optical bistability in this system has been observed by ramping the drive power up and down during a single shot of constant atom flux through the cavity. Panel (a) of Figure 7.22 shows the window of nearly constant resonator transmission that we infer is due to constant atom flux. We use this window to take the transmission curves that are shown in panel (b). This situation is similar to atomic beam experiments since the atoms—confined only in two-dimensions—have a non-zero mean velocity as they pass through the resonator (see Reference [166] and citations within). One disconcerting feature of the supposed optical bistability curves in panel (b) is the absence of a sharp field switching. This could be caused by the stochastic variations of the atom number and atomic motion which would induce g_{eff} to fluctuate, or by the random fluctuations of the cavity detuning which is not locked during the atom transit portion of the experiment. It would be interesting to see if a model incorporating the stochastic fluctuation of g_{eff} and θ accounts for this washing-away of the normally sharp hysteresis cycle. To rule-out thermal effects, an experiment should be done

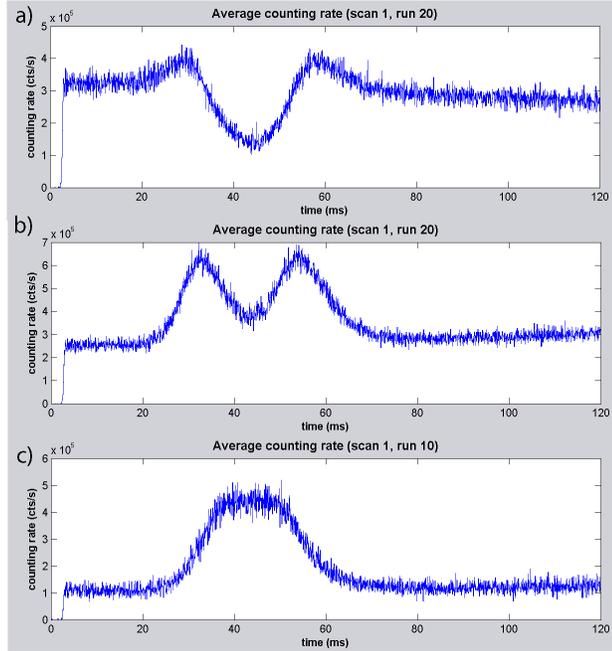


Figure 7.21: a) $[\Delta, \theta] \approx [-3.5\gamma, -\kappa]$; b) $[\Delta, \theta] \approx [2.3\gamma, \kappa]$; c) Δ detuning is the same as in (b) but θ is slightly larger. The data are averaged 20 times in (a) and (b) and 10 times in (c). $\Delta t = 100 \mu\text{s}$, $[g_0, \kappa]/2\pi = [186, 2640] \text{ MHz}$.

that tries to ramp the drive back and forth several times to see if the optical bistability curves are retraced. Trapping atoms in a Z-trap and/or dimple trap at the resonator would aid in eliminating the loss of atoms that is presumably the cause of the disparity in the high drive power transmissions of the black and red curves in panel (b). Instead of taking one data point per drive power in a single-shot, a future experiment should generate the optical bistability curves by ramping the drive to a series of set values and taking several data points at each value of the drive.

Our efforts to see single atom transits by diluting the waveguided atoms—via extra RF cooling—have been unsuccessful so far. The dip in transmission vanishes rather than revealing discrete spikes. One possible explanation is that the single atom \mathcal{S}/\mathcal{N} is not as high as it should be given the calculated g_0 and κ . This may be due to a non-optimal polarization of the intracavity field and/or to optical pumping of the atom that renders it dark before the end of the detection period. To test this, we shut-off the microwave waveguide magnetic field as the atom cloud passes the resonator. The $B_z = -16 \text{ G}$ is maintained, which provides a large quantization field parallel to the resonator axis. In the

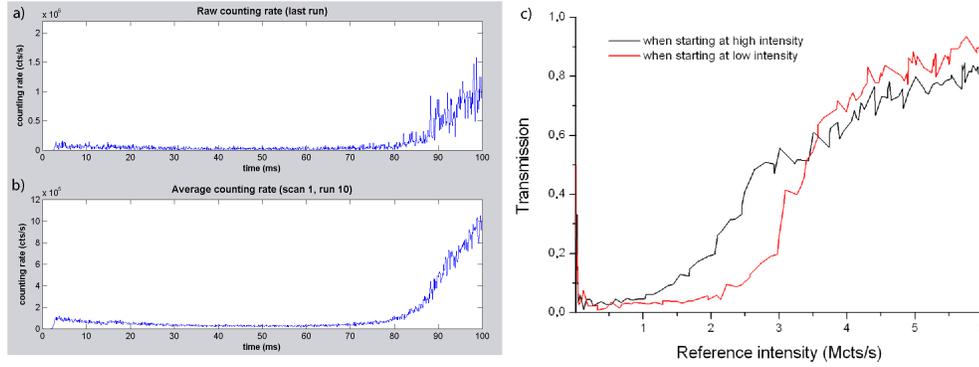


Figure 7.22: a) Regime of nearly constant intracavity atom flux. $\Delta t = 100 \mu\text{s}$. b) Same as (a) but averaged 10 times. c) Possible optical bistability signature. $[g_0, \kappa]/2\pi = [186, 2640]$ MHz.

first trial only a small atom transit signal was observed. This is expected since the large bias field Zeeman shifts the atom out of resonance with the cavity and laser probe (the detuning is 1.4 MHz/G). We recover some of the transmission signal depth by offsetting the laser diode lock to compensate for this detuning. The fiber-gap cavity—whether due to the epoxy, the mirror coating, or fiber itself—has an unknown birefringence. We rotated the $\lambda/2$ and $\lambda/4$ waveplates to attempt to obtain purely circularly polarized light in the cavity mode. Soon after adjusting these waveplates, we observed a sharp dip in transmission immediately after the waveguide turn-off time at $t = 20$ ms. The detection resolution was reduced to $\Delta t = 50 \mu\text{s}$ without reducing the depth of the down-transit dip (see Figure 7.23). While this seems like a promising single atom detection signal, the transient behavior of the system at this detection time casts doubts upon this conclusion. Further measurements are required to fully understand what is occurring. We can be sure, however, that adjusting the intracavity polarization greatly improves detection \mathcal{S}/\mathcal{N} . Ideally, we would like to trap the atoms in a Z-trap and/or dimple trap at the resonator and see several time-resolved transmission dips that become ever more rare as the trap density is reduced. Towards these ends, we are currently trying to improve the loading efficiency of the Z-trap and dimple trap at [GU5, GL7] and [GU6, GL6], respectively. The dimple trap is advantageous since the field at its minimum can be rotated [50] to be nearly in line with the resonator axis which would provide an in-trap quantization field. To improve the trap loading, a quadrupole trap—generated by external anti-Helmholtz coils—is superimposed on the waveguide field and biased to rapidly shift (within less than 100 ms) the atoms down the waveguide while

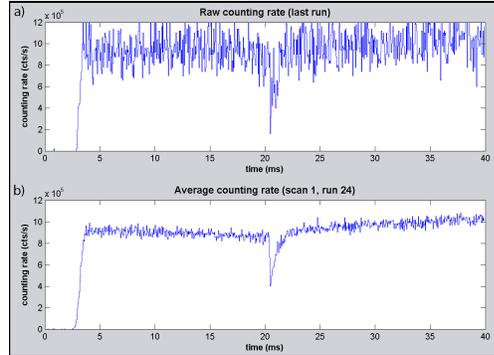


Figure 7.23: a) Transmission dip after turning-off the magnetic waveguide at $t = 20$ ms. The $B_z = -16$ G bias field is maintained after the waveguide is turned-off and the probe field detuned to compensate for the atom’s Zeeman shift. The sharp signal appeared after adjusting the intracavity polarization. $\Delta t = 50 \mu\text{s}$. b) Same as (a) but averaged over 24 trials. $[g_0, \kappa]/2\pi = [186, 2640]$ MHz.

maintaining high phase-space density. So far we have only been able to hold atoms in the Z-trap and dimple trap in the cavity mode for 20 ms. Once improved, we hope to have a well-controlled system with which to pursue signatures of single atom transits.

7.5 Outlook

In addition to single atom detection with this novel atom-cavity chip, we are actively pursuing the on-chip production of a BEC for the purpose of insertion into the mode of the fiber-gap cavity. On-chip BEC production has already been achieved in an identical atom chip system in our labs in Munich (see Chapter 10). The current plan is to pre-cool the atoms at the location of the initial trap at the beginning of the waveguide. Then the waveguide will open and by using either the assistance of a dimple trap or the external quadrupole trap, shift the atoms to the position of the resonator. Once in position, final evaporative cooling will be performed, producing a BEC in an Ioffe trap that is brought into the cavity mode. Optical Fock-state generation is one potential application of having many cold and well-localized atoms in a cavity [177, 178]. Studies of the superfluid to Mott insulator transition would be interesting in this experimental situation, as would be the demonstration of a “quantum tweezer” [179]. Injecting a secondary light field into the cavity could form dipole traps for the creation of the Mott insulator phase [180] while the primary cavity probe couples to the atoms for cavity QED measurements. Similarly, an intracavity, attrac-

tive optical dipole potential could be used for extracting single atoms from a BEC passing through the cavity as it travels along the magnetic waveguide. This quantum tweezer [181] could be useful for delivering one and only one atom to the ground state of an intracavity trapping potential. Measurements of the photon counting statistics (for instance, $g_2(\tau)$) for a stationary BEC coupled to a cavity would be interesting to compare to that measured from a coupled cloud of thermal atoms [182].

Chapter 8

Atom Mirror Etched from a Common Hard Drive

This chapter includes the material contained in publication [12] and augments it with additional figures and experimental details.

8.1 Overview

Laser cooling and trapping techniques have made possible the preparation of extremely cold samples of atoms. Atom optics employs elements such as mirrors, lenses, gratings, and beam-splitters to manipulate these cold atoms in a fashion similar to the familiar photon optics [46]. The advent of atom lasers from the Bose-Einstein condensation of neutral atoms [183] has enhanced the importance of developing atom optical elements. In particular, atom mirrors—surfaces that reflect atoms—play a crucial role in the field of atom optics, and it is of keen interest to develop mirrors that are simple to fabricate yet highly specular. In this chapter we demonstrate a straightforward technique to produce large area, high resolution permanent-magnetic structures on flat, rigid, and inexpensive substrates.

Several types of atom mirrors have been fabricated using evanescent light fields [184], dynamic magnetic fields [185], and static magnetic fields [10]. Evanescent mirrors repulse atoms from a prism surface using a potential created by a blue-detuned light field. Although magnetic mirrors cannot generally be modulated as easily as evanescent mirrors, they do offer many advantages: passive operation, compactness (no laser access is needed), and much larger repulsive areas.

Magnetic mirrors employ a sheet of alternating current or magnetization to create an exponentially increasing potential near the mirror surface [186]. To lowest order, this potential is proportional to $B_0 e^{-ky}$. The surface field, B_0 , sets the maximum atom energy that can be reflected, and the spatial period of the current or magnetization, $a = 2\pi/k$, determines the amount of time the atoms interact with the mirror. The magnetic mirror approximates a perfectly flat mirror as B_0 increases and a decreases. For example, if $B_0 = 1$ kG and $a = 1 \mu\text{m}$, a cesium atom in the $6^2S_{1/2} F = 4, m_F = 4$ state will be reflected when dropped from a height of 0.4 m, and will only interact with the mirror for $5 \mu\text{s}$ if dropped from 2 cm.

Mirrors made from serpentine patterns of wires can produce time-dependent reflection potentials. However, they have not been fabricated with periods smaller than $10 \mu\text{m}$, and the power dissipated by the small wires requires cooling by liquid nitrogen and pulsed operation [53]. Sinusoidal magnetization of audio-tape, floppy disks, and videotape can produce magnetic mirrors with magnetization periods down to $12 \mu\text{m}$ [10]. Mirrors made from millimeter-sized arrays of permanent magnets have been demonstrated, as have mirrors produced by 1 to $4 \mu\text{m}$ periodic structures fabricated by sputtering ferromagnetic material onto a grooved substrate patterned by electron-beam lithography [11].

We recently fabricated a magnetic mirror by etching a common hard drive, and we used this mirror to retroreflect a cold cloud of 10^6 cesium atoms. Hard drives offer several advantages for making and using atom mirrors. The common hard drive provides a large surface area of thin magnetic film whose surface is specifically designed to be very flat, smooth, and rigid. Furthermore, the film's remnant magnetic field and coercivity can be as large as 7 kG and 3 kG, respectively [187]. An atom mirror could in principle be fabricated with a $2 \mu\text{m}$ periodicity over the entire surface of the hard drive. Old or discarded hard drives may be used: an Apple hard drive from the mid-1990's was used for the experiment presented here.

8.2 Design and fabrication of the hard drive atom mirror

We fabricate the mirror by etching $2 \mu\text{m}$ wide, ~ 100 nm deep trenches into a 1 cm^2 section of the surface of the hard drive. These 100 nm trenches extend past the magnetic layer to form a periodic array of $1 \mu\text{m}$ wide, 30 nm thick, and 1 cm long stripes of cobalt alloy

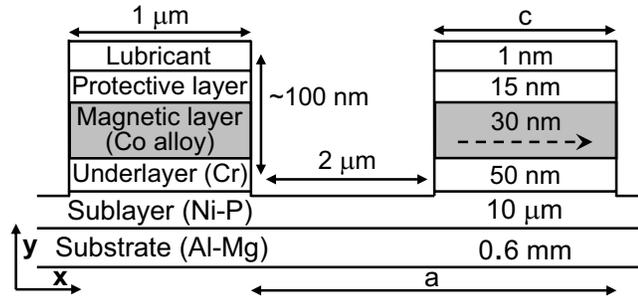


Figure 8.1: A cross-section of the etched hard drive. The magnetization is in-plane. See Reference [187] for a description of the hard drive layers.

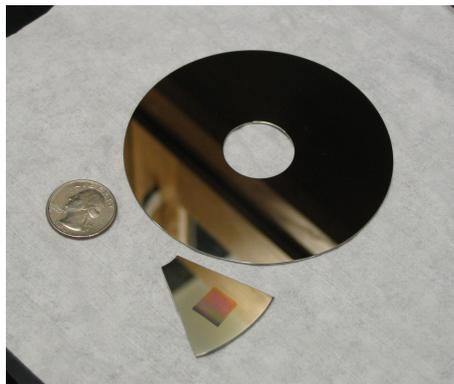


Figure 8.2: The relative size of the hard drive sliver used in this experiment. The square discolored region is the 1 cm^2 etched pattern.

(see Figure 8.1). The cobalt alloy is granular, which enhances the coercivity and allows us to magnetize the material in plane and parallel to the short axis of the magnetic strips. The typical grain size is 20 to 50 nm [187], and we expect the magnetization to be uniform for our much larger features. We do not know the exact materials and thicknesses of the layers of the proprietary hard drive. However, etching $\sim 100 \text{ nm}$ is sufficient to remove the magnetic layer.

Standard photolithography is used to create the etch mask. After cutting the hard drive into 2 to 3 cm^2 sections (see Figure 8.2), positive photoresist (TSMR-8900 from Tokyo Ohka Kogyo Co.¹) is spun onto the cleaned hard drive surface for 40 s at 4200 rpm. A 5 min bake at 98°C followed by a 15 s UV exposure and 65 s in the developer (NMD-W 2.38%) maps the photomask lines into the resist. To make the one micron periodic features, it is important to have the photomask perfectly flush with the photoresist on the hard drive. For this to be

¹Asa Hopkins has since found that AZ1518 works as well.

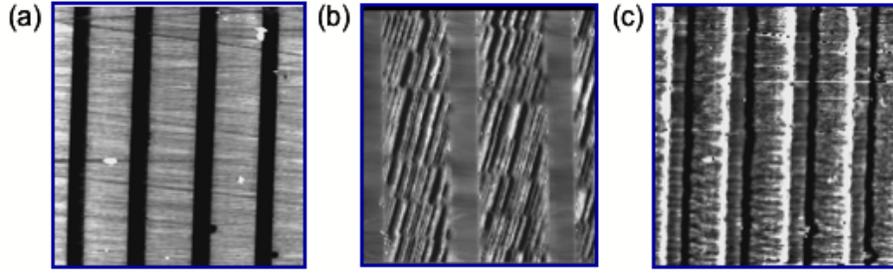


Figure 8.3: Example of etched hard drive with $8 \mu\text{m}$ wide magnetic strips and $4 \mu\text{m}$ wide trenches for viewing convenience. (a) AFM of etched drive. (b) Black and white stripes are the magnetic bits. The magnetic layer is removed in the neutral grey vertical stripes. (c) An 8 kG magnetic field erases the bits and creates north/south pole stripes defined by the vertical etch edges.

accomplished, it is imperative that one removes the beaded photoresist at the edge of the substrate with a second UV exposure. The sample is ion milled with argon in an inductively coupled plasma (ICP) system. We etch for 8 minutes at a forward power of 100 W, ICP power of 400 W, and an argon flow of 40 sccm. The remaining photoresist is removed with acetone and, if necessary, a soft swab. To erase the hard drive's bits and magnetize it as a mirror, we insert the hard drive section into the field of an 8 kG electromagnet whose field is parallel to the surface and perpendicular to the magnetic stripes (Figure 8.3 illustrates this process). This is done by holding the hard drive in the correct orientation as the electromagnet is slowly ramped up and back down.

The magnetic field from the etched hard drive, with in-plane magnetization, M_0 , parallel to the short axis of the magnetic stipes, is analogous to a periodic sheet of alternating in-plane magnetization $+M_0/2$ and $-M_0/2$. In the infinite array limit, the magnetic field above the surface is

$$B^2 = B_1^2 e^{-2ky} + 2B_1 B_3 \cos(2kx) e^{-4ky} + B_3^2 e^{-9ky} + \dots, \quad (8.1)$$

where $B_1 = \mu_0 M_0 (1 - e^{-kb})/\pi$, $B_3 = \mu_0 M_0 (1 - e^{-3kb})/3\pi$, and $b = 30 \text{ nm}$ is the thickness of the magnetic layer. The field has no components in the z -direction, and rotates with a period equal to a in the x - y plane. Cesium atoms in the $F = 4$, $m_F = 4$ state, which has the largest weak-field seeking magnetic moment, would have to be dropped from a height of 25 cm to penetrate to a height at which the second term in the expansion is equal the

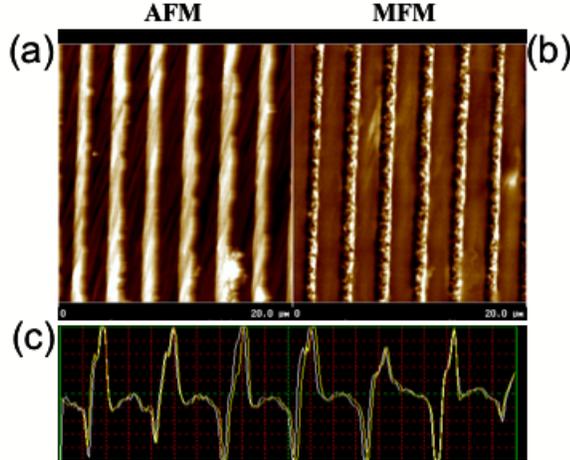


Figure 8.4: Twenty micron wide (a) AFM scan, (b) MFM scan, and (c) MFM cross-section of the etched hard drive surface.

first, so to a good approximation the field may be written as

$$B \approx B_1 e^{-ky} + B_3 e^{-3ky} \cos(2kx). \quad (8.2)$$

For our hard drive mirror, B_1 is equal to 2 to 4 kG depending on the specific cobalt alloy. When $a = 1 \mu\text{m}$, the ratio of the first harmonic term to the purely exponential term for a cesium atom dropped from a height of 2 mm (20 mm) is 1×10^{-6} (1×10^{-3}) at the turning point $y = 0.8 \mu\text{m}$ ($y = 0.4 \mu\text{m}$). Subsequent to this experiment, we learned that custom hard drive platters could be made for us with B_1 as large as 20 kG [15], enabling the reflection of much higher energy atoms or the creation of tighter traps (see Chapter 9).

The etched hard drive used for the experiment has $a \approx 3 \mu\text{m}$ and $c \approx 1 \mu\text{m}$ resulting in a ratio of magnetic layer to gap that is approximately 1:2. Figures 8.4 (a) and (b) show $20 \mu\text{m}$ wide AFM and MFM scans of the hard drive surface. The trenches in the AFM scan are dark, and the light to dark variation of magnetic strips shows the north and south poles of the magnetization. Figure 8.4 (c) shows a $20 \mu\text{m}$ cross-section of the MFM scan: peaks represent the north and south poles. To describe the field above our etched hard drive, Equation 8.2 can be modified to account for the deviation from a 1:1 width ratio by multiplying B_1 by $\sin(\pi c/a)$ and B_3 by $\sin(3\pi c/a)$. In our device, the ratio of $c/a \approx 1/3$ decreases the B_1 term by 0.9, but causes the corrugation term to nearly vanish.

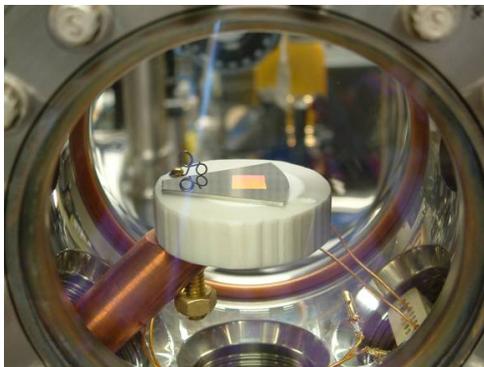


Figure 8.5: The hard drive is clamped-down onto a teflon block in the middle of the vacuum chamber. The diffraction grating character of the etch hard drive is apparent in this image which is taken through one of the AR coated 6" viewports.

8.3 Experimental details

The atom mirror is placed facing upwards in a vacuum chamber pumped to 5×10^{-9} Torr (see Figure 8.5). In contrast to the experiments that use a standard magneto-optical trap (MOT) to trap and cool the atoms ~ 2 cm above the mirror, we use the mirror MOT technique to collect the atoms 1.5 to 4 mm above the surface [2]. A MOT requires the zero of a magnetic quadrupole field to be centered at the intersection of six circularly polarized laser beams coming from all cardinal directions. To satisfy this configuration near the hard drive surface, two 1 cm diameter beams of opposite circular polarization reflect at 45° from the 1 cm^2 etched region (see Figure 8.6). A retroreflected beam is positioned perpendicular to the 45° beams and grazes the surface of the hard drive. Aligning the axis of the quadrupole field with one of the 45° beams completes the mirror MOT configuration. The trapping lasers, each with an intensity of 4 mW/cm^2 and 1 cm wide, are detuned by 10 MHz from cesium's $F = 4$, $F' = 5$ cycling transition. A repumping beam tuned to the $F = 3$, $F' = 4$ transition is superimposed onto both the grazing beam and a 45° beam. The atoms are loaded from a thermal vapor.

In previous experiments using a perfectly reflecting gold mirror, we have been able to trap 2×10^6 cesium atoms in a mirror MOT and cool them to $3 \mu\text{K}$. One might expect trapping and cooling to be much less effective with the etched hard drive due to its poor qualities as an optical mirror: the reflectivity is only $\sim 50\%$ in the etched area and 60-70% in the un-etched, it is a good optical grating with up to 3 diffracted orders visible,

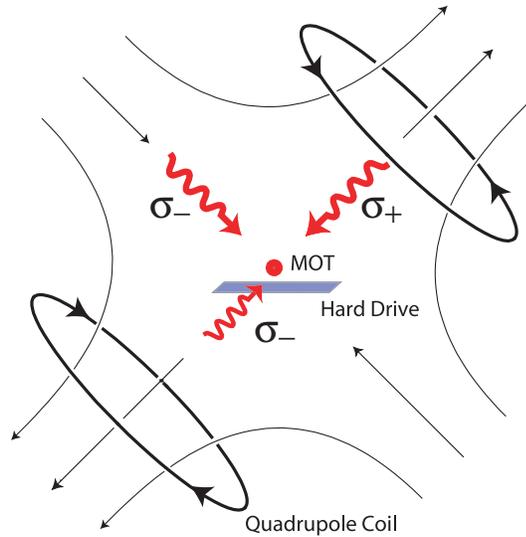


Figure 8.6: Diagram of the experimental set-up. A quadrupole field and two 45° laser beams and one retroreflected grazing beam form a mirror MOT 1.5 to 4 mm above the etched hard drive.

and the magneto-optical Kerr effect degrades the circularity of the reflected 45° beams. Nevertheless, we have been able to collect 1×10^6 atoms and sub-doppler cool them to $11 \mu\text{K}$. Achieving this low temperature is crucial because the atoms released directly from the mirror MOT, at a temperature of $\sim 120 \mu\text{K}$, expand too quickly and become too diffuse to detect by the time they reach the hard drive surface.

The poor optical reflectivity of the mirror does slightly complicate the sub-doppler cooling procedure; however with careful zeroing of the magnetic field it is still possible to achieve polarization-gradient cooling to $11 \mu\text{K}$ in a (downwards) moving reference frame. The problem arises because the reflected 45° beams are too attenuated to properly counterbalance the incoming beams, forcing the atoms downwards in the absence of a quadrupole field. For optimal sub-doppler cooling, one should adjust the bias fields until the atoms drop straight downwards and expand as slowly as possible. This can be achieved by taking fluorescence images from all three angles, adjusting bias fields between image sets, and gradually increasing the image delay time. The atoms are optically pumped into the $F = 4$, $m_F = 4$ Zeeman sub-state just before being dropped, and we apply a 100 mG bias field parallel to the magnetic stripes in order to maintain alignment of the atomic spins while they are falling/bouncing.

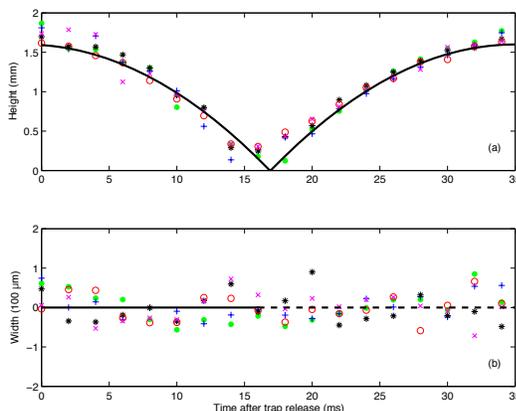


Figure 8.7: Panel (a) shows the mean height of the atoms above the hard drive surface during the first bounce. Panel (b) shows the residuals from a linear fit to the width of the expanding atom cloud.

8.4 Results

We have been able to detect two full bounces of the atoms from the hard drive atom mirror². Figures 8.7 (a) and (b) show data from five runs of the experiment. The top panel shows the mean position of the atoms above the hard drive surface as a function of time. Superimposed is a curve depicting the expected trajectory of a particle falling under gravity and bouncing from a hard wall. The slope of a line fit to the lateral expansion of the falling atom cloud provides a measure of the atoms' rms velocity. A non-specular mirror would heat and diffusely scatter the reflected atoms as they bounce, resulting in a sharp increase of the cloud expansion rate. We made a linear fit to pre-reflection ($t < 15$ ms) data in each of the data sets, and deviation from this line, post-reflection, would be evidence of non-specularity. The dashed segment demarcates the region of unfitted data, and we do not see any increase or offset of the residuals in this post-reflection region: to within the experimental resolution, we do not detect any deviation from specular reflection.

We have realized a specular atom mirror built by etching a common hard drive. Magnetization periodicity of $3 \mu\text{m}$ has been achieved, and we believe it would be straightforward to reduce this to $2 \mu\text{m}$ with photolithography and to $\sim 1 \mu\text{m}$ using a large area electron-beam writer. The hard drive atom mirror is compact, passive, relatively simple to fabricate, and possesses a large remanent magnetic field. Moreover, it has several desirable properties for

²A movie of the bouncing atoms—taken by fluorescence imaging—may be found on either the author's or the MabuchiLab's website [188].

applications beyond the simple reflection of atoms. The hard drive's large coercivity should allow one to use wires fabricated directly on its surface to augment the mirror's ability to manipulate atoms. Likewise, electric pads could be printed on the surface. These pads would allow state-independent forces to act in concert with the state-dependent forces from the mirror's magnetic field to perform quantum logic gates necessary for quantum computation [13]. The mirror can trap cold atom gases in 2D, and can act as an adjustable grating when used in conjunction with a magnetic bias field [14, 189]. Large area mirrors can be fabricated, and it seems possible that these mirrors could be useful for guiding or confining cold neutrons [190]. As hard drive platters are expected to have good surface flatness and substrate rigidity, it may be possible to create 2D waveguides by holding an opposing pair of atom mirrors a few microns apart. Corner cubes and other such atom optical devices are possible.

Chapter 9

A 1-D Magneto-electrostatic Ring Trap for Neutral Atoms

This Chapter is adapted from Reference [45] and contains more references and a few extra pieces of information.

We propose a novel trap for confining cold neutral atoms in a microscopic ring using a magneto-electrostatic potential. The trapping potential is derived from a combination of a repulsive magnetic field from a hard drive atom mirror and the attractive potential produced by a charged disk patterned on the hard drive surface. We calculate a trap frequency of [29.7, 42.6, 62.8] kHz and a depth of [16.1, 21.8, 21.8] MHz for [^{133}Cs , ^{87}Rb , ^{40}K], and discuss a simple loading scheme and a method for fabrication. This device provides a one-dimensional potential in a ring geometry that may be of interest to the study of trapped quantum degenerate one-dimensional gases. With custom hard drive materials, trap frequencies in excess of 100 kHz may be attainable.

9.1 Overview

Creating ever more sophisticated trapping potentials has become a standard method for the study and manipulation of cold neutral atoms, allowing the investigation of fundamental quantum dynamics as well as providing a basis for quantum information processing. The manipulation of trapped atoms on atom chips allows the implementation of many different atom optics elements for trapping, waveguiding, interferometry, etc. [50, 68, 51]. Most atom chips use micron-sized current-carrying wires to generate the magnetic trapping fields. We propose to construct a magneto-electrostatic ring trap, consisting of a hard drive atom

mirror that provides a repulsive force on low-field seeking atoms [12] and electric pads that attract polarizable atoms via the Stark effect [70, 10, 80]. Schmiedmayer and Hinds and Hughes have proposed a range of such traps, including large-area two-dimensional traps, wire-based waveguides, and quantum-dot-like single state traps. Such traps could be used to construct beam splitters or to implement collisional quantum gates [13]. Here we propose a novel ring trap for cold neutral atoms constructed from a conducting disk placed above the atom mirror surface, which produces a trap with a deep ring potential around the edge of the disk. Unlike the ring-shaped waveguides for neutral atoms recently proposed [191], and demonstrated [192], this ring trap would create a tight enough potential to confine a degenerate gas in a 1D regime.

9.2 Ring trap design

Let us first examine the trapping potential from a charged conducting disk above a hard drive atom mirror. The hard drive's sinusoidal pattern of magnetization results in a repulsive potential—for atoms in weak-field seeking states—in the form of a decaying exponential [186]

$$U_{mag} = m_F g_F \mu_B B_0 \exp[-2\pi z/a]. \quad (9.1)$$

The amplitude, B_0 , depends on the remnant magnetization of the mirror as well as the magnetic sublevel m_F and Landé g_F -factor of the atomic ground state. The decay length is proportional to the periodicity a of the magnetization pattern. A small externally applied magnetic field perpendicular to the magnetization of the hard disk eliminates zones of zero magnetic field which would allow Majorana spin-flip losses. The atom's low velocity allows the spin adiabatically to follow the magnetic field and thus the trapping potential depends only on the field magnitude.

In order to create a trap, the repulsive force from the mirror is balanced by an attractive force due to the DC Stark effect. The atomic potential due to an electric field is

$$U_{Stark} = -\frac{1}{2}\alpha |E|^2, \quad (9.2)$$

where we assume that we are working with atoms such as cesium or rubidium which possess only a scalar polarizability in the ground state. A charged conducting disk creates high

electric fields near its edge, resulting in a strong short-range attractive potential.

The mirror is made out of an etched hard drive whose aluminum substrate is grounded. The boundary conditions consist of a ground at the mirror surface, and a constant potential on the surface of the thin conducting disk which is placed a distance d , typically on the order of a micron, above the mirror. The electric fields are calculated from the solution to the Poisson equation with these boundary conditions. The combined atomic potential due to the charged disk and mirror creates a trap above the conducting disk, which is deepest near the edge of the disk.

As an example, consider a conducting disk of radius $10\ \mu\text{m}$, placed $d = 0.6\ \mu\text{m}$ above a hard drive atom mirror. Let the hard drive have a field at its surface of $2\ \text{kG}$ (a typical number for a commercial hard drive), and a periodicity of $3\ \mu\text{m}$ in the magnetization. The trapping potential for cesium in the $F = 3$, $m_F = -3$ state near the edge of the disk has a depth of $16.1\ \text{MHz}$ ($770\ \mu\text{K}$) when the potential on the conducting disk is $14.2\ \text{V}$. For ^{87}Rb in the $F = 2$, $m_F = 2$ state, the trap has a depth of $21.8\ \text{MHz}$ ($1.05\ \text{mK}$) when $18.5\ \text{V}$ is applied to the disk. These two atomic states will be used in all examples for the remainder of the paper. See Figure 9.1 for the ^{133}Cs potential. The ^{87}Rb potential looks qualitatively the same, with a slightly deeper minimum. See Table 9.1 for trap parameters for a range of geometries for ^{133}Cs and ^{87}Rb , respectively. For ^{40}K , the optimal applied voltage is 4% larger than that for the ^{87}Rb trap, and trap frequencies scale up by a factor of $(m_{\text{Rb}}/m_{\text{K}})^{1/2} = 1.48$ relative to the ^{87}Rb case.

The potential applied to the conducting disk is chosen to create the deepest trap while still maintaining a barrier between the trap and the disk surface. If the potential applied to the disk is too large, atoms will simply be forced directly into the disk and lost. The separation between the mirror and the conducting disk must be chosen carefully. The trap becomes shallower as d is increased, due to the decay of the atom mirror field. As d is decreased the trap becomes deeper because a higher voltage can be used on the disk while maintaining the potential barrier between the trap and the surface.

The curvature of the trap is large enough that the atom is confined in the Lamb-Dicke regime. The Lamb-Dicke regime is defined as the regime in which $\eta = (E_{\text{recoil}}/E_{\text{trap}})^{1/2} < 1$. For the parameters of Figure 9.1, the effective harmonic frequencies for [^{133}Cs , ^{87}Rb , ^{40}K] in the radial direction are [29.7 , 42.6 , 62.8] kHz, and [40.6 , 56.8 , 83.8] kHz in the direction perpendicular to the substrate. We obtain a Lamb-Dicke

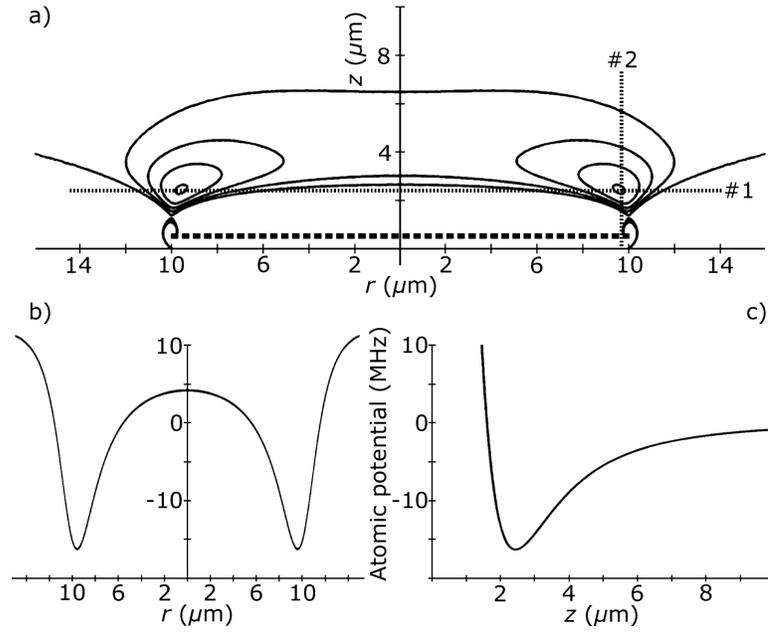


Figure 9.1: The atomic potential for Cs with 14.2 V on the disk. a) A cross section of the atomic potential in the plane containing the axis of the disk. The contour lines are spaced 4 MHz apart. The distance r along a diameter of the disk and the distance z above the disk are plotted on the horizontal and vertical axes, respectively. b) The potential along slice #1 in (a). c) The potential along slice #2 in (a).

Table 9.1: Cs¹³³ and Rb⁸⁷ trap parameters for several disk radii r and disk-hard drive separations d .

d (μm)	r (μm)	V	Trap depth	Trap frequencies	
			(MHz)	$\omega_r/2\pi$ (kHz)	$\omega_\perp/2\pi$ (kHz)
¹³³ Cs					
0.6	5	13.3	17.0	24.4	44.1
0.6	10	14.2	16.1	29.7	40.6
0.6	20	14.8	15.4	30.6	37.0
1.0	5	9.4	8.5	18.0	31.1
1.0	10	10.2	8.2	21.2	28.0
1.0	20	10.8	8.1	22.1	26.3
⁸⁷ Rb					
0.6	5	17.3	22.9	36.0	63.4
0.6	10	18.5	21.8	42.6	56.8
0.6	20	19.2	20.6	43.7	52.7
1.0	5	12.2	11.4	25.7	44.0
1.0	10	13.3	11.2	30.8	40.4
1.0	20	14.0	10.7	31.5	37.5

parameter of $\eta \leq 0.26$ for ^{133}Cs , $\eta \leq 0.30$ for ^{87}Rb , and $\eta \leq 0.37$ for ^{40}K . Significantly higher trap frequencies are possible with the use of custom magnetic materials, which can have remnant magnetic fields of up to 2.4 T [15]. For the same trap geometry as Figure 9.1, but using this custom magnetic material with a correspondingly higher applied voltage, the harmonic frequencies for ^{133}Cs , for instance, are 103 kHz in the radial direction and 137 kHz in the perpendicular direction. The higher remnant magnetic field also allows the disk to be placed further from the hard drive while maintaining significant trap depth.

9.3 Device electrical leads and trap perturbations

A thin lead running along the hard drive surface may be used to connect the disk to a voltage source. The maximum possible voltage on the disk is limited by the breakdown electric field of the dielectric material separating the lead from the conducting hard drive surface. An insulator which can support a field of 10^6 V/cm is sufficient to enable the application of ~ 20 V on a lead ~ 200 nm from the hard drive. In order to minimize the perturbation that the lead produces on the atomic potential from the disk, the lead should be as narrow as is practical (~ 1 μm) and placed much closer to the hard drive surface than to the disk. At this location, the repulsive force from the mirror is much stronger and no trap forms around the charged lead. In order to connect the lead to the disk, the disk is placed on a thin stem, with the lead connected to the bottom of the stem (see Figure 9.2).

Three dimensional solutions to the Poisson equation indicate that the effect from the lead on the trapping potential is minimized if the stem connecting the lead to the disk is located at the center of the disk. For a 10 μm disk and a lead placed 0.25 μm above the hard drive surface (0.35 μm below the surface of the disk), the trap depth for a ^{133}Cs atom rises to ~ 11.5 MHz above the lead, which is a $\sim 30\%$ loss of trapping potential compared to the unperturbed trap. The width of the perturbation is a few μm , which is slightly wider than the lead. A shallower trap in which the electric pad is placed further from the mirror surface is perturbed less by the lead. Use of custom magnetic materials would allow deeper traps to be constructed further from the lead, thereby minimizing the height of the perturbation [15].

There are several possibilities for minimizing or eliminating the perturbation due to the lead, or tuning it to be of a particular height, other than simply adjusting the trap

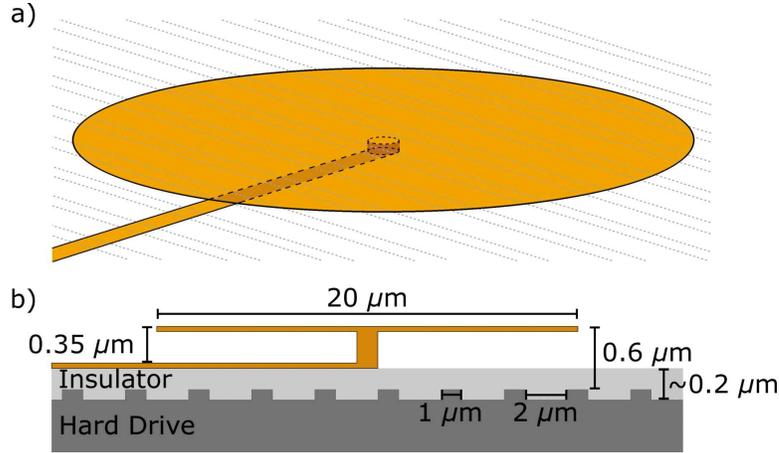


Figure 9.2: a) Schematic of the magneto-electrostatic ring trap drawn to scale. The disk is $20\ \mu\text{m}$ in diameter, with a $1\ \mu\text{m}$ wide lead connected via a central stem. The dotted lines show the hard drive atom mirror's 2:1 etch pattern with a $3\ \mu\text{m}$ periodicity. b) Cross section of the disk (with the vertical direction scaled up by a factor of 5), showing—from top to bottom—the disk, stem, lead, insulating layer, and etched hard drive.

geometry. The most versatile possibility is to add an additional photolithography step to insert another electric pad directly above the lead, separated by a thin insulation layer. The voltage applied to this separate pad can be used to compensate for the effects of the lead. In particular, the voltage on a pad the same width as the lead and placed $100\ \text{nm}$ above it could be tuned to completely eliminate the perturbation (to within the percent-level accuracy of our calculations) or turn it into a dip rather than a bump. Complete elimination of any perturbation is possible by expanding such a pad to cover the entire surface, with a hole to allow the stem to reach from the lead to the disk. Another possibility is to charge the disk not with a lead but with an intrachamber electron beam. Such a system would be hard to charge and discharge quickly, requiring a loading scheme that does not require a rapid change to the charge distribution.

9.4 Fabrication

We intend to fabricate the device as follows. The hard drive atom mirror is etched in the manner described in Reference [12], maintaining the 2:1 ratio of magnetization stripe spacing to minimize higher harmonics. The stripe periodicity will be $\leq 3\ \mu\text{m}$. A deposition of a $\sim 200\ \text{nm}$ thick insulating layer of silicon dioxide or silicon nitride is necessary to prevent shorting between the electric pads and the hard drive surface (see Figure 9.2). This layer

is thick enough to both support the voltage difference between the pads and underlying surface, and to help planarize the 100 nm deep corrugations of the etched hard drive. The ~ 50 nm tall, ~ 1 μm wide gold leads are patterned on the insulator surface using standard photolithography and thermal evaporation of the adhesion metal and gold layers [8, 90]. To create the stems, the surface is spin-coated with photoresist to a predetermined thickness to achieve optimal disk to atom mirror spacing. Photolithography is again used to create vertical, cylindrical holes of 1 μm diameter in the photoresist located at the terminals of the gold leads. The gold stems are electroplated from the gold leads through the cylindrical guide holes to the top of the photoresist. A third photolithographic process and thermal evaporation patterns the 20 μm diameter gold disks attached to the tops of the stems. Finally, the photoresist is removed using standard techniques, leaving behind the mushroom-like structures. It is not entirely necessary to remove the photoresist underneath the disks, as small amounts of photoresist will not impair the vacuum too greatly: field simulations show that to below the percent level, leaving the photoresist under the disk does not disturb the electric field. Moreover, perturbations to the trap due to disk edge roughness or due to the hard drive trench corrugations are both negligible to the percent-level resolution of our calculations.

9.5 Trap loading and surface effects

The trap is conservative once the voltage is established, and the kinetic energy of the atoms must be lowered for them to stay in the trap. A simple, but inefficient, method of loading this trap is to drop a cloud of cold atoms from a magneto-optical trap (MOT)—sub-doppler cooled to 10 μK —onto the device. The atoms are captured by turning on the voltage on the electric pads as the atoms are passing through their classical turning point above the atom mirror. Simulations indicate that this scheme can capture 1 to 2% of the dropped atoms. The fraction is small because the voltage ramp must be quite fast ($\sim 2 \times 10^{-4}$ seconds) in order to remove enough energy from the atoms to trap them, while the atom cloud takes roughly 2×10^{-2} seconds to pass through the trapping volume. This scheme has many different parameters over which loading can be optimized, including the initial position, size and density of the MOT before it is dropped, and the shape and speed of the voltage ramp. Ramping up the voltage on the conducting disk is the simplest scheme for

trapping the atoms, but it is possible that another procedure, involving atomic transitions or other degrees of freedom in the system, could be more effective and is currently being investigated. Using a procedure similar in spirit to that employed in Reference [14] but involving Raman transitions could prove to be more efficient. The following is a rough sketch of such a scheme. For the example of Cs atoms, the ring trap potential would be designed so that the lower $F = 3$, $m_F = -3$ hyperfine state is trapped whereas the upper state, $F = 4$, $m_F = 4$, is untrapped. A series of Raman transitions combined with optical pumping would transfer the population from the untrapped to the trapped state at the moment when the atoms enter the position above the mirror that corresponds to the final trap minimum. Although the Raman and optical pumping beams are always on during this loading sequence, this population transfer is spatially selective due to the high magnetic field gradient of the atom mirror which Zeeman shifts the states into resonance with the Raman beams only at a detuning given by the specific position corresponding to the final trap minimum. If the atom mirror is oriented upwards, than falling atoms that are not trapped during the first bounce could be trapped during subsequent bounces. A more involved, but perhaps ultimately more adaptable method would be to use microwires on the chip surface to guide a cold, dense cloud of atoms to the disk, which is how one might deliver a BEC to the ring trap.

Given a loading efficiency of 2%, a 10 μm radius ring trap will capture roughly 30-50 atoms from an uncompressed, dropped cloud of 10^7 atoms and temperature 10 μK . In order to capture more atoms, disks can be arranged in an array covering a larger surface area. The volume of the trap deeper than 200 μK is 1 to 2×10^{-9} cm^3 . Simulations indicate that these traps can be placed roughly 20 μm apart without significantly disturbing each other. Therefore, roughly 20% of the surface can be covered with the traps. Combining the loading efficiency with this surface coverage, roughly 10^3 atoms can be trapped. The leads can be routed though spaces between the disks with either a separate lead for each disk or a shared network of leads. Instead of disks, a pattern of concentric rings could also densely cover the surface.

Several undesired effects, such as heating, fragmentation of Bose-Einstein condensates, and a reduction of trap lifetimes have been detected in microtrap experiments involving atoms near room-temperature surfaces. The trap proposed here is not susceptible to heating due to technical noise on currents in microwires and to the fragmentation problems caused

by the spatial variation of these currents [93, 62, 97]. However, the trap remains susceptible to atom loss due to spin flips induced by magnetic field fluctuations from thermal currents in the metal forming the electric pads, as detected in several experiments [96, 193, 170]. Surface effects in this system will most closely resemble those in Lin *et al.*, wherein the skin depth for the transition frequency between trapped and untrapped magnetic sublevels of the atoms is much larger than both the distance of the atoms from the metal surface and the thickness of the metal conductor. As reported in Lin *et al.*, at a distance of $2\ \mu\text{m}$ this Johnson noise limits the lifetime of ^{87}Rb atoms above a $2\ \mu\text{m}$ thick copper conductor to a few 100 ms—ample time for detecting atoms in the ring trap. The metal film used for the electric disk pad in the ring trap will be ten to a hundred times thinner than that used for the above experiment, and we expect this to further minimize the trap’s loss rate [194, 195, 170, 196]. Specifically, the thickness of the disk can be adjusted between 50 to 200 nm without significant effect on the trap parameters. An additional surface effect was recently found in an experiment by Cornell’s group which measured a perturbation to the potential of a magnetic trap less than $30\ \mu\text{m}$ from a surface containing ^{87}Rb adsorbates [197]. Modification of the disk potential might be required to adjust for this perturbation.

9.6 Discussion

Several future improvements or extensions of this trapping concept are possible. For example, the decoherence effects due to the proximity of a conductor could be mitigated with the use of a dielectric magnetic film in place of the hard drive, and dielectric pads charged via an electron beam in place of the conducting disks. In addition, disks, rings, wires, and other shapes could be used to trap and manipulate the atoms just above the surface, and voltages adjusted to shift the atoms from one potential into another. Integration of these traps with magnetic microtraps based on current-carrying wires on the surface is also possible. Small single-atom traps with additional electrostatic pads to control the barrier heights could produce a system capable of performing quantum logic gates [13].

In the past several years, there has been much experimental and theoretical interest in trapped one-dimensional (1D) quantum degenerate gases (see References [198, 199, 200, 201, 202, 203, 204] and the citations within). Trapped 1D gases require $k_B T, \mu \ll \hbar\omega_\perp$, where T is the temperature, μ is the chemical potential, and ω_\perp is the transverse trapping

frequency. Various regimes of quantum degeneracy—of which a 1D gas of impenetrable bosons, the Tonks-Girardeau (TG) regime, is of particular interest—can be explored by changing the density of trapped atoms or by modifying the interactions between atoms via Feshbach resonances. In the latter case, a magnetic bias field for adjusting the s-wave scattering length, a , can be added parallel to the magnetization stripes of the atom mirror without affecting the potential of the magneto-electrostatic ring trap. In a 1-D trap, the chemical potential is

$$\mu = 2\hbar\omega_{\perp}an_{1D}, \quad (9.3)$$

where $n_{1D} = N/L$ is the number density and L is the length of the trap. With respect to ^{87}Rb , a common alkali used for BEC, $k_B T/\hbar\omega_{\perp}$ is smaller than 0.05 for temperatures below 100 nK. The TG regime requires that the mean interparticle separation, $1/n$, be much larger than correlation length, $l_c = (\hbar/2mn\omega_{\perp}a)^{1/2}$, where m is the atom's mass, $n = N/L$ is the number density [201]. This constraint limits the number of ^{87}Rb atoms in the ring trap to $N \ll 2m\omega_{\perp}aL/\hbar = 490$ [1600] atoms for a device of circumference $L = 2\pi \cdot 20 \mu\text{m}$, $\omega_{\perp} = 2\pi \cdot 40$ [$2\pi \cdot 130$] kHz, and an a unmodified by Feshbach resonances (the field at the trap minimum is ~ 12 G). Overcoming the challenge of detecting so few atoms may be possible through the incorporation of microwire traps [202].

The ring geometry adds a unique element to the many-body physics of the 1D trap. Josephson effects in trapped BECs have been investigated theoretically for the case of a double-well potential (see Reference [19] and citations within) and investigated experimentally in an optical standing wave [205]. A BEC in this magneto-electrostatic ring trap system with interspersed Josephson junctions formed from the addition of micron-sized perturbations to the trapping potential—such as those caused by wire leads, possibly tuned using additional pads—is reminiscent of superconducting electronic systems. The ratio of the chemical potential to the perturbation barrier height can be adjusted with the trap parameters such as d , r , atom number, and disk potential, as well as the use of additional electric pads, to cause the perturbation to act as either an impenetrable wall, a tunnel junction, or a scattering center. For example, taking the conceivable limits of the ring trap, to form a tunnel junction in a ring trap of $r = 5 \mu\text{m}$, $\omega_{\perp} = 130$ kHz, and containing 1000 Rb atoms, the perturbation would have to be adjusted slightly above $\mu = 47$ kHz which could be accomplished via the secondary electric pads described in Section 9.3. The utility

of this 1-D ring trap is highlighted by recent proposals for using a BEC in a double ring to create a SQUID-like device for neutral atoms [206] and for investigating quantum chaos in the system of the quantum kicked rotor [207]. Matter wave interferometry is one potential application.

This magnetoelectrostatic trap for cold neutral atoms—derived from balancing the repulsive force of an atom mirror with the attractive force from a charged disk—introduces a novel ring trapping geometry for cold neutral atoms. Fabrication of this trap is straightforward, and an array of such traps can trap a significant number of atoms. Furthermore, such a trap may allow the exploration of interesting many-body physics in a one-dimensional ring trap. This device is an example of the rich potential for developing novel atom optical elements through the integration of a hard drive atom mirror, charged pads, and microwires.

Chapter 10

Splitting a BEC in a Magnetic Double-Well Potential: Atom interferometry and Josephson effects on an atom chip

This chapter discusses an atom chip experiment whose goal is to coherently split a Bose-Einstein condensate (BEC) in a magnetic double-well potential. Such a device can be used for atom interferometry [208] or for studying Josephson effects in a matter wave system [209]. This experiment is being performed in collaboration with the group of Professors Theodore Hänsch and Jakob Reichel at the MPQ/LMU in Munich. The purpose of this chapter is to describe the experimental apparatus, present the simulations we performed regarding the Josephson effects observable in our specific double-well trap, and finally relate the non-ideality of the current chip design that obscures these effects. We will not discuss the physics of matter waves in a double-well potential in detail, but rather refer the reader to several papers on the topic (see References [210, 19] and citations within).

10.1 The double-well chip

The magnetic double-well potential in our experiment is formed from a pattern of rectilinear microwires on an atom chip. Figure 10.1 (a) shows a sketch of the relevant wires. The experiment begins with a BEC trapped above the intersection of wires #2 and #4. The BEC could be either condensed at this intersection point or waveguided to this location from the initial trapping zone. The BEC is confined using a dimple trap. Trapping in \hat{y} and \hat{z} is provided by the wireguide formed from current, I_4 , in wire #4 and a bias field, $B_{\hat{y}}$, in

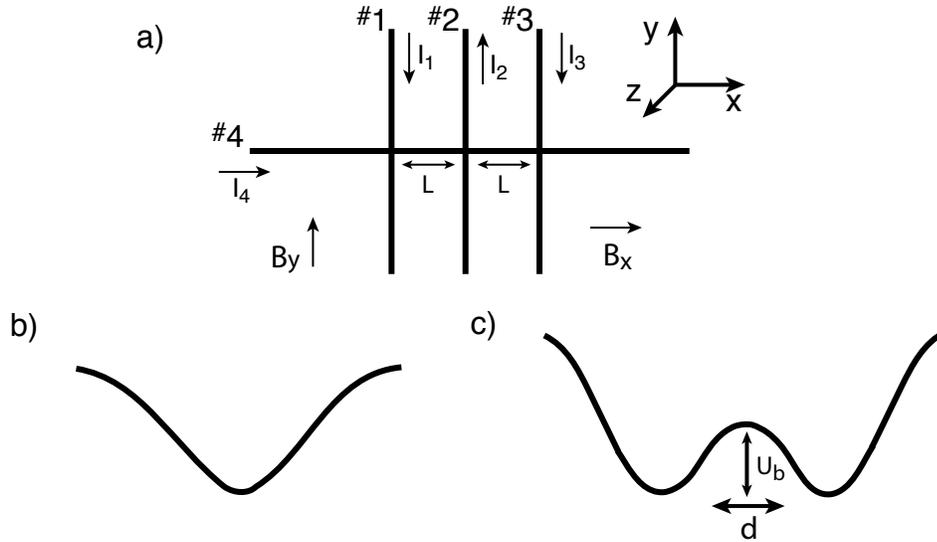


Figure 10.1: a) Microwire layout for producing a magnetic double-well potential. b) Single-well potential in \hat{x} formed with wire currents I_4 , I_1 , and I_3 with bias fields B_y in $+\hat{y}$ and B_x in $+\hat{x}$. The potential minimum is centered above wire #2. c) Double-well potential in \hat{x} formed from the same currents as in (b) but with the addition of I_2 . The peak of the barrier is centered above the wire with current I_2 and the wells are spaced by slightly less than $2L$. U_b is the magnitude of the barrier's energy, and d is the barrier's full width at half maximum.

$+\hat{y}$. Confinement in \hat{x} is provided by currents I_1 and I_3 in wires #1 and #3, respectively, combined with a bias field, $B_{\hat{x}}$, in $+\hat{x}$. The magnetic field from these wires and bias fields form an approximately harmonic potential well above the substrate. Panel (b) of Figure 10.1 shows a slice of the harmonic potential in the \hat{x} - \hat{z} plane. This potential is deformed into a double-well by the field from the current I_2 in wire #2. In contrast to the currents I_1 and I_3 , the oppositely flowing current I_2 produces a field that adds to $B_{\hat{x}}$ rather than subtracts from it. This produces a potential barrier in the harmonic potential well, splitting it into two as shown in Panel (c). In this manner a double-well potential is formed whose axis is along \hat{x} and is roughly cylindrically symmetric in the \hat{y} - \hat{z} plane. The barrier height, U_b , and width, d , is most readily adjusted with I_2 , but this height and the trap frequencies may also be tuned by manipulating the relative magnitudes of I_1 , I_3 , and $B_{\hat{x}}$. In general, $I_4 > [I_1, I_2, I_3]$ and $B_{\hat{y}} > B_{\hat{x}}$ so that the fields in \hat{x} may be viewed as perturbations on the otherwise homogenous waveguide formed from the I_4 and $B_{\hat{y}}$ ¹. To first order, it is the

¹This fails to be true in the case of the non-ideal wire layout of our chip (see Section 10.1.2). $B_{\hat{y}} < B_{\hat{x}}$ is required to overcome the stray fields produced by the current perturbations at the wire intersections.

2-D confinement due to I_4 and $B_{\hat{y}}$ that sets the height of the double-well trap above the substrate. The double-well potential, U , may be approximated by a harmonic potential in \hat{y} and \hat{z} with a Gaussian barrier along \hat{x} and centered at $x = 0$:

$$U(\rho, x) = m\omega_x^2(\lambda^2\rho^2 + x^2)/2 + U_b \exp(-x^2/2d^2), \quad (10.1)$$

where $\rho^2 = y^2 + z^2$, and $\lambda = \omega_p/\omega_x$. The λ parameter characterizes the degree to which the two lobes of the double-well are of a “pancake-shape” ($\lambda \ll 1$) or “cigar-shape” ($\lambda \gg 1$) [19].

When the barrier width is large and $U_b \gg \mu_c$, the two separated BECs are decoupled—tunneling is suppressed—and the phases evolve independently. An atom chip-based atom interferometer may be formed with this device by coherently splitting and recombining the BEC using the I_2 current to manipulate the barrier height [208]. Several other groups are also pursuing atom interferometry with BECs using microwire traps [211, 212, 73] and a few have demonstrated BEC splitting and the detection of interference fringes upon trap release [81, 213, 214]. (See Reference [215] and citations within for information regarding the optical dipole-based double-well traps.)

10.1.1 Josephson effects

Josephson effects can be explored in this double-well system when the barrier is no longer impenetrable, but rather allows weak coupling between the BECs in each lobe. The oscillation rate and atom number current amplitude are exponentially sensitive to the overlap of the BEC wavefunctions in this barrier region. Great care must be taken in designing the tunneling junction geometry to ensure that the Josephson plasma oscillation frequency, ω_{JP} , and the population oscillation amplitude, η_{max} , are large enough to be detected. We define this amplitude to be $\eta(t) = [N_l(t) - N_r(t)]/N$, where N_l and N_r are the atom population in the left and right wells, respectively, and N is the total number of atoms in the condensate. Reference [19] provides a rough estimate of the maximum amplitude of the oscillation:²

$$\eta_{max} \approx \frac{\exp[-\sqrt{md^2(U_b - \mu_c)/4\hbar^2}]}{(ma^2\lambda^4 N^2 \omega_x/\hbar)^{1/10}}. \quad (10.2)$$

To maximize the exponential in η_{max} , d should be minimized and U_b —while always having to be larger than μ_c —should be as close to the chemical potential as possible. The number

²Note: References [19] and [210] use different and sometimes contradictory notations.

current may be further maximized by decreasing m , the atomic mass; a , the s -wave scattering length; λ ; and N . The atom species is usually fixed, though a might be modifiable by Feshbach resonances, so λ and N remain as the most easily adjustable parameters for maximizing η_{max} . The λ ratio is minimized with the ‘‘pancake’’ geometry which increases the junction area. A large condensate population, N , exacerbates the mean-field interaction which inhibits tunneling. Quantum and thermal phase fluctuations between the BECs in the two wells must also be considered when designing the double-well potential to ensure the coherence of the system [210].

Following the method of the Williams paper [19], we have calculated η_{max} and ω_{JP} for the double-well potential produced by our atom chip. Matlab scripts were written to solve the necessary equations listed in this paper. These scripts are quite lengthy and have not been included in Appendix B. Please contact the author to obtain a copy of this Matlab code. This system is described by the Gross-Pitaevskii equation,

$$i\hbar \frac{\partial}{\partial t} \Phi(\mathbf{r}, t) = \left[-\frac{\hbar^2 \nabla^2}{2m} + V_{ext}(\mathbf{r}) + g|\Phi(\mathbf{r}, t)|^2 \right] \Phi(\mathbf{r}, t), \quad (10.3)$$

where $|\Phi(\mathbf{r}, t)|^2$ is the condensate density, $V_{ext}(\mathbf{r}) = U(\rho, x)$ for this double-well, and $g = 4\pi\hbar^2 a/m$ is related to the s -wave scattering length and is the coupling constant characterizing the strength of the non-linear term arising from the interacting atoms [216]. It is this mean-field interaction that limits $\eta_{max} < 1$ and introduces into this system the phenomenon of quantum ‘‘self-trapping’’ wherein tunneling is suppressed for initial population differences larger than η_{max} [19]. The Josephson and the self-trapping effects were recently observed experimentally in a double-well experiment using optical dipole traps [209].

An approximate time-dependent solution to this double-well system is found in the case of $U_b > \mu_c$ by solving the Gross-Pitaevskii equation using a two-mode ansatz [217]. Specifically, the Gross-Pitaevskii dynamics in the double-well may be described by a superposition of wavefunctions localized in the left (l) and right (r) wells:

$$\Phi(\mathbf{r}, t) = \psi_l(t)\Phi_l(\mathbf{r}) + \psi_r(t)\Phi_r(\mathbf{r}), \quad (10.4)$$

where for case of the left well $\psi_l(t) = \sqrt{N_l(t)} \exp(i\phi_l(t))$ and ϕ_l is the phase of the left-localized condensate. A variational procedure using a Gaussian ansatz for the symmetric

and anti-symmetric combinations of Φ_l and Φ_r provides an accurate description of the system (see Reference [19] and the citations within). These wavefunctions are found by minimizing the Gross-Pitaevskii energy functional:

$$E[\Phi] = \int d\mathbf{r} \left[\frac{\hbar^2}{2m} |\nabla\Phi|^2 + U|\Phi|^2 + \frac{g}{2} |\Phi|^4 \right]. \quad (10.5)$$

Once Φ is found, $\eta(t)$ and $\phi(t)$ —and consequently η_{max} and ω_{JP} —are obtained by substituting Φ back into the Gross-Pitaevskii equation. This is the procedure we implemented numerically for our specific double-well potential in the Matlab code discussed previously. Using the notation of Reference [210], the solutions to Equation 10.3 may be used to find the “capacitive-like” energy, E_C , due to the atomic interactions, and the tunneling energy, E_J :

$$E_C = 2 \frac{d\mu_i}{dN_i}, \quad (10.6)$$

$$E_J = \frac{\hbar^2}{m} \int dy dz \left[\Phi_l \frac{\partial \Phi_r}{\partial x} - \Phi_r \frac{\partial \Phi_l}{\partial x} \right]_{x=0}, \quad (10.7)$$

where i equals either l or r (i.e. μ_i is the chemical potential of the BEC in the left well), and this should be evaluated for N divided equally between the two wells. In the Thomas-Fermi limit, $E_C = (4/5)\mu_i/N_i$ [216]. In the limit of strong tunneling, $\omega_{JP} \approx \sqrt{E_C E_J}/\hbar$ and to have quantum phase fluctuations be small, the condition $E_C/E_J \ll 1$ must be satisfied. The BEC temperature should be kept lower than approximately E_J to negate the effects of thermal fluctuations. See Reference [210] for details regarding these issues.

10.1.2 The magnetic double-well

Before we can solve for η_{max} in our system, we must first calculate and characterize the magnetic double-well potential produced by the actual wire layout on our atom chip. A photograph of the chip wires is in Figure 4.5 of Chapter 4. It is evident that the wires are not perfectly thin like those in the model of Figure 10.1. This is due to the fact that the wires must be able to support several hundred mA of current for 100 ms or more. In particular, the central wire must support a current $I_4 \gtrsim 0.5$ A. The actual device, shown in Figure 4.5, has five splitting wires instead of the three shown in Figure 10.1. This was done to add redundancy, and to allow the creation of a larger barrier width for atom interferometry by

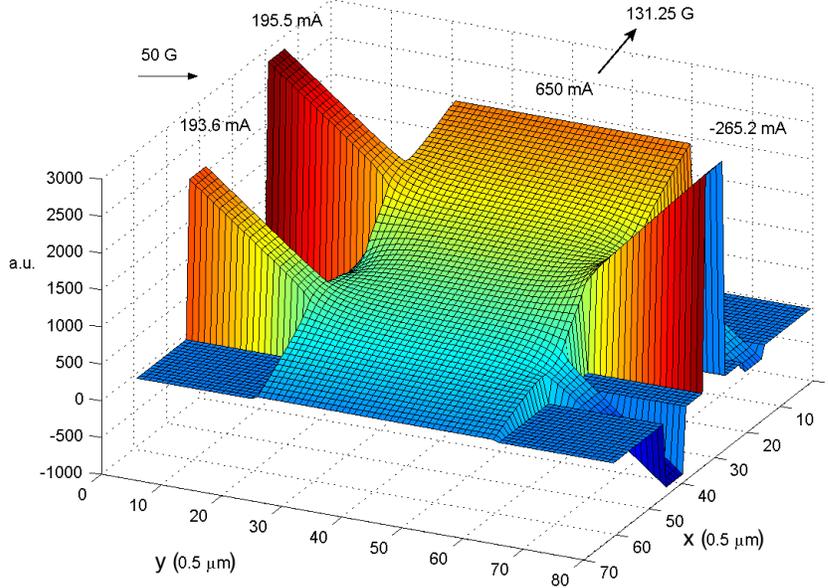


Figure 10.2: Potential on the wires of the double-well atom chip.

using the two or three central wires for the I_2 current. The wires are $4 \mu\text{m}$ tall, and the guiding wire, #4 with I_4 , is of width $20 \mu\text{m}$, and the three (or possibly five) splitting wires, #1, #2, and #3 with I_1 , I_2 , and I_3 , are $2\text{-}3 \mu\text{m}$ wide and spaced by $L = 2 \mu\text{m}$.

A Laplace solver was written in Matlab to calculate the actual current flow in this pattern of finite-width wires. This calculation was undertaken because it was only after the fabrication of the chip that we realized that the non-ideal current flow at the intersection of the wires might perturb the formation of the double-well potential. Specifically, since the guiding wire is so much wider than the splitting wires, the current from the splitting wires flows into the guiding wire rather than straight across: The current makes a 90° turn from \hat{y} to \hat{x} which produces undesirable stray magnetic fields at the wire intersection. Figures 10.2 and 10.3 show the electric potential and current flow in our wire pattern (only three splitting wires are considered). The deviation from \hat{y} for the currents in the splitting wires would not affect the trap potential if the atoms were confined further away from the surface than the width of the guiding wire, $20 \mu\text{m}$. However, for this experiment the atoms need to be within $20 \mu\text{m}$ from the surface to ensure the proper formation of a double-well potential. The double-well trap maximizes η_{max} and ω_{JP} the closer the trap minimum is to the wire surface (i.e. a smaller d and λ can be attained when the trap minimum is closer to the chip). While the field from the segments of the splitting wires that effectively extend to

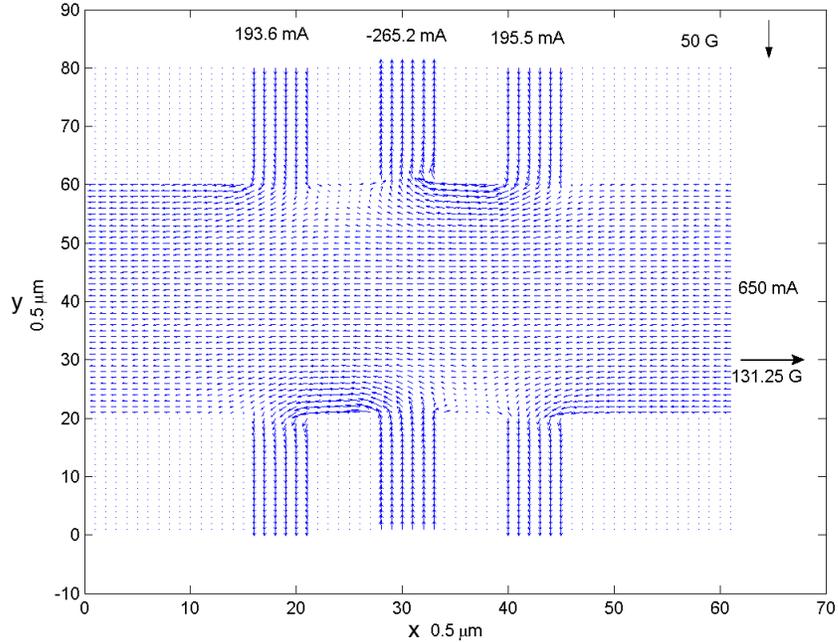


Figure 10.3: Currents flowing through the wires of the double-well atom chip.

$+\infty \hat{y}$ and $-\infty \hat{y}$ allows the formation of a double-well potential even at these small trap heights, the distortion of the current flow in the guide wire intersection skews the axis of the double-well away from \hat{x} . This prevents the attainment of as low a d and λ as we would expect from the ideal case of non-intersecting wires.

We wrote a 3-D Biot-Savart solver in Matlab for computing the magnetic field from the current pattern found from the numerical solutions to the Laplace equation. Future investigations should use a commercial software package for doing this more efficiently. The following figures show the 2-D slices of the magnetic field forming the double-well in the \hat{x} - \hat{z} plane taken along the axis of the double-well trap (Figure 10.4); in the \hat{x} - \hat{y} plane at z_{min} , the minimum of the trap in \hat{z} (Figure 10.5); and in the \hat{y} - \hat{z} plane at the center of a well (Figure 10.6).³ In this experiment, the mirror coating on the atom chip is roughly $15 \mu\text{m}$ tall, and this prevents the trap minimum from being much less than $\sim 20 \mu\text{m}$ from the wires if surface effects are to be minimized: in these calculations $z_{min} = 20 \mu\text{m}$ ⁴.

We found that currents of $I_1 = -193.6 \text{ mA}$, $I_2 = 265.2 \text{ mA}$, $I_3 = -195.5 \text{ mA}$, and $I_4 = 650 \text{ mA}$ and bias fields of $B_{\hat{y}} = 50 \text{ G}$ and $B_{\hat{x}} = 131.25 \text{ G}$ maximize η_{max} and ω_{JP} for

³The origin of these axes do not correspond to those in Figures 10.2 and 10.3. However, the axis labels are consistent throughout Figures 10.4, 10.5, and 10.6.

⁴With respect to the \hat{z} origin in these plots, the position of the wires is $z \approx -18 \mu\text{m}$.

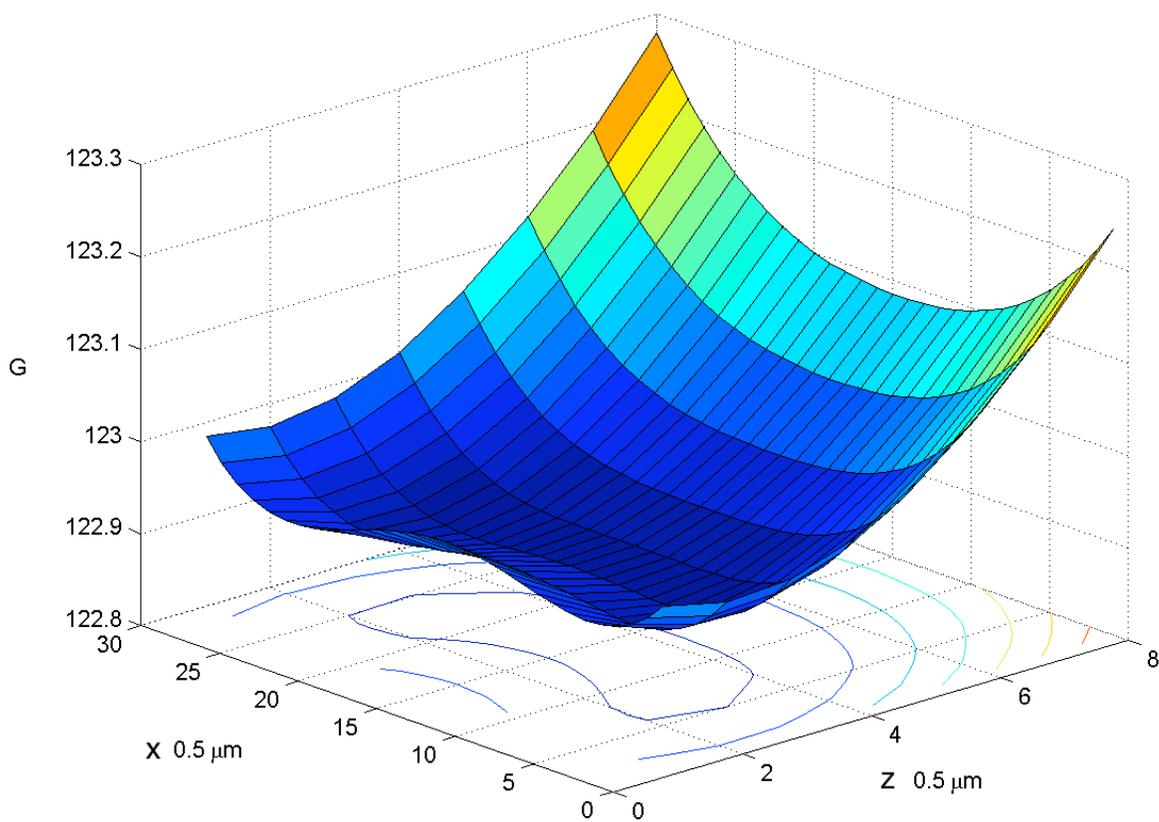


Figure 10.4: Magnetic field in the \hat{x} - \hat{z} plane taken along the axis of the double-well trap. The field is produced from the wire currents and bias fields shown in Figures 10.2 and 10.3. The barrier height is almost imperceptible due to the effort to make $U_b = 1.5\mu_c$.

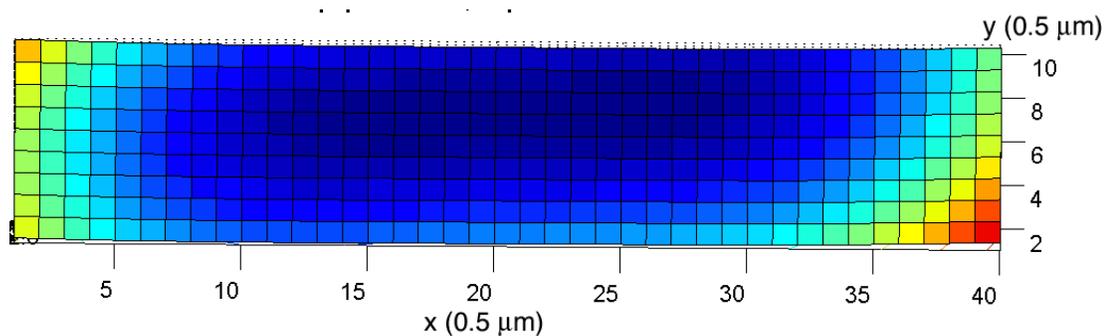


Figure 10.5: Magnetic field in the \hat{x} - \hat{y} plane at $z_{min} = 20 \mu\text{m}$ above the wires. The field is produced from the wire currents and bias fields shown in Figures 10.2 and 10.3. The barrier height is almost imperceptible due to the effort to make $U_b = 1.5\mu_c$. Darker regions (blue) show a smaller magnetic field than the lighter (red) regions.

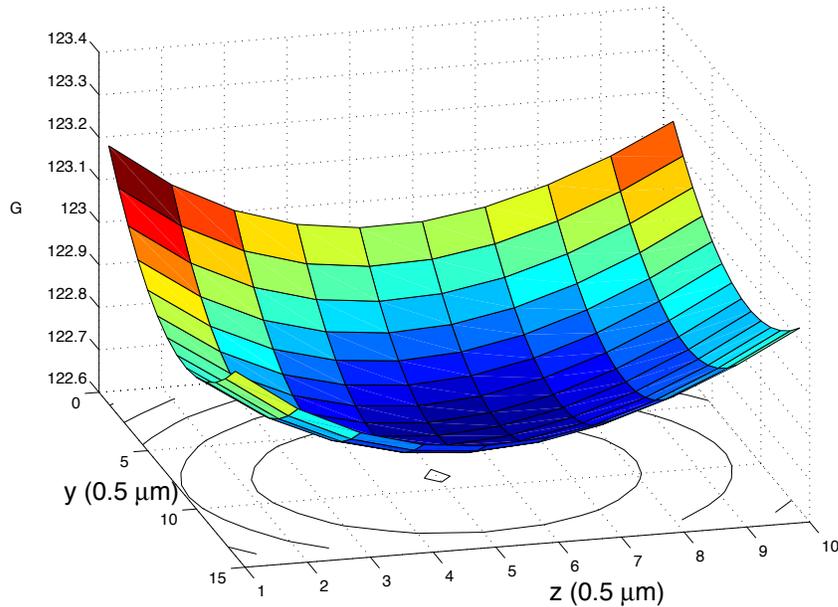


Figure 10.6: Magnetic field in the \hat{y} - \hat{z} plane at the center of a well. The field is produced from the wire currents and bias fields shown in Figures 10.2 and 10.3.

our actual wire pattern and a trap minimum at $z_{min} = 20 \mu\text{m}$. We fit these 2-D slices of the magnetic field to find the trap frequencies, ω_x , ω_y , and ω_z , and d and U_b .⁵ These values are inserted into the Gross-Pitaevskii equation solver, which then allows the calculation of η_{max} and ω_{JP} for our trap parameters. We found $\eta_{max} = 1\%$ and $\omega_{JP} = 2\pi \cdot 14 \text{ Hz}$ for a double-well system of $N = 1000$, $U_b = 1.5\mu_c$, $z_{min} = 20 \mu\text{m}$, $\omega_x = 1.8 \text{ kHz}$, $\omega_\rho = 3 \text{ kHz}$, and $d = 6.5 \mu\text{m}$. If we had used a different atom chip mirror that allowed the trap minimum to be closer to the surface, then for $z_{min} = 8 \mu\text{m}$ we could have improved the amplitude and frequency of the Josephson oscillations to $\eta_{max} = 1.5\%$ and $\omega_{JP} = 2\pi \cdot 72 \text{ Hz}$ with a double well system of $N = 1000$, $U_b = 1.5\mu_c$, $z_{min} = 8 \mu\text{m}$, $\omega_x = 5.2 \text{ kHz}$, $\omega_\rho = 8 \text{ kHz}$, and $d = 3.8 \mu\text{m}$. The η_{max} is not spectacular for either of these cases, and number current oscillations of this amplitude would be hard to detect experimentally. The skewing of the double-well and the large trap height are the main limiting factors in this experiment. It would be difficult to shrink the wire spacing, L , by much more than a factor of two, and so large improvements to d would be hard to realize. In both of the above cases, $\lambda \approx 1.6$

⁵ ω_y and ω_z are approximately equal, and we take ω_ρ to be their average for solving the Gross-Pitaevskii equation.

which is not yet into the high junction area, “pancake” trap regime which would greatly improve the η_{max} .

Experimentally, we have been able to create and trap BECs above the wire intersections used for the double-well trap. A cold cloud of ^{87}Rb atoms are collected in a Z-trap — after a mirror MOT, U-MOT, and P-trap (see Chapter 2 for loading details)—before being waveguided to the double-well region. The waveguiding is accomplished by RF-cooling the trapped atoms before transferring them from the Z-trap to a long wireguide. The guide wire, #4 in Figure 10.1, delivers the atoms to the intersection of wires #4 and #2. As the atoms pass down the guide wire with current $I_4 \approx 2$ A, the wires with I_1 and I_3 form a dimple trap that confines the cold atoms in the splitting region. Final BEC production is performed by RF evaporation at this location. Unfortunately, an unintentionally large current in the guide wire coupled too much current into the splitting wires and overheated and broke some of them before we were able to attempt a double-well trap.

The next generation of the double-well atom chip takes into account the Josephson junction engineering intuition we have gained by performing the calculations presented above. Three improvements will be made: the atom chip will be bilayer; the splitting wire spacings, L , will be smaller; and the detection system improved to allow a lower BEC population, N . The small splitting wires will be on a layer separated from the large guiding wire by a $10\ \mu\text{m}$ thick insulating layer of polyimide. Since the splitting wires and the guide wire no longer intersect, we will not have to worry about either burning-out the splitting wires or the stray magnetic field from the non-ideal current flow at the intersection of the wires. A smaller trap population, N , increases the maximum atom current oscillation between the wells (see Equation 10.2 for η_{max}). With a similar wire pattern and a trap height of $z_{min} = 7.8\ \mu\text{m}$, we should be able to produce a double-well trap of $d = 3.2\ \mu\text{m}$, $\omega_x = 0.5$ kHz, $\omega_y = 1.8$ kHz, and $\omega_z = 2.1$ kHz. While λ is larger in this trap than before, by using only $N = 50$, we can achieve a $\eta_{max} = 30\%$ while keeping the quantum phase fluctuations below 0.05 (see Reference [210] for information regarding this latter quantity). A major challenge will be in ensuring that ambient magnetic field fluctuations or current supply noise do not jiggle the double-well trap and dephase the oscillations. Low-noise, custom power supplies are being employed and μ -metal shielding installed to help mitigate these unwanted effects.

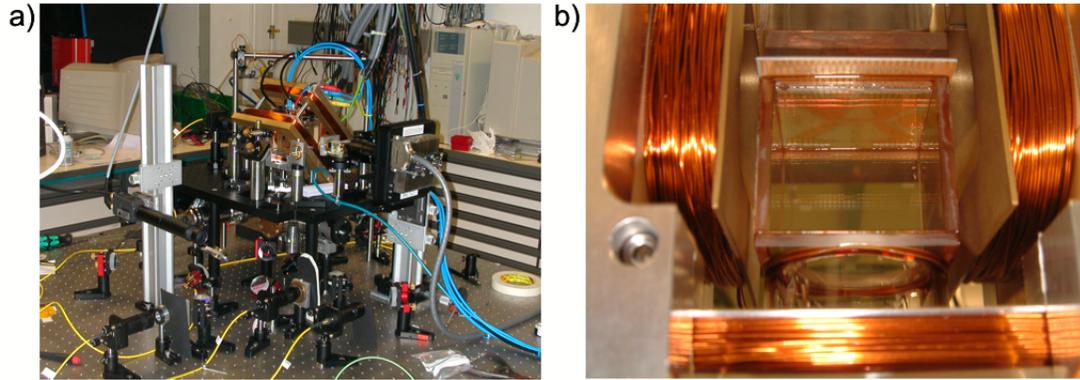


Figure 10.7: a) Glued-cell vacuum chamber for atom chip BEC production. The vacuum chamber piping and ion pump are obscured by the black breadboard. b) Close-up of the glued-cell mounted inside the magnetic field coils. Notice the extension of the base chip outside of the vacuum chamber.

10.2 The glued-cell chamber

The high gradients produced by magnetic microtraps evaporatively cool atomic gases to degeneracy more rapidly than in more standard magnetic traps. Consequently, the vacuum pressure in the trapping chamber can be higher since the trap lifetime doesn't need to be as long. This allows one to relax the vacuum chamber design requirements, thereby enabling the use of a much more compact BEC production system [17]. A major practical difficulty of atom chip experiments is the connection of typically 20 to 40 electrical contacts from the chip to the outside of the vacuum chamber without blocking optical access or introducing incompatible vacuum materials. Fortunately, the less stringent vacuum requirements allow us to solve all these problems in a relatively elegant manner by allowing the use of relatively vacuum-safe epoxy: the atom chip substrate itself forms one of the walls of the vacuum chamber by being glued to an open face of a glass cell. This enables great miniaturization and forms a natural way to create an electrical feedthrough since the edge of the chip can extend outside the vacuum chamber. The entire volume of the vacuum chamber can fit inside a roughly 1 m^3 space. Combined with a compact diode laser system, this forms a space-efficient apparatus for BEC production (see Figure 10.7 (a)).

Figures 7.3 and 10.7 (b) show this chip assembly. Attached to a standard UHV (ultra-high vacuum) vacuum system—comprised of a valve to a turbo pump, a compact ion pump, titanium sublimator, and a vacuum gauge—is a glass-to-metal seal built on a 2.75" Conflat

piece. The glass part is an open-ended, roughly 1" diameter cylinder. This is pictured in the lower half of Figure 7.3. The face of the top edge of the cylinder is polished in-house to enable smooth mating with the face of a glass cell glued on top of it. The cubic glass cell is purchased from Hellma for about \$50. We order it without one of the glass faces, and with a diamond-tipped hole-cutter we drill-out a hole in the face that mates with the glass-to-metal seal. The polished edge of the glass cylinder is glued to the face with the hole and the atom chip is glued to the polished edge of the missing face of the Hellma cell. The final assembly is shown in Figure 7.3. The dark amber areas are the glue seals. We use the UHV compatible glue, Epotek 353. This epoxy cures at 80° C for 30 min and makes a good vacuum seal at these joints without too much out-gassing. We have been able to achieve pressures as low as 1.5×10^{-10} Torr in these glued-cell chambers. The cells are reasonably strong and can withstand a $\sim 150^\circ$ C chamber bake. However, we have noticed that they are prone to develop micro-fissures at the glue joints after 6 months to a year under vacuum. Nearly all of these cracks have been satisfactorily sealed with the addition of small glue patches.

The vacuum cell with chip occupies a cubic volume of only ~ 30 mm on each side. This allows the close placement of magnetic bias field coils, which in turn minimizes the power dissipation required to produce a field of tens of Gauss at the atom trap position. The atom chip is larger than the face of the Hellma cell, and the part extending outside this cell is used as the connection point for the chip's microwires. The left panel of Figure 7.4 and Panel (b) of Figure 10.7 show how this is done: the metal pads on the chip extend from inside the vacuum to the air, and the Epotek 353 forms a tight seal over and between these several micron tall wires and the glass Hellma cell. A female PCI slot from a computer is used to couple the wire pads to macroscopic wires. This connector is arranged so as not to block optical access to the trapped atomic cloud. On the top face of the atom chip—which is exposed to air—a water-cooled copper block, thermistor, and thermoelectric cooler are attached to cool and temperature stabilize the atom chip. The atom chip that is glued to the cell may not necessarily be used as the primary chip device. In our experiments, a secondary chip that contains the smallest microwire features is glued to this “base chip.” This secondary atom chip is smaller than the surface area of the cell's face and is entirely contained within the vacuum chamber. Figure 10.8 shows these two chips glued together. The wire pads from the smaller, main atom chip are connected via wire bonds to the lower

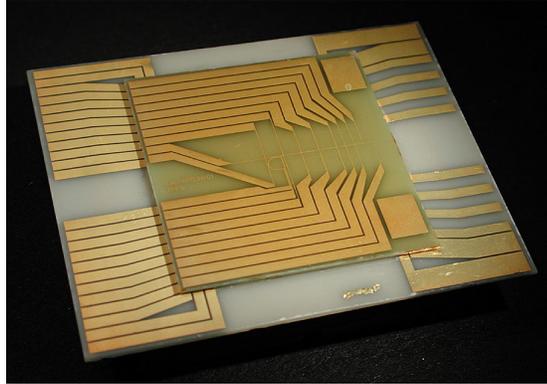


Figure 10.8: Atom chip glued to “base chip” with the thermally conductive epoxy, Epotek H77. The top chip (roughly 30×30 mm) contains the microwires, and these wires are wire bonded down to the contact pads on the base chip (wire bonds not shown). The base chip also contains a several hundred micron wide wire for creating a U-MOT. This wire is located underneath the upper atom chip. The wire pads of the base chip shown in the left and rightmost extremes of the picture extend outside the vacuum and form the electrical feedthroughs

base chip. Up to 10 wire bonds are required per pad for redundancy in case of bond breaking and to enable the conductance of up to an amp of current. To improve thermal conductance between the two chips, a thin layer of thermally conductive epoxy, Epotek H77, attaches the chips to one another. This epoxy is not as UHV compatible as the Epotek 353 and care should be taken to minimize its vacuum-exposed surface area.

Atom chip and base chip assembly. The two AlN substrates are glued together with thermally conductive epoxy (Epotek H77). The wire pads of the base chip shown in the left and rightmost extremes of the picture extend outside the vacuum and form the electrical feedthroughs.

The cell walls and the atom chip’s surface are coated with a dielectric film that forms an anti-reflection (AR) coating and a mirror, respectively. The film is of the detachable kind described in Chapter 7, Section 7.2.2. The coating is glued to the atom chip and cell wall with the Epotek 353 epoxy, which in thin films is transparent in the NIR. This is a nice technique as it enables the AR coating of hard-to-reach places such as the inside of the vacuum cell.

In the interest of compactness, we trap atoms in the initial mirror MOT using a vapor of ^{87}Rb atoms in the glass cell that is attached to the atom chip. The most difficult aspect of BEC production in the glued-cell chamber is the balance between large trap

atom number and long trap lifetime. Although evaporative cooling occurs more rapidly in the high gradients of the atom chip, one still needs to start with around 10^7 atoms so that at the end of the cooling cycle, which takes 5-6 seconds, one has more than 10^5 remaining. A high vapor pressure of ^{87}Rb during MOT loading is required for high trap populations, but during the magnetic confinement thereafter, the vapor pressure should be low to eliminate the background collisions that quench trap lifetime. These seemingly contradictory requirements are satisfied by time-dependently controlling the vapor pressure. This is accomplished by the combined use of a Rb dispenser [218, 219] and a UV desorption light [220].

The dispenser is the source of the Rb vapor pressure and operates by resistively heating a Rb compound to several hundred degrees Celsius. The Rb compound is contained in a metal jacket a centimeter or two long that has a thin ~ 1 mm wide slit. When ≥ 2.5 A of current flows through the metal jacket the Rb compound is heated and pure Rb is released. The Rb vapor pressure can be precisely tuned with the dispenser current. The Rb dispenser and Kapton-coated wire contacts are shown in Figure 7.3. The metal jacket is positioned vertically with the slit facing towards the opposite side of the glass cylinder. It is important that the slit not face the atom chip surface or else hot atoms will bombard the MOT and magnetic trap. The dispenser must be properly conditioned before operation to prevent “dirt” from being emitted along with the Rb. The general procedure is to heat the dispenser with higher than normal currents (~ 5 A) that one increases every few hours. The vacuum pressure spikes each time the current steps to a higher value, but reduces after a half-hour as the dirt is removed from the dispenser. References [83, 218, 219] contain a more thorough discussion of this procedure.

A UV source is flashed-on during the MOT loading phase to transiently increase the vapor pressure of Rb in the vacuum cell which increases the MOT atom number. The UV light desorbes Rb from the glass cell walls, allowing a factor of 5 to 20 increase in MOT number depending on the desorption efficiency and ambient vapor pressure of Rb (i.e. there is a larger fractional trap population increase for lower initial vacuum pressures). Certain types of glass work better than others for absorbing Rb, and it seems like Pyrex is superior to quartz and fused silica (the group of J. Thywissen at the University of Toronto has been quite successful with the technique using Pyrex). Generally, the first flash of the UV light after a prolonged pause in the experimental cycle desorbes more Rb than in

subsequent flashes. Halogen bulbs from overhead projectors, though white light sources, contain enough UV for this purpose. Unfortunately, they become quite hot and take a few seconds to completely turn-off when the drive current is shunted to ground. The heat of the bulb can disturb the glued cell and thermally detune cavities within the cell. Royal blue and UV LEDs are now on the market and can be used instead. The UV LED is said to be superior even though it only emits a few tens of mW. We have used the 50 W UV light source from the Dymax UV curing set-up for this desorption to great effect. Unfortunately, this 50 W source is quite expensive and prone to breakdown.

We use the Rb dispenser at a very low current, 2.5 to 2.8 A, to introduce a constant, low level of Rb into the chamber. The UV light source is turned on for the first ~ 5 s of the ~ 8 s MOT loading sequence, during which the added Rb increases the vacuum pressure to the high 10^{-10} to low 10^{-9} Torr range. The UV is off during the last 3 s to allow the chamber to decrease back down to a steady-state vacuum pressure in the low 10^{-10} Torr range. After the MOT phase, the magnetic traps are loaded and the combination of high initial atom number and low chamber pressure allows evaporative cooling to quantum degeneracy without losing all of the trapped atoms. Some groups only run the dispenser once a day at a high current to coat the cell walls instead of continuously running it at low current. Care must be taken to balance the dispenser current and the UV light source operation time and bulb position to optimally provide high atom number and long lifetimes. We place at least two Rb dispensers into the chamber in case one is extinguished. Their lifetime at normal operating currents and conditions seems to be greater than one year. Future experiments will use a chilled copper block to cool the Rb dispenser rapidly after shutting off the current, thereby allowing the Rb emission to be quickly quenched. This will allow the pulsed operation of the dispenser as well as the UV light.

Appendix A

Magnetic Fields, Gradients, and Trap Minima of U- and Z-Traps

This appendix lists the analytic expressions for the field calculation of the U-trap. All wires are assumed to be infinitely thin, and this approximation breaks down only when the trap minimum, y_0 , is less than a wire width, w , above the substrate surface [2, 50, 68]. The magnetic field at a distance y from the finite-width wire is:

$$B(y) = \frac{\mu_0 I}{\pi w} \left(\frac{\pi}{2} - \arctan \frac{2y}{w} \right). \quad (\text{A.1})$$

This reduces to

$$B(y) \approx \frac{\mu_0 I}{\pi w} \left(\frac{\pi}{2} - \frac{2y}{w} \right), \quad (\text{A.2})$$

for $y \leq w$. It is important to note that because the near-wire gradient is

$$\nabla B(y) = -\frac{\mu_0}{2\pi} \frac{1}{y^2 + (w/2)^2}, \quad (\text{A.3})$$

the wire width must be decreased to achieve arbitrarily high gradients for reasonable wire currents. An analytic expression exists for a wide wire, but is too cumbersome to list here and numerical simulations should be performed instead. Moreover, in these near-wire situations, numerical calculations involving the Laplace equation are often required to account for the actual current flow between wire intersections. The double-well experiment discussed in Chapter 10 was one such case. Although we have written MATLAB code for these numerical simulations, we recommend the use of FEMLAB for all future calculations due to its flexibility and optimized coding. For quick simulations of the U-traps, Z-traps,

and microwire waveguides, the following analytic expressions for the fields are sufficient and allow the fast searching of parameter space. Appendix B, Section B.3 contains MATLAB code incorporating the expressions for easily running simulations of arbitrary combinations of rectilinearly arranged U-traps, Z-traps, dimple traps, microwire waveguides, and external quadrupole and single coil fields.

A.1 Infinitely thin linear wires

We first list expressions for the fields from wires that form the building blocks for U-traps, Z-traps, and microwire waveguides with gates: infinite wires, wire segments, and half-infinite wires. All coordinates are as they appear in Figure 5.4. The field from an infinite straight wire in \hat{z} :

$$\mathbf{B}_{InfWire} = \frac{\mu_0 I}{2\pi} \frac{1}{x^2 + y^2} [-y\hat{x} + x\hat{y}]. \quad (\text{A.4})$$

The integral for finding the field from a wire segment (given here with the wire extended along \hat{x}) is

$$\frac{\mu_0 I}{4\pi} \int \frac{[y\hat{z} + z\hat{y}] dx'}{[(x - x')^2 + y^2 + z^2]^{\frac{3}{2}}} = \frac{\mu_0 I}{4\pi} \frac{x' - x}{\alpha \sqrt{\alpha + (x - x')^2}}, \quad (\text{A.5})$$

where $\alpha = y^2 + z^2$. For a wire segment symmetrically positioned about $x = 0$ and of length L , this expression becomes:

$$\mathbf{B}_{SymSegment} = \frac{-\mu_0 I}{4\pi\alpha} \left[\frac{L_-}{\sqrt{4\alpha + L_-^2}} + \frac{L_+}{\sqrt{4\alpha + L_+^2}} \right] [y\hat{z} + z\hat{y}], \quad (\text{A.6})$$

where $L_{\pm} = L \pm 2x$. For the U- and Z-trap we need to have expressions for a half infinite wire from both 0 to $+\infty$ and $-\infty$ to 0. For wires positioned along \hat{z} , we have for the $+\infty$ case:

$$\begin{aligned} \mathbf{B}_{+\infty} &= \frac{\mu_0 I}{4\pi} \left[\frac{\beta}{(x - x')^2 + y^2} \right] [-y\hat{x} + (x - x')\hat{y}], \\ \beta &= 1 + \frac{z}{\gamma}, \\ \gamma &= \sqrt{(x - x')^2 + y^2 + z^2}. \end{aligned} \quad (\text{A.7})$$

For the $-\infty$ case:

$$\mathbf{B}_{-\infty} = \frac{\mu_0 I}{4\pi} \left[\frac{-y\hat{x} + (x - x')\hat{y}}{\gamma^2 + z\gamma} \right]. \quad (\text{A.8})$$

The U-trap, centered at $x' = 0$ as in Figure 5.4, is formed in the following manner:

$$\mathbf{B}_U = \mathbf{B}_{+\infty}(x' = -L/2) + \mathbf{B}_{+\infty}(x' = L/2) + \mathbf{B}_{SymSegment} + B_{bias}\hat{z}. \quad (\text{A.9})$$

Similarly, the Z-trap is

$$\mathbf{B}_Z = \mathbf{B}_{+\infty}(x' = -L/2) + \mathbf{B}_{-\infty}(x' = L/2) + \mathbf{B}_{SymSegment} + B_{Zbias}\hat{z} + B_{Xbias}\hat{x}, \quad (\text{A.10})$$

where the last term is commonly used to manipulate the trap curvature and minimum field.

A.2 Force on an atom in a U-trap

The force on an atom with a magnetic moment, \mathbf{m} , is

$$\begin{aligned} \mathbf{F} &= (\mathbf{m} \cdot \nabla)\mathbf{B} \\ &= m_x \frac{\partial \mathbf{B}}{\partial x} + m_y \frac{\partial \mathbf{B}}{\partial y} + m_z \frac{\partial \mathbf{B}}{\partial z}. \end{aligned} \quad (\text{A.11})$$

Let us assume that away from the center of the trap the atom's spin always follows the magnetic field:

$$m_i = g\mu_0 B_i / \sqrt{B_x^2 + B_y^2 + B_z^2}. \quad (\text{A.12})$$

Near the center, where the field vanishes, the atom's spin no longer follows the field, but let us write the force as

$$F_i \propto \frac{\partial B_i}{\partial x} + \frac{\partial B_i}{\partial y} + \frac{\partial B_i}{\partial z}. \quad (\text{A.13})$$

The relevant partial derivatives for the U-trap field expressed in Equation A.9 are listed below. Though messy, we write them here due to their importance in the fitting routine used in Section 5.2.1.

$$\begin{aligned} A_{\pm} &= x \pm L/2, \\ \beta_{\pm} &= 1 + \frac{z}{\sqrt{A_{\pm}^2 + y^2 + z^2}}, \end{aligned}$$

$$\begin{aligned}
\mathbf{B}_{Utrap} &= \frac{\mu_0 I}{4\pi} [B_x \hat{x} + B_y \hat{y} + B_z \hat{z}], \\
B_x &= \frac{y\beta_-}{A_-^2 + y^2} - \frac{y\beta_+}{A_+^2 + y^2}, \\
B_y &= \frac{2z}{y^2 + z^2} + \frac{A_+\beta_+}{A_+^2 + y^2} - \frac{A_-\beta_-}{A_-^2 + y^2}, \\
B_z &= \frac{4\pi B_{bias}}{\mu_0 I} - \frac{2y}{y^2 + z^2}, \\
\frac{\partial B_x}{\partial x} &= y \left[\frac{\partial \beta_-}{\partial x} \frac{1}{A_-^2 + y^2} - \frac{\partial \beta_+}{\partial x} \frac{1}{A_+^2 + y^2} - \frac{2\beta_- A_-}{(A_-^2 + y^2)^2} + \frac{2\beta_+ A_+}{(A_+^2 + y^2)^2} \right], \\
\frac{\partial B_y}{\partial x} &= \frac{\partial \beta_+}{\partial x} \frac{A_+}{A_+^2 + y^2} - \frac{\partial \beta_-}{\partial x} \frac{A_-}{A_-^2 + y^2} + \frac{\beta_+}{A_+^2 + y^2} - \frac{\beta_-}{A_-^2 + y^2} + \frac{2A_-^2 \beta_-}{(A_-^2 + y^2)^2} - \dots \\
&\quad - \frac{2A_+^2 \beta_+}{(A_+^2 + y^2)^2}, \\
\frac{\partial B_x}{\partial y} &= \frac{\partial \beta_-}{\partial y} \frac{y}{A_-^2 + y^2} - \frac{\partial \beta_+}{\partial y} \frac{y}{A_+^2 + y^2} + \frac{\beta_-}{A_-^2 + y^2} - \frac{\beta_+}{A_+^2 + y^2} + \frac{2y^2 \beta_+}{(A_+^2 + y^2)^2} - \dots \\
&\quad - \frac{2y^2 \beta_-}{(A_-^2 + y^2)^2}, \\
\frac{\partial B_y}{\partial y} &= \frac{\partial \beta_+}{\partial y} \frac{A_+}{A_+^2 + y^2} - \frac{\partial \beta_-}{\partial y} \frac{A_-}{A_-^2 + y^2} + \frac{2A_- \beta_- y}{(A_-^2 + y^2)^2} - \frac{2A_+ \beta_+ y}{(A_+^2 + y^2)^2} - \frac{4zy}{(y^2 + z^2)^2}, \\
\frac{\partial B_z}{\partial y} &= \frac{4y^2}{(y^2 + z^2)^2} - \frac{2}{y^2 + z^2}, \\
\frac{\partial B_x}{\partial z} &= \frac{\partial \beta_-}{\partial z} \frac{y}{A_-^2 + y^2} - \frac{\partial \beta_+}{\partial z} \frac{y}{A_+^2 + y^2}, \\
\frac{\partial B_y}{\partial z} &= \frac{\partial \beta_+}{\partial z} \frac{A_+}{A_+^2 + y^2} - \frac{\partial \beta_-}{\partial z} \frac{A_-}{A_-^2 + y^2} + \frac{2}{y^2 + z^2} - \frac{4z^2}{(y^2 + z^2)^2}, \\
\frac{\partial B_z}{\partial z} &= \frac{4yz}{(y^2 + z^2)^2}, \\
\frac{\partial \beta_{\pm}}{\partial x} &= \frac{A_{\pm} z}{1 - z/\sqrt{A_{\pm}^2 + y^2}} \frac{1}{(A_{\pm}^2 + y^2)^{3/2}} \cos \left[\arctan \left(\frac{-z}{\sqrt{A_{\pm}^2 + y^2}} \right) \right], \\
\frac{\partial \beta_{\pm}}{\partial y} &= \frac{yz}{1 - z/\sqrt{A_{\pm}^2 + y^2}} \frac{1}{(A_{\pm}^2 + y^2)^{3/2}} \cos \left[\arctan \left(\frac{-z}{\sqrt{A_{\pm}^2 + y^2}} \right) \right], \\
\frac{\partial \beta_{\pm}}{\partial z} &= \frac{1}{z - \sqrt{A_{\pm}^2 + y^2}} \cos \left[\arctan \left(\frac{-z}{\sqrt{A_{\pm}^2 + y^2}} \right) \right]. \tag{A.14}
\end{aligned}$$

A.3 U-trap minimum and gradients

To find the minimum of the U-trap, we must find x_0 , y_0 , and z_0 so that B_x , B_y , and B_z separately vanish. From the symmetry of the U-wire layout in Figure 5.4 we know that

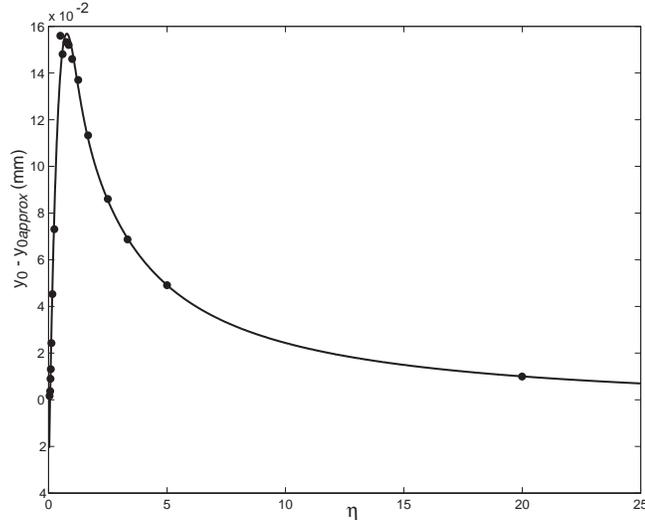


Figure A.1: Difference between y_0 and the simplified expression, $y_0 \text{ approx} = 2\pi\eta/\mu_0$, as a function of $\eta = I/B_{\text{bias}}$.

$x_0 = 0$. The expressions for $B_y(x = 0)$ and $B_z(x = 0)$ give us the following equations for y_0 and z_0 :

$$\begin{aligned}
 0 &= \frac{L\beta_0}{L^2/4 + y_0^2} + \frac{2z_0}{y_0^2 + z_0^2}, \\
 y_0 &= \frac{\mu_0 I}{4\pi B_{\text{bias}}} + \sqrt{\left(\frac{\mu_0 I}{4\pi B_{\text{bias}}}\right)^2 + z_0^2}, \\
 \beta_0 &= 1 + \frac{z_0}{\sqrt{L^2/4 + y_0^2 + z_0^2}},
 \end{aligned} \tag{A.15}$$

From the first equation we see that $z_0 \leq 0$, and we have verified this experimentally for both sets of wire current and bias field orientations.

The remainder of this appendix contains plots relevant to understanding how the U-trap's field minimum and gradients evolve during the compression and lowering of the trap center. For all plots, $x = 0$, and unless otherwise noted, $L = 1.5$ mm.

The numerical solution of Equations A.15 for y_0 and z_0 is plotted in Figures A.1 and A.2. Note that the curve in Figure A.1 represents the difference between y_0 calculated with the side wires and the value for y_0 that ignores the field from the side wires, defined as $y_0 \text{ approx}$ (see Equations 5.22). The fits are to a 4th order rational that is only a function of $\eta = I/B_{\text{bias}}$.

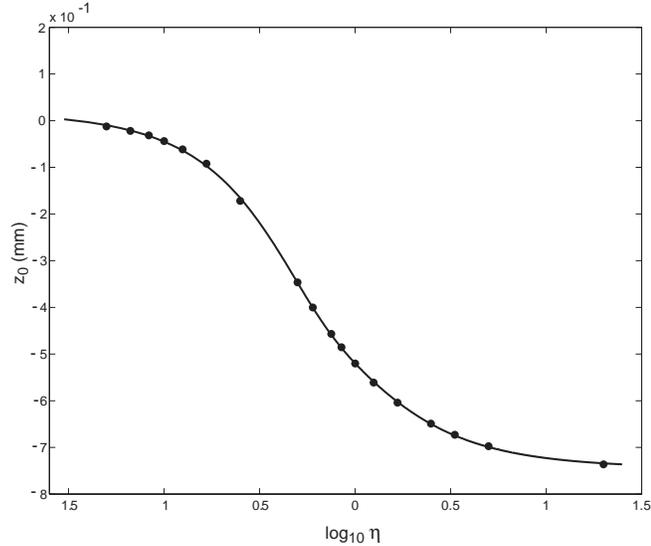


Figure A.2: Log plot of z_0 as a function of η . With no side wires present, z_0 would equal 0.

Using the y_0 and z_0 fits, we calculate the U-trap field gradients at the trap minimum and the results are shown in Figures A.3, A.4, and A.5. The plot of $\nabla_z \mathbf{B}$ is strange in that the gradient passes through zero to become negative for some values of $y_0(\eta)$. This can be understood by noting that the magnetic field changes from being positive along $+\hat{z}$ to being negative in along the $+\hat{z}$. This switch occurs because as z_0 becomes more negative, the quadrupole field pattern rotates with respect to the z-axis. Consequently, the z-axis passes through collinearity with an axis of the quadrupole field, and the gradient changes sign. The gradient sign change shifts back to positive once y_0 becomes small and the field rotates back again. This process does not cause the atoms to be lost in the \hat{z} direction because we assume the spins follow the field adiabatically. However, the trapping laser orientations will not be perfectly aligned with the quadrupole field, and this will transiently degrade the performance of the U-MOT.

Figure A.6 shows the difference between the full calculation for y_0 and the expression, $y_0 \text{ approx}$, that does not account for the field of the side wires. The simpler expression, $y_0 \text{ approx}$, overestimates the gradient by as much as 25 to 50 G/cm.

Not surprisingly, decreasing L increases the \hat{x} gradient. Figure A.7 plots $\nabla_x \mathbf{B}$ for $L = 0.5, 1, \text{ and } 1.5$ mm. Finally, Figure A.8 plots the ratio of the gradient in \hat{y} to that in \hat{x} . At 0.5 mm above the substrate, the trap anisotropy increases dramatically.

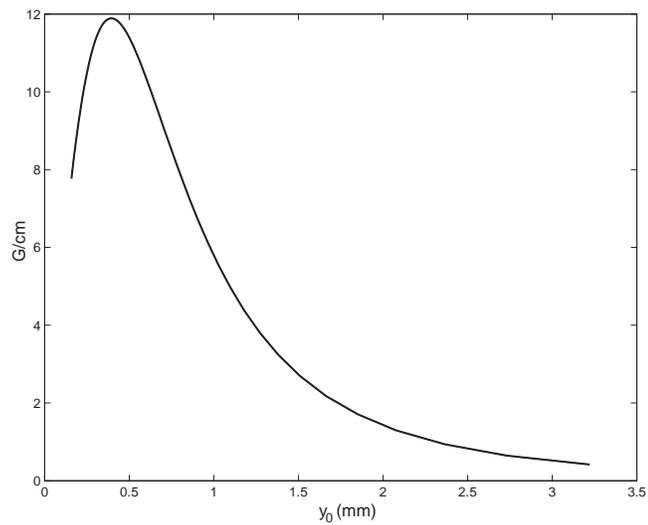


Figure A.3: The U-trap gradient in \hat{x} at the trap minimum for B_{bias} varying between 0.1 to 12 G. $I = 1$ A.

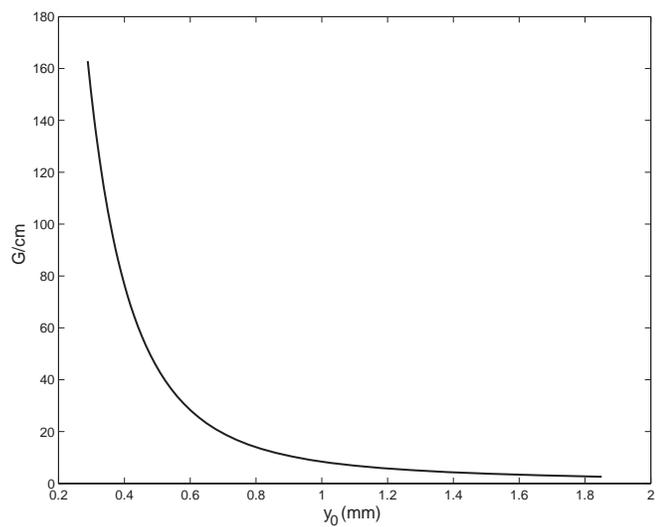


Figure A.4: The U-trap gradient in \hat{y} at the trap minimum for B_{bias} varying between 0.1 to 6 G. $I = 1$ A.

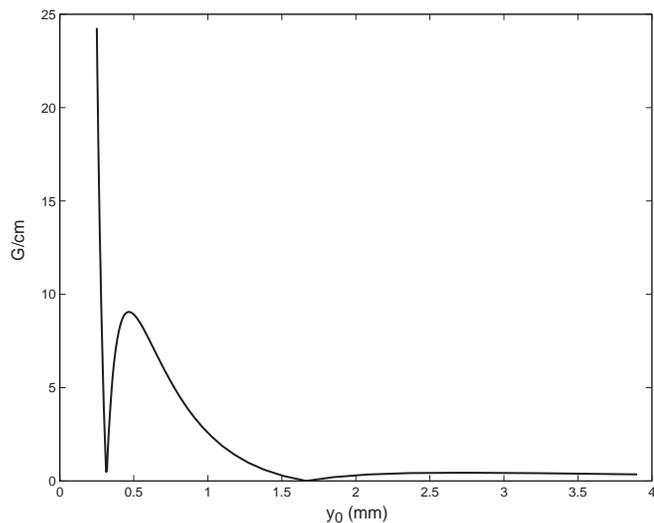


Figure A.5: The absolute value of the U-trap gradient in \hat{z} at the trap minimum for B_{bias} varying between 0.1 to 7 G. $I = 1$ A.

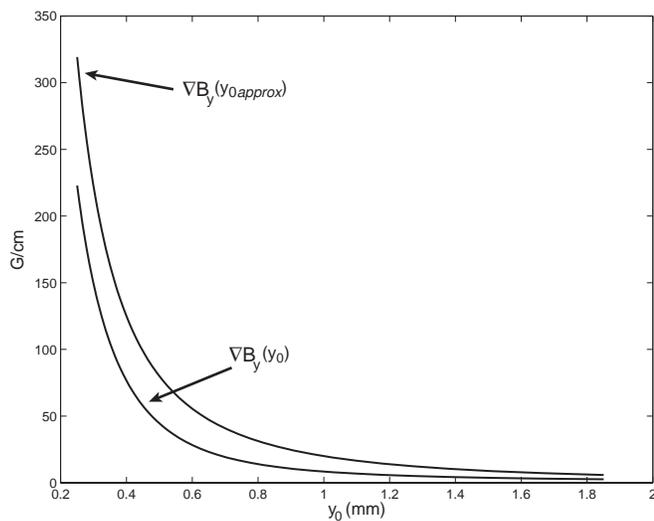


Figure A.6: The trap gradient in \hat{y} using the full calculation for y_0 versus the simpler $y_{0approx} = 2\pi\eta$ expression. Note that the simpler expression overestimates the gradient. B_{bias} varies between 0.1 to 7 G. $I = 1$ A.

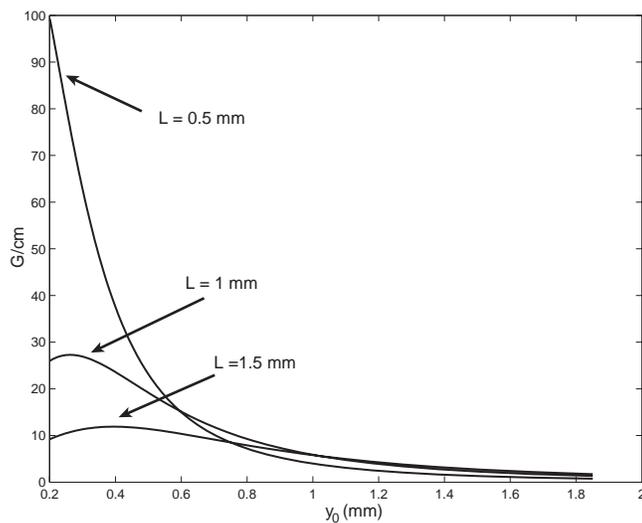


Figure A.7: The U-trap gradient in \hat{x} plotted for various values of L . B_{bias} varies between 0.1 to 9 G. $I = 1$ A.

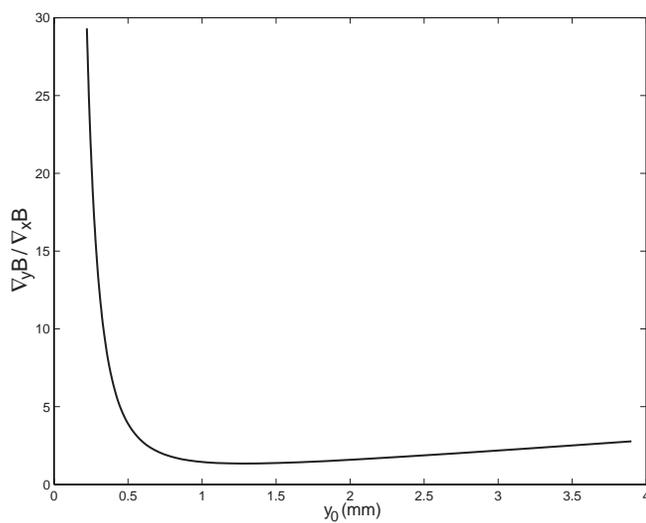


Figure A.8: The anisotropy of the U-trap in \hat{y} versus \hat{x} as a function of trap height. $\eta = I/B_{bias} = 0.125$ to 10 A/G.

Appendix B

Matlab Code

This Appendix includes the template Matlab codes used for simulations in this thesis and for computer control of the atom chip waveguiding experiment.

B.1 Cavity QED: transmission, signal-to-noise, and spontaneous emission

This Matlab code, `RunMastEqn.m`, calls the `MastEqn.m` function—presented in the next section—to calculate the cavity transmission, the signal, and the signal-to-noise, \mathcal{S}/\mathcal{N} , for arbitrary cavity QED parameters $[g_0, \kappa, \gamma_{\perp}, \Delta, \theta]$. The code presented here calculates these for varying drive power but can easily be modified to calculate them as a function of g , Δ , and θ as desired. For comparisons, the code solves both the optical bistability equation 6.5 and the two-level atom master equation 6.6. See Kevin Birnbaum’s thesis [83] for code that extends beyond the two-level approximation. The code presented here creates a plot of transmission versus drive power, signal versus drive power, and \mathcal{S}/\mathcal{N} versus drive power. The Fock-state basis N can be assigned for each drive power, which minimizes computational overhead. For instance, using $N = 100$ for a drive intracavity photon number of $n = 0.001$ is overkill, but is absolutely necessary for $n \geq 40$. We ran this code on a double 3 GHz Xeon processor under Linux, and found that we could reach a Fock state basis of 107 before crashing the calculation. For most parameter regimes, an intracavity drive of $n = 50$ to 70 could be reliably simulated.

```
clear
figure_series1 = 1;
figure_series2=figure_series1+1;
figure_scale = 1;
```

```

%-----constants
hbar=(6.626e-34)/(2*pi); %J s
c=3e8; % m/s
e0=8.85e-12; %coul^2/N-m^2 permitivity of free space
mcs=2.207e-25; %kg mass of Cs
mud=2.686e-29; %C m dipole moment D2 for Cs cycling transition
Lambda=852.4e-9; %m cesium D2
gmapara=2*2.6e6*2*pi; % for cesium D2
% mrb87=1.443e-25; %kg mass of Rb87
% mud=2.534e-29; %C m dipole moment D2 for Rb87 cycling transition
% Lambda=780.2e-9; %m Rb87 D2
% gmapara=2*3e6*2*pi; % for Rb87
wl=2*pi*c/Lambda;
%-----cavity QED parameters
% for Fabry-Perot let cavity_parm=0, for others let cavity_parm=1
cavity_parm=0;
if cavity_parm==0;
F = 1050;
R = 1e-3;
L = 27e-6;
w0=sqrt((L*Lambda/6.3)*sqrt((2*R-L)/L));
V=pi*w0^2*L/4;
kappa=pi*c/(2*L*F); % divide by 2pi for Hz
g0=mud*sqrt(hbar*wl/(2*hbar^2*e0*V)); % divide by 2pi for Hz
else cavity_parm==1;
g0=16e9*2*pi;
kappa=4.4e9*2*pi;
end
% The cavity detunings
Delta=0e9*2*pi; % divide by 2pi for Hz
theta=0e9*2*pi; % divide by 2pi for Hz
% Detection parameters
eta = 0.44; % detection efficiency
Int_time = 10e-6; % seconds, 1/bandwidth of detector
% number of atoms
N_atom=1;
% drive power quoted in intracavity photon number in empty, resonant
% cavity
Drive=[1e-5 1e-4 4e-4 1e-3 4e-3 1e-2 2e-2 4e-2 6e-2 0.1 ...
0.2 0.4 0.6 1 2 6 10 15 20 30];
% Fock basis
N=[5 5 5 5 5 5 5 5 5 10 10 10 10 15 20 30 40 60 80];
if length(Drive)~=length(N)
error('Drive and Fock basis mismatched')
end
%-----master eqn soln
% The cavity QED parameters are normalized to gmapara(llel) for

```

```

% numerical stability
g=g0/gammapara;
k=kappa/gammapara;
energydecay=2*kappa;
gammapara_norm=1;
Delta_norm=Delta/gammapara;
theta_norm=theta/gammapara;
DriveField = k.*sqrt(Drive);
input_vec1=[g k gammapara_norm Delta_norm theta_norm];
input_vec0=[0 k gammapara_norm Delta_norm theta_norm];
for i=1:length(Drive);
    tic
    [ss_photon_numbers(i),var(i),Rho_ee(i)]= ...
        MastEqn(DriveField(i),N(i),input_vec1);
    [ss_photon_numbers0(i),var0(i),Rho_ee0(i)] = ...
        MastEqn(DriveField(i),N(i),input_vec0);
    toc
end
% transmission vs. drive plot
Counts=real(ss_photon_numbers);
EmptyCounts=real(ss_photon_numbers0);
Trans=Counts./Drive;
figure(figure_series1)
loglog(Drive,Trans,'k','MarkerSize',16,'MarkerFaceColor','k');
hold
%-----optical bistability eqn soln
gammaperp=0.5*gammapara;
n_0=gammaperp^2/(2*g0^2);
maxx=10;
nn=1000;
DriveOBE=sqrt(nn)/(40*sqrt(n_0)); %nn is scaled drive
C=g0.^2.*N_atom/(2.*kappa.*gammaperp);
Delta_norm2=Delta./gammaperp;
theta_norm2=theta./kappa;
x=(0:.00001:maxx).*DriveOBE;
a=(1+2.*C./(1+x.^2+Delta_norm2.^2));
b=(theta_norm2-2.*C.*Delta_norm2./(1+x.^2+Delta_norm2.^2));
y=sqrt(x.^2.*(a.^2+b.^2));
X=(x.^2).*n_0;
Y=(y.^2).*n_0;
T=(x./y).^2;
loglog(Y,T,'k','LineWidth',2);
% without atom
CC=0;
xnoatom=(.01:.1:maxx).*DriveOBE;
a=(1+2.*CC./(1+xnoatom.^2+Delta_norm2.^2));
b=(theta_norm2-2.*CC.*Delta_norm2./(1+xnoatom.^2+Delta_norm2.^2));

```

```

ynoatom=sqrt(xnoatom.^2.*(a.^2+b.^2));
X1=(xnoatom.^2).*n_0;
Y1=(ynoatom.^2).*n_0;
Tnoatom=(xnoatom./ynoatom).^2;
loglog(Y1,Tnoatom,'--r','LineWidth',2);
%-----calculation of spontaneous emission per integration time
signal=eta.*abs(EmptyCounts-Counts).*energydecay*Int_time;
err=sqrt(eta.*(EmptyCounts+Counts).*energydecay*Int_time);
% DriveNew=Drive.*energydecay*Int_time;
DriveNew=Drive.*energydecay.*hbar.*wl.*1e9; % in nW
SN=signal./err;
[signalmax signalmaxindex] = max(signal);
[SNmax SNmaxindex] = max(SN);
[drivemax maxdriveindex] = max(DriveNew);
%-----calculation of spontaneous emission per integration time
SponE = real((Rho_ee.*gammapara).*Int_time);
Sig=- (real(ss_photon_numbers)-real(ss_photon_numbers0))...
    *Int_time*energydecay;
RatioE=SponE./Sig;
Ratio_Optimum=RatioE(SNmaxindex)
%-----transmission vs. drive plot
title(['L = ',num2str(L*1e6),...
    '\mu m, F = ',num2str(F),', g_0 = ',num2str(round(g0*1e-6/6.28)),...
    ' MHz, \kappa = ',num2str(round(100*kappa*1e-6/6.28)/100),...
    ' MHz']; ['\Delta = ',num2str(Delta*1e-6/(2*pi)), ' MHz; \Theta = '...
    ,num2str(theta*1e-6/(2*pi)), ' MHz; SponE = '...
    ,num2str(round(Ratio_Optimum*1000)/10),'% of signal']...
    , 'FontSize',12,'FontWeight','Bold')
ylabel('Transmission','FontWeight','Bold')
xlabel('Intracavity Photon Number','FontWeight','Bold')
axis([0.5*Drive(1) nn.*10 1e-4 1.0])
hold
%-----signal and S/N figures
figure(figure_series2)
subplot(2,1,1)
errorbar(DriveNew,signal,err,'.')
Int_time2=round(Int_time*1e6*100)/100;
MaxPhoton=DriveNew(SNmaxindex);
ylabel(['Signal per ', num2str(Int_time2), ' \mu s integration']...
    , 'FontWeight','Bold')
title(['N = ',num2str(N_atom),', L = ',num2str(L*1e6),'; \eta = '...
    ,num2str(eta),'; {\mu}m, F = ',num2str(F),...
    ', g_0 = ',num2str(round(g0*1e-6/6.28)), ' MHz, \kappa = '...
    ,num2str(round(100*kappa*1e-6/6.28)/100), ' MHz; \Delta = '...
    ,num2str(Delta*1e-6/(2*pi)), ' MHz; \Theta = '...
    ,num2str(theta*1e-6/(2*pi)), ' MHz'];...
    ['Int time = ',num2str(Int_time2),'; SponE/Int time = '...

```

```

,num2str(round(SponE(SNmaxindex)*10)/10),' Photons; Max SN =',...
num2str(round(10*SNmax/SN(SNmaxindex))/10)...
,'; Optimal Drive/BW = ',num2str(round(MaxPhoton)), ' Photons']},...
'FontSize',10,'FontWeight','Bold')
axis([-1e-3*drivemax (drivemax*figure_scale+100) ...
0 (signal(SNmaxindex)+signal(SNmaxindex)*0.5)])
subplot(2,1,2)
plot(DriveNew,SN,'.k')
axis([-1e-3*drivemax (drivemax*figure_scale+100)...
0 (SNmax/SN(SNmaxindex)+10)])
ylabel('S/N','FontWeight','Bold')
xlabel('Drive (nW)','FontWeight','Bold')

```

B.2 Master equation for a two-level atom

Using the Quantum Optics Toolbox [160, 161], this Matlab function calculates the steady-state density operator ρ_{ss} by solving the cavity QED master equation 6.6 presented in Section 6.4. This function is to be used with the program `RunMasterEqn.m` listed in the previous section (or a variation of this program).

```

function [ss_photon_numbers,var,Rho_ee]=MastEqn(DriveField,N,input_vec)
wl = 0;      % define laser frequency as zero detuning
E = DriveField;
g = input_vec(1);
kappa = input_vec(2);
gamma_par = input_vec(3);
wa = input_vec(4);
wc = input_vec(5);
ida = identity(N); idatom = identity(2);
% Define cavity field and atomic operators
a = tensor(destroy(N),idatom);
sm = tensor(ida,sigmam);
% Define projection operators into excited/ground state
proj_e = tensor(ida,1/2*(idatom+sigmaz));
proj_g = tensor(ida,1/2*(idatom-sigmaz));
% Hamiltonian
H = (wa-wl)*sm'*sm + (wc-wl)*a'*a + i*g*(a'*sm - sm'*a) + i*E*(a'-a);
% Collapse operators
C1 = sqrt(2*kappa)*a;
C2 = sqrt(gamma_par)*sm;
C1dC1 = C1'*C1;
C2dC2 = C2'*C2;
% Calculate the Liouvillian
LH = -i * (spre(H) - spost(H));
L1 = spre(C1)*spost(C1')-0.5*spre(C1dC1)-0.5*spost(C1dC1);

```

```

L2 = spre(C2)*spost(C2')-0.5*spre(C2dC2)-0.5*spost(C2dC2);
L = LH+L1+L2;
% Find steady state density matrix
rho_ss = steady(L);
% Calculated the steady state intracavity photon_number and it's
% variance. Rho_ee is the excited state population which allows
% one to calculate the number of spontaneously emitted photons
ss_photon_numbers = expect(a'*a,rho_ss);
n2 =expect(a'*a*a'*a,rho_ss);
var =sqrt(n2-(ss_photon_numbers).^2);
Rho_ee =expect(sm'*sm,rho_ss);

```

This function may also be used for simulating the presence of more than one atom by defining an atomic operator, \mathbf{sm}_i , and collapse operator, $\mathbf{C2}_i$, for each atom i . Up to 10 atoms were introduced without noticing any computational slowdown. The Hamiltonian and the Liouvillian then become (for the case of $i = 2$):

```

H = (wa-wl)*sm_1'*sm_1+(wa-wl)*sm_2'*sm_2 +...
    i*g*(a'*sm_1-sm_1'*a)+i*g*(a'*sm_2-sm_2'*a)+(wc-wl)*a'*a+i*E*(a'-a);
L2_1 = spre(C2_1)*spost(C2_1')-0.5*spre(C2dC2_1)-0.5*spost(C2dC2_1);
L2_2 = spre(C2_2)*spost(C2_2')-0.5*spre(C2dC2_2)-0.5*spost(C2dC2_2);
L = LH+L1+L2_1+L2_2;

```

B.3 Magnetic field, gradient, and curvature for U- and Z-traps and waveguides

This Matlab script, `Waveguide.m`, uses the function `Field.m`—provided below—to calculate and plot any slice through the 3D magnetic field above an arbitrary arrangement of rectilinear wire segments, both finite and semi-infinite. Bias fields in \hat{x} , \hat{y} , and \hat{z} are included as well as a quadrupole field from coils whose axis is along the waveguide. The code presented here is for the case of the Munich atom chip as described in Chapter 7 and pictured in Figure 7.5. It is easily adjusted for any other rectilinear wire pattern, specifically that used in the microcavity experiment at Caltech (see Chapters 2, 3, and 6). The `GLeft` and `GRight` designate the beginning and end of each segment of the wireguide designated as the wire G1 to the intersection of GU8 in Figure 7.5. The `GU` and `GL` are the semi-infinite wires that form the sides of the the U- and Z-traps, the waveguide gates, and the dimple wires. They are also listed in Figure 7.5. The wireguide is positioned along \hat{x} while the side wires are along \hat{z} . \hat{y} is perpendicular to the atom chip. `Field.m` calculates the field

analytically using expressions listed in Appendix A¹. The fields from finite wire widths can be calculated analytically, but the expression is too unwieldy to list here or include in this code. In cases where this is needed, full numerical solutions of the Laplace equation are usually required as well: Experiments in which atoms are closer to the wire than the wire width are usually performed for purposes requiring more precision than waveguiding (for example, BEC splitting in a double well as discussed in Chapter 10). In these cases, a full numerical simulation should be performed using commercial packages such as Femlab.

```
clear
%plot freq (1) or G/cm^2 (0)?
%plot gradient of x (1) or z (0)?
freq=1; xplot=0;
% all distance in units of microns
resolution=50; % resolution of the plots
plotleft=-2000; % boundary of plots
plotright=9000;
% shifting and quadrupole coil radii
R1=50e3; R2=50e3; RQ=50e3;
% Number of turns in quadrupole or shifting coils
N1=87; N2=N1; N=87; shim=0.85;
% Coil currents
I_Quad=0*N*shim;
I_Coil1=-0*N1*shim;
I_Coil2=0*N2*shim;
% Coil positions (half the distance between coils)
Q_Coil=40e3;
OffsetY=0e3;
D_Coil1=-40e3;
D_Coil2=40e3;
% Bias fields (in Gauss)
B_xbias=-5*1; % (- for positive wire currents)
B_ybias=-1.8*0;
B_zbias=16.3;
% wire layout (1) for wire presence and (0) for wire absence
Guide_Wires= [ 1 1 1 1 1 1 1 0 0];
Guide_Current = -2; %needs to be negative if sides are positive
GateU_Wires = [ 1 0 0 0 0 1 1 0];
GateU_Current= [2.5 1 1 1 1 -1 2 1];
GateD_Wires = [ 1 0 0 0 0 1 0 0];
GateD_Current= [2.5 1 1 1 1 -1 2 1];
Guide_Off=0;
max_guide=length(Guide_Wires);
```

¹Note: these expressions are valid only as long as the atoms are no closer the chip surface than the microwire's width (see chapter 2)

```

for i=1:max_guide
    if Guide_Wires(i) == 1;
        I_Guide(i)=Guide_Current; %amps
    else
        I_Guide(i)=Guide_Off; %amps
    end
end
% wire guide segment layouts:this is for the Munich guide, adjust
% accordingly
GLeft(1)=-5e3;GRight(1)=-1e3;
GLeft(2)=GRight(1);GRight(2)=1e3;
GLeft(3)=GRight(2);GRight(3)=3e3;
GLeft(4)=GRight(3);GRight(4)=5e3;
GLeft(5)=GRight(4);GRight(5)=7e3;
GLeft(6)=GRight(5);GRight(6)=8e3;
GLeft(7)=GRight(6);GRight(7)=9e3;
GLeft(8)=GRight(7);GRight(8)=10e3;
GLeft(9)=GRight(8);GRight(9)=11e3;
% gate wires, upper segments
GateU_Off=0;
max_gate_u=length(GateU_Wires);
for i=1:max_gate_u
    if GateU_Wires(i) == 1;
        I_GateU(i)=GateU_Current(i); %amps
    else
        I_GateU(i)=GateU_Off; %amps
    end
end
% positions are along the wire guide
GU(1)=-1e3;
GU(2)=1e3;
GU(3)=3e3;
GU(4)=5e3;
GU(5)=7e3;
GU(6)=8e3;
GU(7)=9e3;
GU(8)=10e3;
% gate wires, lower segments
GateD_Off=0;
max_gate_d=length(GateD_Wires);
for i=1:max_gate_d
    if GateD_Wires(i) == 1;
        I_GateD(i)=GateD_Current(i); %amps
    else
        I_GateD(i)=GateD_Off; %amps
    end
end
end

```

```

% positions are along the wire guide
GL(1)=-1e3;
GL(2)=1e3;
GL(3)=3e3;
GL(4)=5e3;
GL(5)=7e3;
GL(6)=8e3;
GL(7)=9e3;
GL(8)=11e3;
% unit conversions
GLeft=GLeft.*1e-6;
GRight=GRight.*1e-6;
GU=GU.*1e-6;
GL=GL.*1e-6;
Q_Coil=Q_Coil.*1e-6;
OffsetY=OffsetY.*1e-6;
D_Coil1=D_Coil1.*1e-6;
D_Coil2=D_Coil2.*1e-6;
RQ=RQ.*1e-6;
R1=R1.*1e-6;
R2=R2.*1e-6;
% set the trap min for calculating the slice in 3D
y=abs(2000*I_Guide(2)/B_zbias)
% permute coordinates to create slices in other planes.
[x,z]=meshgrid(plotleft:resolution:plotright);
x=x.*1e-6;y=y.*1e-6;z=z.*1e-6;
B_xbias=B_xbias*1e-4;
B_ybias=B_ybias*1e-4;
B_zbias=B_zbias*1e-4;
[B_tot,QuadField_x]=Field(I_Guide,GLeft,GRight,x,y,z,R1,R2,D_Coil1,...
    D_Coil2,I_Coil1,I_Coil2,B_xbias,B_ybias,B_zbias,I_Quad,Q_Coil,...
    RQ,OffsetY,max_guide,GU,max_gate_u,I_GateU,GL,max_gate_d,I_GateD);
B_tot=B_tot.*1e4;
x=x.*1e6;y=y.*1e6;z=z.*1e6;
figure(1)
meshc(x,z,B_tot)
xlabel('X axis (\mum)')
ylabel('Z axis (\mum)')
zlabel('B field (G)')
% axis([-5000 11000 -2000 2000 0 max(max(B_tot))+10])
% If there is a quadrupole field, this prints its gradient
gradQuad=100*(diff(QuadField_x(round(length(QuadField_x)/2),:))./...
    100e-6);
Quad_Gradient=abs(min(gradQuad))
Bmin=min(min(B_tot))
% plots the gradient and curvature of the wire field
[GradBx,GradBz]=gradient(B_tot,resolution*1e-4);

```

```

[GGradBx,GGradBzx]=gradient(GradBx,resolution*1e-4);
[GGradBxz,GGradBz]=gradient(GradBz,resolution*1e-4);
if xplot==1
    figure(11)
    meshc(x,z,abs(GradBx));
    xlabel('X axis (\mum)')
    ylabel('Z axis (\mum)')
    zlabel('Gradient of B field, X-axis (G/cm)')
    % axis([-5000 11000 -2000 2000 0 max(max(GradBy))])
    if freq==1
        figure(12);
        meshc(x,z,12.7.*sqrt(GGradBx));
        xlabel('X axis (\mum)')
        ylabel('Z axis (\mum)')
        zlabel('Frequency of Trap in X-axis (Hz)')
        % axis([-5000 11000 -2000 2000 0 max(max(GGradBy))])
    else
        figure(12);
        meshc(x,z,GGradBx);
        xlabel('X axis (\mum)')
        ylabel('Z axis (\mum)')
        zlabel('Curvature of B Field, X-axis (G/cm^2)')
        % axis([-5000 11000 -2000 2000 0 max(max(GGradBy))])
    end
else
    figure(11)
    meshc(x,z,abs(GradBz));
    xlabel('X axis (\mum)')
    ylabel('Z axis (\mum)')
    zlabel('Gradient of B field, Z-axis (G/cm)')
    % axis([-5000 11000 -2000 2000 0 max(max(GradBy))])
    if freq==1
        figure(12);
        meshc(x,z,12.7.*sqrt(GGradBz));
        xlabel('X axis (\mum)')
        ylabel('Z axis (\mum)')
        zlabel('Frequency of Trap in Z-axis (Hz)')
        % axis([-5000 11000 -2000 2000 0 max(max(GGradBy))])
    else
        figure(12);
        meshc(x,z,GGradBz);
        xlabel('X axis (\mum)')
        ylabel('Z axis (\mum)')
        zlabel('Curvature of B Field, Z-axis (G/cm^2)')
        % axis([-5000 11000 -2000 2000 0 max(max(GGradBy))])
    end
end
end

```

```

function [B_tot,QuadField_x]=Field(I_Guide,LL,LR,x,y,z,R1,R2,...
    D_Coil1,D_Coil2,I_Coil1,I_Coil2,B_xbias,B_ybias,B_zbias,I_Quad,...
    Q_Coil,RQ,OffsetY,max_guide,LU,max_gate_u,I_GateU,LD,...
    max_gate_d,I_GateD)
%-----Defs
mu=(4*pi)*10^-7; %N/A^2
Beta=y.^2+z.^2;
for i=1:max_guide
    const_G(i)=mu.*I_Guide(i)./(4*pi);
    B_G(:,:,i)=-const_G(i).*((LL(i)-x)./(Beta.*sqrt(Beta+...
        (x-LL(i)).^2))-(LR(i)-x)./(Beta.*sqrt(Beta+(x-LR(i)).^2)));
    B_Gz(:,:,i)=B_G(:,:,i).*y;
    B_Gy(:,:,i)=B_G(:,:,i).*z;
end
for i=1:max_gate_u
    AU(:,:,i)=x-LU(i);
    Alpha_U(:,:,i)=sqrt(AU(:,:,i).^2+y.^2+z.^2);
    ZU(:,:,i)=1+z./Alpha_U(:,:,i);
    const_GateU(i)=mu.*I_GateU(i)./(4*pi);
    B_GateU(:,:,i)=const_GateU(i).*(ZU(:,:,i)./(AU(:,:,i).^2+y.^2));
    B_GateUx(:,:,i)=B_GateU(:,:,i).*(-y);
    B_GateUy(:,:,i)=B_GateU(:,:,i).*AU(:,:,i);
end
for i=1:max_gate_d
    AD(:,:,i)=x-LD(i);
    Alpha_D(:,:,i)=sqrt(AD(:,:,i).^2+y.^2+z.^2);
    const_GateD(i)=mu.*I_GateD(i)./(4*pi);
    B_GateD(:,:,i)=const_GateD(i).*(1./(Alpha_D(:,:,i).^2+...
        z.*Alpha_D(:,:,i)));
    B_GateDx(:,:,i)=B_GateD(:,:,i).*(-y);
    B_GateDy(:,:,i)=B_GateD(:,:,i).*AD(:,:,i);
end
B_Gz_tot=sum(B_Gz,3);
B_Gy_tot=sum(B_Gy,3);
B_GateUx_tot=sum(B_GateUx,3);
B_GateUy_tot=sum(B_GateUy,3);
B_GateDx_tot=sum(B_GateDx,3);
B_GateDy_tot=sum(B_GateDy,3);
Coil1x=(mu/2)*R1.^2.*I_Coil1./(((x-D_Coil1).^2+R1^2).^(3/2));
Coil2x=(mu/2)*R2.^2.*I_Coil2./(((x-D_Coil2).^2+R2^2).^(3/2));
const_QuadField=3*mu.*I_Quad.*Q_Coil.*RQ^2./(RQ^2+Q_Coil^2).^(5/2);
QuadField_x=const_QuadField.*x;
QuadField_y=-0.5.*const_QuadField.*(y-OffsetY);
QuadField_z=-0.5.*const_QuadField.*z;
B_tot=sqrt((B_GateUx_tot+B_GateDx_tot+B_xbias+Coil1x+Coil2x...
    +QuadField_x).^2+(B_Gy_tot+B_GateUy_tot+B_GateDy_tot...
    +B_ybias+QuadField_y).^2+(B_Gz_tot+B_zbias+QuadField_z).^2);

```

B.4 Computer control code for the atom chip experiment

This is the Matlab code used for producing the .dat files that control each of the analog and digital output channels. Chapter 3, Section 3.4 contains more information regarding the program's use. This program calls a "Calibration" function that contains a list of the computer output voltage-to-power supply current conversions. $I = V/\alpha$, where α is the conversion factor. The conversions are: $\alpha = 0.5$ for the Kepco 20-20; $\alpha = 0.25$ for the Kepco 25-4, current controlled; and $\alpha = 0.09$ for the Power10. For the voltage controlled Kepco 25-4, the computer voltage-to-supply voltage conversion is 0.4.

```

%%%%%%%% Update rate should be 10000 updates/sec.
%%%%%%%% "Bias" is the Utrap bias field made by the intra-chamber
%%%%%%%% coils. "Top" coil is in the z-axis
%%%%%%%%%%%%%% Supply Key:
% % Supply Key = 1      Kepco 25-4 Current Control: label: #1, 2, 4, 5
% % Supply Key = 2      Kepco 20-20 Current Control: Digital,Analog
% % Supply Key = 3      Kepco 36-30
% % Supply Key = 4      Power10 Lower Supply
% % Supply Key = 5 and higher Kepco 25-4 Voltage Control: label: #3
clear
warning('off','Matlab:dividebyzero');
%-----Initalization-----
Bias_Supply=2;MacroU_Supply=4;Y_Supply=2;Top_Supply=5;
DTrap_Supply=1;GuideWire_Supply=1;HLeft_Supply=1;HRight_Supply=1;
[I_Bias, I_MacroU, I_Y, I_Top, I_DTrap, I_GuideWire, I_HLeft,...
 I_HRight]=Calibration(Bias_Supply,MacroU_Supply,Y_Supply,...
 Top_Supply,DTrap_Supply,GuideWire_Supply,HLeft_Supply,...
 HRight_Supply);
DointsperTime=10; %points/ms
ms=1*DointsperTime;
on=0;          %Laser TTL
off=1;        %Laser TTL 5 for analog, 1 for digital
short=1;      %Relay TTL 5 for analog, 1 for digital
disconnect=0; %Relay TTL
onAbs=0;     %AOM 1 TTL
offAbs=1;    %AOM 1 TTL 5 for analog, 1 for digital
%-----Detection-----
det=1;        %0 for Fluor, 1 for Abs
%-----Times-----
ShiftTime     = 1*20*ms;
HoldTime      = 0*0*10*ms; % rarely used
DMOTime      = 1*20*ms;
DTransTime   = 0*0*5*ms;
CoolTime     = 1*2.7*ms;
DecayTime    = 1*0.3*ms;

```

```

ODTime          = 1*0.6*ms;
DTrapTime       = 1*0*5*ms;
MagShiftTime    = 1*8*ms;
RotateTime      = 1*20.4*ms;
RevRotateTime   = 0*0*40.0*ms; % used for diagnostics
HTransTime      = 1*15*ms;
HTrapTime       = 1*2*ms;
FreeExpand      = 1*1.0*ms;    % min of 0.3 ms for abs imaging
ImageTime       = 1*40*ms;
CaptureTime     = 1*ms;
TimeToImage     = ShiftTime+HoldTime+DMOTime+DTransTime+CoolTime+...
                DecayTime+ODTime+DTrapTime+MagShiftTime+RotateTime+...
                RevRotateTime+HTransTime+HTrapTime+FreeExpand;
TimeToImage = TimeToImage/ms
TimeTotal = 10*(TimeToImage+(ImageTime+CaptureTime)/ms);
if mod(TimeTotal,2) == 1
    CaptureTime = 1*ms;
else
    CaptureTime = 1.1*ms;
end
%-----Current and Field Final Setpoints-----
BiasFin         = 0.9;    %A
MacroUFin       = 22;    %A
YOffsetFin      = 0.5;    %A
AbsDetectFin    = offAbs;
TopOffsetFin    = 0.1;    %V
DTrapFin        = 0.0;    %A
GuideWireFin    = 0.0;    %A
HLeftFin        = 0.0;    %A
HRightFin       = 0.0;    %A
SubDopplerFin   = off;
TrapLaserFin    = on;
HLeftFinRelay   = disconnect;
HRightFinRelay  = disconnect;
%-----Image and Ramp Setpoints-----
trunc           = 1;
YOffsetAbsImage = 1*(YOffsetFin + 4);
%-----Special Current Setpoints-----
MacroURampFin   = 0;
DTrapHold       = 1.5;    %A rarely used
DTrapLoad       = 2.20;   %A
DTrapTrans      = .5;    %A rarely used
DMTrapStart     = 3.5;    %A
DMTrap          = 3.5;    %A
DMShift         = 3.5;    %A
DMRotate        = 3.5;    %A
BiasShift       = .5;    %A

```

```

BiasHold           = 0.5;    %A rarely used
BiasDLoad          = 1.2;    %A
BiasDTrans         = 0.0;    %A rarely used
BiasOD             = 2;      %A
BiasDMTrapStart    = 3.0;    %A
BiasDMTrap         = 3.0;    %A
BiasDMShift        = 3.0;    %A
BiasRotate         = 1.0;    %A
BiasHTrans         = 1.0;    %A
BiasH              = 1.0;    %A
TopOffsetShift     = 0.0;    %V
TopOffsetHold      = 0;      %V
TopOffsetDMOT      = 0.0;    %V
TopOffsetDTrans    = 0.0;    %V
TopOffsetDTrap     = 0.0;    %V
TopOffsetMShift    = 0.0;    %V
TopOffsetR         = 0.4;    %V
TopOffsetH         = 0.4;    %V
YOffsetShift       = YOffsetFin + 0.0; %A
YOffsetHold        = YOffsetFin + 0.0; %A
YOffsetDMOT        = YOffsetFin - 0.1; %A
YOffsetDTrans      = YOffsetFin + 0.0; %A
YOffsetDTrap       = YOffsetFin + 0.4; %A
YOffsetMShift      = YOffsetFin - 1.5; %A
YOffsetRStart      = YOffsetFin - 1.5; %A
YOffsetRFin        = YOffsetFin - 4;   %A
YOffsetHTrans      = YOffsetFin - 4.0; %A
YOffsetHTrap       = YOffsetFin - 4;   %A
YTransferMax       = 3.5;    %A may be used instead of H-trap
YTrapHold          = 3.5;    %A may be used instead of H-trap
YTrapRampFin       = 0;      %A may be used instead of H-trap
GuideWireTrans     = 3.5*1; %A
HLeftTrans         = 2.0*0; %A
HRightTrans        = 2.0*0; %A
GuideWireHold      = 3.5*1; %A
HLeftHold          = 2.0*0; %A
HRightHold         = 2.0*0; %A
GuideWireRampFin   = 0;      %A
HLeftRampFin       = 0;      %A
HRightRampFin      = 0;      %A
HLeftTransRelay    = short;
HRightTransRelay   = short;
HLeftHoldRelay     = short;
HRightHoldRelay    = short;
%-----Events-----
Shifting           = 1 + ShiftTime;
Holding            = Shifting + HoldTime;

```

```

HoldDMOT      = Holding + DMOTime;
DTransfer     = HoldDMOT + DTransTime;
Cooling       = DTransfer + CoolTime;
FieldsOff     = Cooling + DecayTime;
OpticalDumping = FieldsOff + OTime;
DTrapping     = OpticalDumping + DTrapTime;
MagShift      = DTrapping + MagShiftTime;
Rotating      = MagShift + RotateTime;
RevRotating   = Rotating + RevRotateTime;
HTransferring = RevRotating + HTransTime;
HTrapping     = HTransferring + HTrapTime;
Drop          = HTrapping + FreeExpand;
FluorImage    = Drop + ImageTime;
CaptureF      = FluorImage + CaptureTime;
AbsImage      = Drop + ImageTime;
CaptureA      = AbsImage + CaptureTime;
%-----Processes-----
%%%%%%%%%%%% Shift
time          = 1:Shifting;
if length(time) ~= 1
x=-2:4/Shifting:2-4/Shifting;
Bias(time)    = BiasShift;
MacroU(time)  = (MacroURampFin-MacroUFin)/2.*...
               (erf(x)-erf(-2))+MacroUFin;
YOffset(time) = YOffsetShift;
TopOffset(time) = TopOffsetShift;
DTrap(time)   = 0;
GuideWire(time) = 0;
HLeft(time)   = 0;
HRight(time)  = 0;
TrapLasers(time) = on;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
else
Bias(time)    = BiasFin;
MacroU(time)  = MacroUFin;
YOffset(time) = YOffsetFin;
TopOffset(time) = TopOffsetFin;
DTrap(time)   = DTrapFin;
GuideWire(time) = 0;
HLeft(time)   = 0;
HRight(time)  = 0;
TrapLasers(time) = TrapLaserFin;
AbsDetect(time) = offAbs;

```

```

Vacant2(time)           = 0;
Vacant3(time)           = 0;
HLeftRelay(time)        = disconnect;
HRightRelay(time)       = disconnect;
end
%%%%%%%%%%%% Holding
x=-2:4/HoldTime:2-4/HoldTime;
time                     = (Shifting+1):Holding;
Bias(time)               = BiasHold;
MacroU(time)             = (MacroURampFin-MacroUFin)/2.*...
    (erf(x)-erf(-2))+MacroUFin;
YOffset(time)           = YOffsetHold;
TopOffset(time)         = TopOffsetHold;
DTrap(time)              = DTrapHold;
GuideWire(time)         = 0;
HLeft(time)              = 0;
HRight(time)             = 0;
TrapLasers(time)        = on;
AbsDetect(time)         = offAbs;
Vacant2(time)           = 0;
Vacant3(time)           = 0;
HLeftRelay(time)        = disconnect;
HRightRelay(time)       = disconnect;
%%%%%%%%%%%% DMOT
time                     = (Holding+1):HoldDMOT;
Bias(time)               = BiasDLoad;
MacroU(time)             = 0;
YOffset(time)           = YOffsetDMOT;
TopOffset(time)         = TopOffsetDMOT;
DTrap(time)              = DTrapLoad;
GuideWire(time)         = 0;
HLeft(time)              = 0;
HRight(time)             = 0;
TrapLasers(time)        = on;
AbsDetect(time)         = offAbs;
Vacant2(time)           = 0;
Vacant3(time)           = 0;
HLeftRelay(time)        = disconnect;
HRightRelay(time)       = disconnect;
%%%%%%%%%%%% DTransfer
y=-2:4/DTransTime:2- 4/DTransTime;
time                     = (HoldDMOT+1):DTransfer;
Bias(time)               = BiasDTrans;
MacroU(time)             = 0;
YOffset(time)           = YOffsetDTrans;
TopOffset(time)         = TopOffsetDTrans;
DTrap(time)              = DTrapTrans;

```

```

GuideWire(time)           = 0;
HLeft(time)               = 0;
HRight(time)              = 0;
TrapLasers(time)         = on;
AbsDetect(time)           = offAbs;
Vacant2(time)             = 0;
Vacant3(time)             = 0;
HLeftRelay(time)         = disconnect;
HRightRelay(time)        = disconnect;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Cooling
time                       = (DTransfer+1):Cooling;
Bias(time)                = 0;
MacroU(time)              = 0;
YOffset(time)             = 0;
TopOffset(time)           = 0;
DTrap(time)               = 0;
GuideWire(time)           = 0;
HLeft(time)               = 0;
HRight(time)              = 0;
TrapLasers(time)         = on;
AbsDetect(time)           = offAbs;
Vacant2(time)             = 0;
Vacant3(time)             = 0;
HLeftRelay(time)         = disconnect;
HRightRelay(time)        = disconnect;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% FieldsOff
time                       = (Cooling+1):FieldsOff;
Bias(time)                = 0;
MacroU(time)              = 0;
YOffset(time)             = 0;
TopOffset(time)           = 0;
DTrap(time)               = 0;
GuideWire(time)           = 0;
HLeft(time)               = 0;
HRight(time)              = 0;
TrapLasers(time)         = off;
AbsDetect(time)           = offAbs;
Vacant2(time)             = 0;
Vacant3(time)             = 0;
HLeftRelay(time)         = disconnect;
HRightRelay(time)        = disconnect;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Optical Dumping
time                       = (FieldsOff+1):OpticalDumping;
Bias(time)                = BiasOD;
MacroU(time)              = 0;
YOffset(time)             = 0;
TopOffset(time)           = 0;

```

```

DTrap(time) = 0;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%%%% DTrap
z=-2:4/DTrapTime:2-4/DTrapTime;
time = (OpticalDumping+1):DTrapping;
Bias(time) =(BiasDMTrap-BiasDMTrapStart)/2.*...
    (erf(z)-erf(-2))+BiasDMTrapStart;
MacroU(time) = 0;
YOffset(time) = YOffsetDTrap;
TopOffset(time) = TopOffsetDTrap;
DTrap(time) = (DMTrap-DMTrapStart)/2.*...
    (erf(z)-erf(-2))+DMTrapStart;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%%%% MagShift
time = (DTrapping+1):MagShift;
Bias(time) = BiasDMShift;
MacroU(time) = 0;
YOffset(time) = YOffsetMShift;
TopOffset(time) = TopOffsetMShift;
DTrap(time) = DMShift;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%%%% Rotate
w=0:RotateTime-1;

```

```

time = (MagShift+1):Rotating;
Bias(time) = (BiasRotate-BiasDMShift).*...
    (1-cos((trunc.*pi.*w)./(2*RotateTime)))+BiasDMShift;
MacroU(time) = 0;
YOffset(time) = (YOffsetRFin-YOffsetRStart).*...
    sin((trunc.*pi.*w)./(2*RotateTime))+YOffsetRStart;
TopOffset(time) = TopOffsetR;
DTrap(time) = DMRotate;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%%%%%% RevRotate
w=0:RevRotateTime-1;
time = (Rotating+1):RevRotating;
Bias(time) = (BiasRotate-BiasDMTrap).*...
    (1-sin((trunc.*pi.*w)./(2*RevRotateTime)))+BiasDMTrap;
MacroU(time) = 0;
YOffset(time) = (YOffsetRFin-YOffsetRStart).*...
    (cos((trunc.*pi.*w)./(2*RevRotateTime)))+YOffsetRStart;
TopOffset(time) = TopOffsetR;
DTrap(time) = DMRotate;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%%%%%% HTransfer
y=-2:4/HTransTime:2-4/HTransTime;
time = (RevRotating+1):HTransferring;
Bias(time) = BiasHTrans;
MacroU(time) = 0;
YOffset(time) = (YOffsetHTrans-YOffsetRFin)/2.*...
    (erf(y)-erf(-2))+YOffsetRFin;
TopOffset(time) = TopOffsetH;
DTrap(time) = (0-DMRotate)/2.*...
    (erf(y)-erf(-2))+DMRotate;
GuideWire(time) = (GuideWireTrans-0)/2.*(erf(y)-erf(-2));

```

```

HLeft(time) = (HLeftTrans-0)/2.*(erf(y)-erf(-2))+0;
HRight(time) = (HRightTrans-0)/2.*(erf(y)-erf(-2))+0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = HLeftTransRelay;
HRightRelay(time) = HRightTransRelay;
%%%%%%%%%% HTrap
y=-2:4/HTrapTime:2-4/HTrapTime;
time = (HTransferring+1):HTrapping;
Bias(time) = BiasH;
MacroU(time) = 0;
YOffset(time) = YOffsetHTrap;
TopOffset(time) = TopOffsetH;
DTrap(time) = 0;
GuideWire(time) = (GuideWireHold-...
    GuideWireTrans)/2.*(erf(y)-erf(-2))+GuideWireTrans;
HLeft(time) = (HLeftHold-HLeftTrans)/2.*...
    (erf(y)-erf(-2))+HLeftTrans;
HRight(time) = (HRightHold-HRightTrans)/2.*...
    (erf(y)-erf(-2))+HRightTrans;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = HLeftHoldRelay;
HRightRelay(time) = HRightHoldRelay;
%%%%%%%%%% Drop
time = (HTrapping+1):Drop;
Bias(time) = 0;
MacroU(time) = 0;
YOffset(time) = 0;
TopOffset(time) = 0;
DTrap(time) = 0;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
if det==0
%%%%%%%%%% FluorImage
time = (Drop+1):FluorImage;

```

```

Bias(time) = 0;
MacroU(time) = 0;
YOffset(time) = 0;
TopOffset(time) = 0;
DTrap(time) = 0;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = on;
AbsDetect(time) = offAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
%%%%%%%% Capture
time = (FluorImage+1):CaptureF;
Bias(time) = BiasFin;
MacroU(time) = MacroUFin;
YOffset(time) = YOffsetFin;
TopOffset(time) = TopOffsetFin;
DTrap(time) = DTrapFin;
GuideWire(time) = GuideWireFin;
HLeft(time) = HLeftFin;
HRight(time) = HRightFin;
TrapLasers(time) = TrapLaserFin;
AbsDetect(time) = AbsDetectFin;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;
HRightRelay(time) = disconnect;
TotalTime = [0:CaptureF-1] ./ms;
elseif det==1
%%%%%%%% AbsImage
time = (Drop+1):AbsImage;
Bias(time) = 0;
MacroU(time) = 0;
YOffset(time) = YOffsetAbsImage;
TopOffset(time) = 0;
DTrap(time) = 0;
GuideWire(time) = 0;
HLeft(time) = 0;
HRight(time) = 0;
TrapLasers(time) = off;
AbsDetect(time) = onAbs;
Vacant2(time) = 0;
Vacant3(time) = 0;
HLeftRelay(time) = disconnect;

```

```

HRightRelay(time)           = disconnect;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Capture
time                        = (AbsImage+1):CaptureA;
Bias(time)                  = BiasFin;
MacroU(time)                = MacroUFin;
YOffset(time)              = YOffsetFin;
TopOffset(time)            = TopOffsetFin;
DTrap(time)                 = DTrapFin;
GuideWire(time)            = GuideWireFin;
HLeft(time)                 = HLeftFin;
HRight(time)                = HRightFin;
TrapLasers(time)           = TrapLaserFin;
AbsDetect(time)            = AbsDetectFin;
Vacant2(time)               = 0;
Vacant3(time)               = 0;
HLeftRelay(time)           = disconnect;
HRightRelay(time)          = disconnect;
TotalTime = [0:CaptureA-1]./ms;
end
%-----Plots-----
figure(1)
plot(TotalTime,2*Bias,'b',TotalTime,MacroU,'k',TotalTime,...
     5*TopOffset,'g')
legend('Bias','MacroU','TopOffset',2);
axis([0 max(TotalTime) -5 35])
figure(2)
plot(TotalTime,-1*YOffset,'r',TotalTime,DTrap,'k',TotalTime,...
     GuideWire,'g',TotalTime,HLeftRelay,'b',TotalTime,HRightRelay,'c');
legend('YOffset','DTrap','GuideWire','HLeft','HRight',3);
axis([0 max(TotalTime) -5 8])
figure(3)
AbsDetectPlot=AbsDetect;
TrapLasersPlot=TrapLasers;
plot(TotalTime,AbsDetectPlot,'y',TotalTime,TrapLasersPlot,'r')
axis([0 max(TotalTime) -0.25 1.25])
legend('AbsDetect','TrapLasers',2)
%-----Conversions to Volts-----
Bias          = Bias.*I_Bias;
MacroU        = MacroU.*I_MacroU;
YOffset       = YOffset.*I_Y;
TopOffset     = TopOffset.*I_Top;
DTrap         = DTrap.*I_DTrap;
GuideWire     = GuideWire.*I_GuideWire;
HLeft         = HLeft.*I_HLeft;
HRight        = HRight.*I_HRight;
%-----Errors-----
if max(MacroU) > 3.2

```

```

        error('MacroU too high')
    end
    if max(DTrap) > 0.875
        error('DTrap too high')
    end
    if max(GuideWire) > 0.875
        error('GuideWire too high')
    end
    if max(HLeft) > 0.875
        error('HLeft too high')
    end
    if max(HRight) > 0.875
        error('HRight too high')
    end
    %-----Binary-----
    BinPortA = TrapLasers.*1 + AbsDetect.*10 + HLeftRelay.*100 + ...
        HRightRelay.*1000 + 0.*1000 + 0.*10000 + 0.*100000 + 1.*1000000;
    for i=1:length(TrapLasers)
        StrPortA = num2str(BinPortA(i));
        PortA(i) = bin2dec(StrPortA);
    end
    %-----Files-----
    fid1 = fopen('D0.dat','w');
    fprintf(fid1,'%f\n',DortA(1));
    fclose(fid1);
    fid1 = fopen('Bias.dat','w');
    fprintf(fid1,'%f\n',Bias(1));
    fclose(fid1);
    fid1 = fopen('MacroU.dat','w');
    fprintf(fid1,'%f\n',MacroU(1));
    fclose(fid1);
    fid1 = fopen('YOffset.dat','w');
    fprintf(fid1,'%f\n',YOffset(1));
    fclose(fid1);
    fid1 = fopen('TopOffset.dat','w');
    fprintf(fid1,'%f\n',TopOffset(1));
    fclose(fid1);
    fid1 = fopen('HLeft.dat','w');
    fprintf(fid1,'%f\n',HLeft(1));
    fclose(fid1);
    fid1 = fopen('Vacant2.dat','w');
    fprintf(fid1,'%f\n',Vacant2(1));
    fclose(fid1);
    fid1 = fopen('GuideWire.dat','w');
    fprintf(fid1,'%f\n',GuideWire(1));
    fclose(fid1);
    fid1 = fopen('DTrap.dat','w');

```

```
fprintf(fid1,'%f\n',DTrap(1));
fclose(fid1);
fid1 = fopen('Vacant3.dat','w');
fprintf(fid1,'%f\n',Vacant3(1));
fclose(fid1);
fid1 = fopen('HRight.dat','w');
fprintf(fid1,'%f\n',HRight(1));
fclose(fid1);
fid1 = fopen('D0.dat','a');
for i=2:length(DortA)
fprintf(fid1,'%f\n',DortA(i));
end
fclose(fid1);
fid1 = fopen('Bias.dat','a');
for i=2:length(Bias)
fprintf(fid1,'%f\n',Bias(i));
end
fclose(fid1);
fid1 = fopen('MacroU.dat','a');
for i=2:length(MacroU)
fprintf(fid1,'%f\n',MacroU(i));
end
fclose(fid1);
fid1 = fopen('YOffset.dat','a');
for i=2:length(YOffset)
fprintf(fid1,'%f\n',YOffset(i));
end
fclose(fid1);
fid1 = fopen('TopOffset.dat','a');
for i=2:length(TopOffset)
fprintf(fid1,'%f\n',TopOffset(i));
end
fclose(fid1);
fid1 = fopen('HLeft.dat','a');
for i=2:length(HLeft)
fprintf(fid1,'%f\n',HLeft(i));
end
fclose(fid1);
fid1 = fopen('Vacant2.dat','a');
for i=2:length(Vacant2)
fprintf(fid1,'%f\n',Vacant2(i));
end
fclose(fid1);
fid1 = fopen('GuideWire.dat','a');
for i=2:length(GuideWire)
fprintf(fid1,'%f\n',GuideWire(i));
end
```

```
fclose(fid1);
fid1 = fopen('DTrap.dat','a');
for i=2:length(DTrap)
fprintf(fid1,'%f\n',DTrap(i));
end
fclose(fid1);
fid1 = fopen('Vacant3.dat','a');
for i=2:length(Vacant3)
fprintf(fid1,'%f\n',Vacant3(i));
end
fclose(fid1);
fid1 = fopen('HRight.dat','a');
for i=2:length(HRight)
fprintf(fid1,'%f\n',HRight(i));
end
fclose(fid1);
```

Bibliography

- [1] H. Mabuchi, M. Armen, B. Lev, M. Loncar, J. Vuckovic, H. J. Kimble, J. Preskill, M. Roukes, and A. Scherer. Quantum networks based on cavity QED. *Quantum Information and Computation* **1**, Special Issue on Implementation of Quantum Computation, 7, (1999).
- [2] J. Reichel, W. Hänsel, and T. W. Hänsch. Atomic micromanipulation with magnetic surface traps. *Phys. Rev. Lett.*, **83** 3398, (1999).
- [3] J. D. Weinstein and K. G. Libbrecht. Microscopic magnetic traps for neutral atoms. *Phys. Rev. A*, **52** 4004, (1995).
- [4] D. Müller, D. Anderson, R. Grow, P. Schwindt, and E. Cornell. Guiding neutral atoms around curves with lithographically patterned current-carrying wires. *Phys. Rev. Lett.*, **83** 5194, (1999).
- [5] D. Cassettari, B. Hessmo, R. Folman, T. Maier, and J. Schmiedmayer. Beam splitter for guided atoms. *Phys. Rev. Lett.*, **85** 5483, (2000).
- [6] 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, (2000).
- [7] R. Folman, P. Krüger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer. Controlling cold atoms using nanofabricated surfaces: Atom chips. *Phys. Rev. Lett.*, **84** 4749, (2000).
- [8] B. Lev. Fabrication of micro-magnetic traps for cold neutral atoms. *Quantum Inf. Comput.*, **3** 450, (2003).

- [9] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi. Adaptive homodyne measurement of optical phase. *Phys. Rev. Lett.*, **89** 133602, (2002).
- [10] E. A. Hinds and I. G. Hughes. Magnetic atom optics: mirrors, guides, traps, and chips for atoms. *J. Phys. D: Appl. Phys.*, **32** R119, (1999).
- [11] A. I. Sidorov, R. J. McLean, F. Scharnberg, D. S. Gough, T. J. Davis, B. J. Sexton, G. I. Opat, and P. Hannaford. Permanent-magnet microstructures for atom optics. *Acta Physica Polonica B*, **33** 2137, (2002).
- [12] B. Lev, Y. Lassailly, C. Lee, A. Scherer, and H. Mabuchi. Atom mirror etched from a hard drive. *Appl. Phys. Lett.*, **83** 395, (2003).
- [13] J. Schmiedmayer, R. Folman, and T. Calarco. Quantum information processing with neutral atoms on an atom chip. *Journal of Modern Optics*, **49** 1375, (2002).
- [14] E.A. Hinds, M.G. Boshier, and I.G. Hughes. Magnetic waveguide for trapping cold atom gases in two dimensions. *Phys. Rev. Lett.*, **80** 645, (1998).
- [15] S. Xue. Seagate Technology, Private communication, February 2004.
- [16] W. Hänsel, J. Reichel, P. Hommelhoff, and T. W. Hänsch. Trapped-atom interferometer in a magnetic microtrap. *Phys. Rev. A*, **64** 063607, (2001).
- [17] S. Du, M. B. Squires, Y. Imai, L. Czaia, R. A. Saravanan, V. Bright, J. Reichel, T. W. Hänsch, and D. Z. Anderson. Atom-chip Bose-Einstein condensation in a portable vacuum cell. *Phys. Rev. A*, **70** 053606, (2004).
- [18] P. Treutlein, P. Hommelhoff, T. Steinmetz, T. W. Hänsch, and J. Reichel. Coherence in microchip traps. *Phys. Rev. Lett.*, **92** 203005, (2004).
- [19] J. E. Williams. Optimal conditions for observing Josephson oscillations in a double-well Bose-Einstein condensate. *Phys. Rev. A*, **64** 013610, (2001).
- [20] D. W. Vernooy, V. S. Ilchenko, H. Mabuchi, E. W. Streed, and H. J. Kimble. High-Q measurements of fused-silica microspheres in the near-infrared. *Opt. Lett.*, **23** 247, (1998).

- [21] D. W. Vernooy. *Cold atoms in cavity QED for quantum information processing*. PhD thesis, California Institute of Technology, 2000.
- [22] J. R. Buck. *Cavity QED in Microsphere and Fabry-Perot Cavities*. PhD thesis, California Institute of Technology, 2003.
- [23] 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** 1, (2003).
- [24] J. Vučković, M Lončar H. Mabuchi, and A. Scherer. Design of photonic crystal microcavities for cavity QED. *Phys. Rev. E*, **65** 016608, (2001).
- [25] H. J. Carmichael. *An Open Systems Approach to Quantum Optics*. Springer-Verlag, Berlin, 1993.
- [26] H. M. Wiseman and G. J. Milburn. Quantum theory of field-quadrature measurements. *Phys. Rev. A*, **47** 642, (1993).
- [27] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan. Quantum feedback control and classical control theory. *Phys. Rev. A*, **62** 012105, (2000).
- [28] D. A. Steck, K. Jacobs, H. Mabuchi, T. Bhattacharya, and S. Habib. Quantum feedback control of atomic motion in an optical cavity. *Phys. Rev. Lett.*, **92** 223004, (2004).
- [29] V. Vuletić, A. T. Black, and J. K. Thompson. External-feedback laser cooling of gases, 2004, quant-ph/0410168.
- [30] J. E. Reiner, H. M. Wiseman, and H. Mabuchi. Quantum jumps between dressed states: A proposed cavity-QED test using feedback. *Phys. Rev. A*, **67** 042106, (2003).
- [31] J. E. Reiner, W. P. Smith, L. A. Orozco, H. M. Wiseman, and J. Gambetta. Quantum feedback in a weakly driven cavity QED system. *Phys. Rev. A*, **51**, (2004).
- [32] M. Armen and H. Mabuchi. To be published, 2005.
- [33] H. Mabuchi and A. C. Doherty. Cavity quantum electrodynamics: Coherence in context. *Science*, **298** 1372, (2002).

- [34] T. Pellizzari, S. A. Gardiner, J. I. Cirac, and P. Zoller. Decoherence, continuous observation, and quantum computing: A cavity QED model. *Phys. Rev. Lett.*, **75** 3788, (1995).
- [35] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Phys. Rev. Lett.*, **78** 3221, (1997).
- [36] S. J. van Enk, J. I. Cirac, and P. Zoller. Photonic channels for quantum communication. *Science*, **279** 205, (1998).
- [37] C. J. Hood, T. W. Lynn, A. C. Doherty, A. S. Parkins, and H. J. Kimble. The atom-cavity microscope: Single atoms bound in orbit by single photons. *Science*, **287** 1447, (2000).
- [38] P. W. H. Pinkse, T. Fischer, P. Maunz, and G. Rempe. Trapping an atom with single photons. *Nature*, **404** 365, (2000).
- [39] J. A. Sauer, K. M. Fortier, M. S. Chang, C. D. Hamley, and M. S. Chapman. Cavity QED with optically transported atoms. *Phys. Rev. A*, **69** 051804, 2004.
- [40] J. McKeever, J. R. Buck, A. D. Boozer, A. Kuzmich, H. C. Nägerl, D. M. Stamper-Kurn, and H. J. Kimble. State-insensitive cooling and trapping of single atoms in an optical cavity. *Phys. Rev. Lett.*, **90** 133602, (2003).
- [41] J. McKeever, A. Boca, A. D. Boozer, J. R. Buck, and H. J. Kimble. Experimental realization of a one-atom laser in the regime of strong coupling. *Nature*, **425** 268, (2003).
- [42] J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, and H. J. Kimble. Deterministic generation of single photons from one atom trapped in a cavity. *Science*, **303** 1992, (2004).
- [43] A. Boca, R. Miller, K. M. Birnbaum, A. D. Boozer, J. McKeever, and H. J. Kimble. Observation of the vacuum Rabi spectrum for one trapped atom. *Phys. Rev. Lett.*, **93** 233603, (2004).

- [44] B. Lev, K. Srinivasan, P. Barclay, O. Painter, and H. Mabuchi. Feasibility of detecting single atoms using photonic bandgap cavities. *Nanotechnology*, **15** S556, (2004).
- [45] A. Hopkins, B. Lev, and H. Mabuchi. Proposed magneto-electrostatic ring trap for neutral atoms. *Phys. Rev. A*, **70** 053616, (2004).
- [46] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, New York, 1999.
- [47] 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, (2000).
- [48] R. Raussendorf and H.-J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, **86** 5188, (2001).
- [49] M. Drndić, K. S. Johnson, J. H. Thywissen, M. Prentiss, and R. M. Westervelt. Micro-electromagnets for atom manipulation. *Appl. Phys. Lett.*, **72** 2906, (1998).
- [50] J. Reichel, W. Hänsel, P. Hommelhoff, and T. W. Hänsch. Applications of integrated magnetic microtraps. *Appl. Phys. B*, **72** 81, (2001).
- [51] R. Folman, P. Krüger, J. Schmiedmayer, J. Denschlag, and C. Henkel. Microscopic atom optics: From wires to an atom chip. *Adv. At. Mol. Opt. Phys.*, **48** 263, (2002).
- [52] D. Müller, E. Cornell, M. Prevedelli, P. Schwindt, Y. Wang, and D. Anderson. Magnetic switch for integrated atom optics. *Phys. Rev. A*, **63** 041602, (2001).
- [53] K. S. Johnson, M. Drndić, J. H. Thywissen, G. Zabow, R. M. Westervelt, and M. Prentiss. Atomic deflection using an adaptive microelectromagnet mirror. *Phys. Rev. Lett.*, **81** 1137, (1998).
- [54] P. Horak, B. Klappauf, A. Haase, R. Folman, J. Schmiedmayer, P. Domokos, and E. A. Hinds. Towards single-atom detection on a chip. *Phys. Rev. A*, **67** 043806, (2003).
- [55] Y. Louyer, D. Meschede, and A. Rauschenbeutel. Tunable whispering gallery mode resonators for cavity quantum electrodynamics, 2005, quant-ph/0501002.

- [56] M. Rosenblit, P. Horak, S. Hellsby, and R. Folman. Single-atom detection using whispering-gallery modes of microdisk resonators. *Phys. Rev. A*, **70** 053808, (2004).
- [57] S. M. Spillane, T. J. Kippenberg, K. J. Vahala, K. W. Goh, E. Wilcut, and H. J. Kimble. Ultrahigh-Q toroidal microresonators for cavity quantum electrodynamics. *Phys. Rev. A*, **71** 013817, (2005).
- [58] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, **417** 709, (2002).
- [59] M. A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. Transport of quantum states and separation of ions in a dual rf ion trap. *Quantum Inf. Comput.*, **2** 257, (2002).
- [60] W. Hänsel, P. Hommelhoff, T. W. Hänsch, and J. Reichel. Bose-Einstein condensation on a microelectronic chip. *Nature*, **413** 498, (2001).
- [61] H. Ott, J. Fortagh, G. Schlotterbeck, A. Grossmann, and C. Zimmermann. Bose-Einstein condensation in a surface microtrap. *Phys. Rev. Lett.*, **87** 230401, (2001).
- [62] A. E. Leanhardt, Y. Shin, A. P. Chikkatur, D. Kielpinski, W. Ketterle, and D. E. Pritchard. Bose-Einstein condensates near a microfabricated surface. *Phys. Rev. Lett.*, **90** 100404, (2003).
- [63] S. Schneider, A. Kasper, Ch. vom Hagen, M. Bartenstein, B. Engeser, T. Schumm, I. Bar-Joseph, R. Folman, L. Feenstra, and J. Schmiedmayer. Bose-Einstein condensation in a simple microtrap. *Phys. Rev. A*, **67** 023612, (2003).
- [64] S. Wildermuth, P. Krüger, C. Becker, M. Brajdic, S. Haupt, A. Kasper, R. Folman, and J. Schmiedmayer. Optimized magneto-optical trap for experiments with ultracold atoms near surfaces. *Phys. Rev. A*, **69** 030901R, (2004).
- [65] T. Bergeman, G. Erez, and H. Metcalf. Magnetostatic trapping fields for neutral atoms. *Phys. Rev. A*, **59** 1535, (1987).
- [66] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. Trapping of neutral sodium atoms with radiation pressure. *Phys. Rev. Lett.*, **35** 2631, (1987).

- [67] C. Cohen-Tannoudji. *Atom-Photon Interactions*. John Wiley & Sons, Inc., 1992.
- [68] J. Reichel. Microchip traps and Bose-Einstein condensation. *Appl. Phys. B*, **75** 469, (2002).
- [69] J. Fortágh, S. Kraft, A. Günther, C. Trüch, P. Wicke, and C. Zimmermann. Perspectives of ultracold atoms trapped in magnetic micro potentials. *Optics Comm.*, **243** 45, (2004).
- [70] J. Schmiedmayer. Quantum wires and quantum dots for neutral atoms. *Eur. Phys. J. D*, **4** 57, (1998).
- [71] J. Denschlag, D. Cassettari, A. Chenet, S. Schneider, and J. Schmiedmayer. A neutral atom and a wire: towards mesoscopic atom optics. *Appl. Phys. B*, **69** 291, (1999).
- [72] D. Cassettari, A. Chenet, R. Folman, A. Haase, B. Hessmo, P. Krüger, T. Maier, S. Schneider, T. Calarco, and J. Schmiedmayer. Micromanipulation of neutral atoms with nanofabricated structures. *Appl. Phys. B*, **70** 721, (2000).
- [73] E. A. Hinds, C. J. Vale, and M. G. Boshier. Two-wire waveguide and interferometer for cold atoms. *Phys. Rev. Lett.*, **86** 1462, (2002).
- [74] X. Lou, P. Krueger, K. Brugger, S. Wildermuth, H. Gimpel, M.W. Klein, S. Groth, R. Folman, I. Bar-Joseph, and J. Schmiedmayer. An atom fiber for guiding cold neutral atoms. *Optics Letters*, **29** 2145, (2004).
- [75] S. Kraft, A. Günther, P. Wicke, B. Kasch, C. Zimmermann, and J. Fortágh. Atom-optical elements on micro chips, 2005, cond-mat/0504242.
- [76] R. Long and J. Reichel. Private communication, 2002.
- [77] W. Hänsel, J. Reichel, P. Hommelhoff, and T. W. Hänsch. Magnetic conveyor belt for transporting and merging trapped atom clouds. *Phys. Rev. Lett.*, **86** 608, (2001).
- [78] V. Vuletic, T. Fischer, M. Praeger, T. W. Hänsch, and C. Zimmermann. Microscopic magnetic quadrupole trap for neutral atoms with extreme adiabatic compression. *Phys. Rev. Lett.*, **80** 1634, (1998).

- [79] M. Vengalattore, W. Rooijackers, and M. Prentiss. Ferromagnetic atom guide with in situ loading. *Phys. Rev. A*, **66** 053403, (2002).
- [80] P. Krüger, X. Luo, M.W. Klein, K. Brugger, A. Haase, S. Wildermuth, S. Groth, I. Bar-Joseph, R. Folman, and J. Schmiedmayer. Trapping and manipulating neutral atoms with electrostatic fields. *Phys. Rev. Lett.*, **91** 233201, (2003).
- [81] 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, 2005, quant-ph/0507047.
- [82] D. A. Steck. Cesium D line data. available at <http://steck.us/alkalidata>, 2003.
- [83] K. Birnbaum. *Cavity QED with Multi-Level Atoms*. PhD thesis, California Institute of Technology, 2005.
- [84] E. R. I. Abraham and E. A. Cornell. Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems. *Applied Optics*, **37** 1762, (1998).
- [85] C. E. Wieman and L. Hollberg. Using diode lasers for atomic physics. *Rev. Sci. Instrum.*, **62** 1, (1991).
- [86] K. G. Libbrecht and J. L. Hall. A low-noise high-speed diode laser current controller. *Rev. Sci. Instrum.*, **64** 2133, (1993).
- [87] J. H. Shirley. Modulation transfer processes in optical heterodyne saturation spectroscopy. *Opt. Lett.*, **7** 537, (1982).
- [88] K. Singer, S. Jochim, M. Mudrich, A. Mosk, and M. Weidemüller. Low-cost mechanical shutter for light beams. *Rev. Sci. Instrum.*, **73** 4402, (2002).
- [89] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski. Simple scheme for tunable frequency offset locking of two lasers. *Rev. Sci. Instrum.*, **70** 242, (1999).
- [90] M. J. Madou. *Fundamentals of microfabrication*. CRC Press, New York, 2001.
- [91] S. Groth, P. Krüger, S. Wildermuth, R. Folman, T. Fernholz, J. Schmiedmayer, D. Mahalu, and I. Bar-Joseph. Atom chips: Fabrication and thermal properties. *Appl. Phys. Lett.*, **85** 2980, (2004).

- [92] M. Drndić and C. S. Lee and R. M. Westervelt. Three-dimensional microelectromagnet traps for neutral and charged particles. *Phys. Rev. B*, **63** 085321, (2001).
- [93] S. Kraft, A. Günther, H. Ott, D. Wharam, C. Zimmermann, and J. Fortágh. Anomalous longitudinal magnetic field near the surface of copper conductors. *J. Phys. B*, **35** 469, (2002).
- [94] J. Fortágh, H. Ott, S. Kraft, A. Günther, and C. Zimmermann. Surface effects in magnetic microtraps. *Phys. Rev. A*, **66** 041604, (2002).
- [95] 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, (2002).
- [96] 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, (2003).
- [97] J. Estève, C. Aussibal, T. Schumm, C. Figl, D. Maily, I. Bouchoule, C. I. Westbrook, and A. Aspect. Role of wire imperfections in micromagnetic traps for atoms. *Phys. Rev. A*, **70** 043629, (2004).
- [98] T. Weber, J. Herbig, and R. Grimm M. Mark H.-C. Nägerl. Bose-Einstein condensation of cesium. *Science*, **305** 232, (2003).
- [99] J. Weiner, V. Bagnato, S. Zilio, and P. Julienne. Experiments and theory in cold and ultracold collisions. *Rev. Mod. Phys.*, **71** 1, (1999).
- [100] D. Guéry-Odelin, J. Söding, P. Desbiolles, and J. Dalibard. Is Bose-Einstein condensation of atomic cesium possible? *Europhys. Lett.*, **44** 25, (1998).
- [101] V. Vuletić, A. Kerman, C. Chin, and S. Chu. Observation of low-field Feshbach resonances in collisions of cesium atoms. *Phys. Rev. Lett.*, **82** 1406, (1999).
- [102] C. Chin, V. Vuletić, A. Kerman, and S. Chu. High resolution feshbach spectroscopy of cesium. *Phys. Rev. Lett.*, **85** 2717, (2000).

- [103] E. Tiesinga, A. J. Moerdijk, B. J. Verhaar, and H. T. C. Stoof. Conditions for Bose-Einstein condensation in magnetically trapped atomic cesium. *Phys. Rev. A*, **46** R1167, (1992).
- [104] V. Vuletić, C. Chin, A. Kerman, and S. Chu. Suppression of atomic radiative collisions by tuning the ground state scattering length. *Phys. Rev. Lett.*, **83** 943, (1999).
- [105] P. Leo, C. Williams, and P. Julienne. Collision properties of ultracold ^{133}Cs atoms. *Phys. Rev. Lett.*, **85** 2721, (2000).
- [106] M. Arndt, M. Ben Dahan, D. Guéry-Odelin, M. W. Reynolds, and J. Dalibard. Observation of a zero-energy resonance in Cs-Cs collisions. *Phys. Rev. Lett.*, **79** 625, (1997).
- [107] C. Chin, V. Vuletić, A. Kerman, S. Chu, E. Tiesinga, P. Leo, and C. Williams. Precision Feshbach spectroscopy of ultracold Cs_2 . *Phys. Rev. A*, **70** 032701, (2004).
- [108] J. M. Vogels, C. C. Tsai, R. S. Freeland, S. J. J. M. F. Kokkelmans, B. J. Verhaar, and D. J. Heinzen. Prediction of Feshbach resonances in collisions of ultracold rubidium atoms. *Phys. Rev. A*, **56** 1067, (1997).
- [109] A. Marte, T. Voltz, J. Schuster and S. Dürr, G. Rempe, E. G. M. van Kempen, and B. J. Verhaar. Feshbach resonances in rubidium 87: Precision measurement and analysis. *Phys. Rev. Lett.*, **89** 283202, (2002).
- [110] M. Greiner, C. A. Regal, and D. S. Jin. Emergence of a molecular Bose-Einstein condensate from a Fermi gas. *Nature*, **426** 537, (2003).
- [111] C. A. Regal, M. Greiner, and D. S. Jin. Observation of resonance condensation of fermionic atom pairs. *Phys. Rev. Lett.*, **92** 040403, (2004).
- [112] J. Herbig, T. Kraemer, M. Mark, T. Weber, C. Chin, H.-C. Nägerl, and R. Grimm. Preparation of a pure molecular quantum gas. *Science*, **301** 1510, (2003).
- [113] D. Sesko, T. Walker, C. Monroe, A. Gallagher, and C. Wieman. Collisional losses from a light-force atom trap. *Phys. Rev. Lett.*, **63** 961, (1989).

- [114] A. Fioretti, J.H. Müller, P. Verkerk, M. Allegrini, E. Arimondo, and P. Julienne. Direct measurement of fine-structure collisional losses from a Cs magneto-optical trap. *Phys. Rev. A*, **55** R3999, (1997).
- [115] P.A. Willems and K.G. Libbrecht. Creating long-lived neutral-atom traps in a cryogenic environment. *Phys. Rev. A*, **51** 1403, (1995).
- [116] P. A. Willems, R. A. Boyd, J. L. Bliss, and K. G. Libbrecht. Stability of magneto-optical traps with large field gradients: Limits on the tight confinement of single atoms. *Phys. Rev. Lett.*, **78** 1660, (1997).
- [117] B. Ueberholz, S. Kuhr, D. Frese, D. Meschede, and V. Gomer. Counting cold collisions. *J. Phys. B*, **33** L135, (2000).
- [118] A. Gallagher and D. Pritchard. Exoergic collisions of cold Na^* -Na. *Phys. Rev. Lett.*, **63** 957, (1989).
- [119] P. Julienne and J. Vigue. Cold collisions of ground-state and excited-state alkali-metal atoms. *Phys. Rev. A*, **44** 4464, (1991).
- [120] C. Monroe, W. Swann, H. Robinson, and C. Wieman. Very cold trapped atoms in a vapor cell. *Phys. Rev. Lett.*, **65** 1571, (1990).
- [121] J. Söding, D. Guéry-Odelin, P. Desbiolles, G. Ferrari, and J. Dalibard. Giant spin relaxation of an ultracold cesium gas. *Phys. Rev. Lett.*, **80** 1869, (1998).
- [122] M. Mudrich, S. Kraft, J. Lange, A. Mosk, M. Weidemüller, and E. Tiesinga. Hyperfine-changing collisions in an optically trapped gas of ultracold cesium and lithium. *Phys. Rev. A*, **70** 062712, (2004).
- [123] D. Boiron, A. Michaud, J. M. Fournier, L. Simard, M. Sprenger, G. Grynberg, and C. Salomon. Cold and dense cesium clouds in far-detuned dipole traps. *Phys. Rev. A*, **57** R4106, (1998).
- [124] M. E. Holmes, M. Tschernack, P. A. Quinto-Su, and N. P. Bigelow. Isotopic difference in the heteronuclear loss rate in a two-species surface trap. *Phys. Rev. A*, **69** 063408, (2004).

- [125] W. Petrich, M. Anderson, J. Ensher, and E. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, **74** 3352, (1995).
- [126] A. Hope, D. Haubrich, G. Muller, W.G. Kaenders, and D. Meschede. Neutral cesium atoms in strong magnetic-quadrupole fields at sub-doppler temperatures. *Europhys. Lett.*, **22** 669, (1993).
- [127] D. Haubrich, A. Hope, and D. Meschede. A simple model for optical capture of atoms in strong magnetic quadrupole fields. *Optics Comm.*, **102** 225, (1993).
- [128] William Press, Saul Teukolsky, William Vetterling, and Brian Flannery. *Numerical Recipes in C*. Cambridge University Press, 1999.
- [129] George Arfken and Hans Weber. *Mathematical Methods for Physicists*. Academic Press, 1995.
- [130] Philip Bevington and D. Keith Robinson. *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill, 1992.
- [131] W. Ketterle and N.J. van Druten. *Advances in Atomic, Molecular, and Optical Physics*, edited by B. Bederson and H. Walther. Academic Press, 1996.
- [132] S. Shresta, E. Tiesinga, and C. Williams. Scattering length determination from trapped pairs of atoms, (2005), quant-ph/0505102.
- [133] P. Berman, editor. *Cavity Quantum Electrodynamics*. Academic Press, San Diego, 1994.
- [134] H. J. Kimble. Strong interactions of single atoms and photons in cavity QED. *Physica Scripta*, **T76** 127, (1998).
- [135] *Special issue on atom optics*, volume **8** of *Quantum Semiclass. Opt.*, (1996).
- [136] S. Nußmann, K. Murr, M. Hijlkema, B. Weber, A. Kuhn, and G. Rempe. Vacuum-stimulated cooling of single atoms in three dimesions, 2005, quant-ph/0506067.
- [137] S. Nußmann, M. Hijlkema, B. Weber, Felix Rohde, G. Rempe, and A. Kuhn. Nano positioning of single atoms in a micro cavity, 2005, quant-ph/0506088.

- [138] P. E. Barclay, K. Srinivasan, and O. Painter. Design of photonic crystal waveguides for evanescent coupling to optical fiber tapers and for integration with high-Q cavities. *J. Opt. Soc. Am. B*, **20** 2274, (2003).
- [139] P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter. Efficient input and output fiber coupling to a photonic crystal waveguide. *Optics Letters*, **29** 697, (2004).
- [140] P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter. Probing the dispersive and spatial properties of planar photonic crystal waveguide modes via highly efficient coupling from optical fiber tapers. *Appl. Phys. Lett.*, **85** 4, (2004).
- [141] P. E. Barclay, K. Srinivasan, and O. Painter. Nonlinear response of silicon photonic crystal micresonators excited via an integrated waveguide and fiber taper. *Optics Express*, **13** 801, (2005).
- [142] K. Srinivasan and O. Painter. Momentum space design of high-Q photonic crystal optical cavities. *Optics Express*, **10** 670, (2002).
- [143] K. Srinivasan, P. E. Barclay, O. Painter, J. Chen, A.Y. Cho, and C. Gmachl. Experimental demonstration of a high-Q photonic crystal microcavity. *Appl. Phys. Lett.*, **83** 1915, (2003).
- [144] K. Srinivasan, P. E. Barclay, M. Borselli, and O. Painter. Optical-fiber-based measurement of an ultrasmall volume high-Q photonic crystal microcavity. *Phys. Rev. B*, **70** 081306R, (2004).
- [145] K. Srinivasan, M. Borselli, T.J. Johnson, P. E. Barclay, O. Painter A. Stintz, and S. Krishna. Optical loss and lasing characteristics of high-quality-factor AlGaAs microdisk resonators with embedded quantum dots. *Appl. Phys. Lett.*, **86** 151106, (2005).
- [146] M. Borselli, T. J. Johnson, and O. Painter. Beyond the Rayleigh scattering limit in high-Q silicon microdisks: theory and experiment. *Optics Express*, **13** 1515, (2005).
- [147] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe. Vacuum Rabi splitting with a single quantum dot in a photonic crystal nanocavity. *Nature*, **432** 200, (2004).

- [148] J. P. Reithmaier, G. Sęk, A. Löffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel. Strong coupling in a single quantum dotsemiconductor microcavity system. *Nature*, **432** 197, (2004).
- [149] E. Peter, P. Senellart, D. Martrou, A. Lemaître, J. Hours, J. M. Gérard, and J. Bloch. Exciton-photon strong-coupling regime for a single quantum dot embedded in a microcavity. *Phys. Rev. Lett.*, **95** 067401, (2005).
- [150] J. M. Raimond, M. Brune, and S. Haroche. Colloquium: Manipulating quantum entanglement with atoms and photons in a cavity. *Rev. Mod. Phys.*, **73** 565, (2001).
- [151] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, **69** 062320, (2004).
- [152] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, **431** 162, (2004).
- [153] C. J. Hood. *Real-time measurement and trapping of single atoms by single photons*. PhD thesis, California Institute of Technology, 2000.
- [154] T. W. Lynn. *Measurement and control of individual quanta in cavity QED*. PhD thesis, California Institute of Technology, 2003.
- [155] W. von Klitzing, R. Long, V. S. Ilchenko, J. Hare, and V. Lefèvre Seguin. Tunable whispering gallery modes for spectroscopy and CQED experiments. *New Journal of Physics*, **3** 14.1, (2001).
- [156] J. Buck and H. J. Kimble. Optimal sizes of dielectric microspheres for cavity QED with strong coupling. *Phys. Rev. A*, **67** 033806, (2003).
- [157] C. J. Hood, H. J. Kimble, and J. Ye. Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity. *Phys. Rev. A*, **64** 033804, (2001).
- [158] B.-S. Song, S. Noda, T. Asano, and Y. Akahane. Ultra-high-Q photonic double-heterostructure nanocavity. *Nature Materials*, **4** 207, (2005).

- [159] J. Reichel and R. Long. Private communication, October 2002.
- [160] S. M. Tan. A quantum optics toolbox for Matlab 5. *J. Opt. B: Quantum Semiclass. Opt.*, **1**, (1999).
- [161] S. M. Tan. A quantum optics toolbox for Matlab 5.
- [162] K. M. Birnbaum, A. S. Parkins, and H. J. Kimble. Cavity QED with multiple hyperfine levels, 2005, to appear on the arXiv and submitted to PRA.
- [163] L. A. Lugiato. *Progress in Optics*. Amsterdam: Elsevier Science Publishers, 1984.
- [164] H. Mabuchi and H. M. Wiseman. Retroactive quantum jumps in a strongly coupled atom-field system. *Phys. Rev. Lett.*, **81** 4620, (1998).
- [165] H. Mabuchi. *Continuous observation of quantum dynamics*. PhD thesis, California Institute of Technology, 1998.
- [166] Q. A. Turchette. *Quantum optics with single atoms and single photons*. PhD thesis, California Institute of Technology, 1997.
- [167] H. Mabuchi, Q. A. Turchette, M. S. Chapman, and H. J. Kimble. Real-time detection of individual atoms falling through a high-finesse optical cavity. *Opt. Lett.*, **21** 1393, (1996).
- [168] A. C. Doherty, A. S. Parkins, S. M. Tan, and D. F. Walls. Motion of a two-level atom in an optical cavity. *Phys. Rev. Lett.*, **56** 833, (1997).
- [169] A. C. Doherty, T. W. Lynn, C. J. Hood, and H. J. Kimble. Trapping of single atoms with single photons in cavity QED. *Phys. Rev. A*, **63** 013401, (2000).
- [170] Y. 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, (2003).
- [171] O. Painter and P. Barclay. Private communication, 2005.
- [172] P. A. Quinto-Su, M. Tscherneck, M. Holmes, and N. P. Bigelow. On-chip optical detection of laser cooled atoms. *Optics Express*, **12** 5098, (2004).

- [173] M. Trupke, E. A. Hinds, S. Eriksson, E. A. Curtis, Z. Moktadir, E. Kukharenya, and M. Kraft. Microfabricated high-finesse optical cavity with open access and small volume, 2005, quant-ph/0506234.
- [174] E. D. Black. An introduction to pound-drever-hall laser frequency stabilization. *Am. J. Phys.*, **69** 1, (2001).
- [175] 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. *Appl. Phys. B*, **31** 97, (1983).
- [176] D. Hunger. MPQ/LMU Diploma thesis, 2005.
- [177] K. R. Brown, K. M. Dani, D. M. Stamper-Kurn, and K. B. Whaley. Deterministic optical fock-state generation. *Phys. Rev. A*, **67** 043818, (2003).
- [178] S. Leslie, N. Shenvi, K. R. Brown, D. M. Stamper-Kurn, and K. B. Whaley. Transmission spectrum of an optical cavity containing n atoms. *Phys. Rev. A*, **69** 043805, (2004).
- [179] H. Mabuchi. Private communication, 2005.
- [180] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, **415** 39, (2002).
- [181] R. B. Diener, B. Wu, M. G. Raizen, and Q. Niu. Quantum tweezer for atoms. *Phys. Rev. Lett.*, **89** 070401, (2002).
- [182] A. Öttl, S. Ritter, M. Köhl, and T. Esslinger. Correlations and counting statistics of an atom laser. *Phys. Rev. Lett.*, **95** 090404, (2005).
- [183] M.-O. Mewes, M. R. Andrews, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. Output coupler for Bose-Einstein condensed atoms. *Phys. Rev. Lett.*, **78** 582, (1997).
- [184] V. Savalli, D. Stevens, J. Estève, P.D. Featonby, V. Josse, N. Westbrook, C.I. Westbrook, and A. Aspect. Specular reflection of matter waves from a rough mirror. *Phys. Rev. Lett.*, **88** 250404, (2002).

- [185] A. Arnold, C. MacCormick, and M. Boshier. Adaptive inelastic magnetic mirror for Bose-Einstein condensates. *Phys. Rev. A*, **65** 031601, (2002).
- [186] G.I. Opat, S.J. Wark, and A. Cimmino. *Appl. Phys. B*, **54** 396, (1992).
- [187] R. Lawrence Comstock. *Introduction to Magnetism and Magnetic Recording*. Wiley, New York, 1999.
- [188] MabuchiLab Website, <http://minty.caltech.edu>.
- [189] P. Rosenbusch, B.V. Hall, I.G. Hughes, C.V. Saba, and E.A. Hinds. Manipulation of cold atoms using a corrugated magnetic reflector. *Phys. Rev. A*, **61** 031404R, (2000).
- [190] V.V. Vladimirovskii. *Sov. Phys. JETP*, **12** 740, (1961).
- [191] A. S. Arnold and E. Riis. Bose-einstein condensates in ‘giant’ toroidal magnetic traps. *Journal of Modern Optics*, **49** 959, (2002).
- [192] J. A. Sauer, M. D. Barrett, and M. S. Chapman. Storage ring for neutral atoms. *Phys. Rev. Lett.*, **87** 270401, (2001).
- [193] 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, (2003).
- [194] C. Henkel and S. Pötting. Coherent transport of matter waves. *Appl. Phys. B*, **72** 73, (2001).
- [195] C. Henkel, P. Krüger, R. Folman, and J. Schmiedmayer. Fundamental limits for coherent manipulation on atom chips. *Appl. Phys. B*, **76** 173, (2003).
- [196] V. Dikovskiy, Y. Japha, C. Henkel, and R. Folman. Reduction of magnetic noise in atom chips by material optimization, (2005), [quant-ph/0504027](https://arxiv.org/abs/quant-ph/0504027).
- [197] J. M. McGuirk, D. M. Harber, J. M. Obrecht, and E. A. Cornell. Alkali-metal adsorbate polarization on conducting and insulating surfaces probed with Bose-Einstein condensates. *Phys. Rev. A*, **69** 062905, (2004).
- [198] H. Moritz, T. Stöferle, Michael Köhl, and T. Esslinger. Exciting collective oscillations in a trapped 1D gas. *Phys. Rev. Lett.*, **91** 250402, (2003).

- [199] B. Laburthe Tolra, K. M. OHara, 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, (2004).
- [200] M. Olshanii. Atomic scattering in the presence of an external confinement and a gas of impenetrable Bosons. *Phys. Rev. Lett.*, **81** 938, (1998).
- [201] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. Regimes of quantum degeneracy in trapped 1D gases. *Phys. Rev. Lett.*, **85** 3745, (2000).
- [202] J. Reichel and J. H. Thywissen. Using magnetic chip traps to study Tonks-Girardeau quantum gases. (2003), cond-mat/0310330.
- [203] B. Paredes, A. Widera, V. Murg, O. Mandel, Simon 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, (2004).
- [204] T. Kinoshita, T. Wenger, and D. S. Weiss. Observation of a one-dimensional Tonks-Girardeau gas. *Science*, **305** 1125, (2004).
- [205] B. P. Anderson and M. A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, **282** 1686, (1998).
- [206] B. P. Anderson, K. Dholakia, and E. M. Wright. Atomic-phase interference devices based on ring-shaped Bose-Einstein condensates: Two-ring case. *Phys. Rev. A*, **67** 033601, (2003).
- [207] C. Zhang, J. Liu, M. G. Raizen, and Q. Niu. Transition to instability in a kicked Bose-Einstein condensate. *Phys. Rev. Lett.*, **92** 054101, (2004).
- [208] W. Hänsel, J. Reichel, P. Hommelhoff, and T. W. Hänsch. Trapped-atom interferometer in a magnetic microtrap. *Phys. Rev. A*, **64** 063607, (2001).
- [209] 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, (2005).
- [210] L. Pitaevskii and S. Stringari. Thermal vs quantum decoherence in double well trapped Bose-Einstein condensates. *Phys. Rev. Lett.*, **87** 180402, (2001).

- [211] J. Estève, T. Schumm, J.-B. Trebbia, I. Bouchoule, A. Aspect, and C. I. Westbrook. Realizing a stable magnetic double-well potential on an atom chip, 2005, quant-ph/0503112.
- [212] E. Andersson, T. Calarco, Ron Folman, M. Andersson, B. Hessmo, and J. Schmiedmayer. Multimode interferometer for guided matter waves. *Phys. Rev. Lett.*, **88** 100401, (2002).
- [213] 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 on an atom chip, 2005, cond-mat/0506464.
- [214] 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, (2005).
- [215] Y. Shin, M. Saba, A. Schirotzek, T. A. Pasquini, A. E. Leanhardt, D. E. Pritchard, and W. Ketterle. Distillation of Bose-Einstein condensates in a double-well potential. *Phys. Rev. Lett.*, **92** 150401, (2004).
- [216] F. Dalfovo, S. Giorgini, L. Pitaevskii, and S. Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.*, **71** 463, (1999).
- [217] S. Raghavan, A. Smerzi, S. Fantoni, and S. R. Shenoy. Coherent oscillations between two weakly coupled Bose-Einstein condensates: Josephson effects, π oscillations, and macroscopic quantum self-trapping. *Phys. Rev. A*, **620** 59, (1999).
- [218] J. Fortagh, A. Grossmann, T. W. Hänsch, and C. Zimmermann. Fast loading of a magneto-optical trap from a pulsed thermal source. *J. App. Phys.*, **84** 6499, (1998).
- [219] U. D. Rapol, A. Wasan, and V. Natarajan. Fast loading of a magneto-optical trap from a pulsed thermal source. *Phys. Rev. A*, **64** 023402, (2001).
- [220] B. P. Anderson and M. A. Kasevich. Loading a vapor-cell magneto-optic trap using light-induced atom desorption. *Phys. Rev. A*, **63** 023404, (2001).