An All Optical BEC Apparatus for the Investigation of Magnetic Order in Spinor Atomic Gases in Ultralow Magnetic Fields

by

Joshua Alan Hughes

(Under the direction of Chad Fertig)

Abstract

Bose-Einstein condensates (BECs) of spinor atomic gases have become powerful tools in quantum simulation experiments. Quantum magnetism, in particular, is a field of study that is difficult to simulate in the solid state but is readily accessible in BECs. However, there are still many quantum magnetic effects predicted to occur in spinor BECs that have yet to be observed in the laboratory. This is due to the presence of typical background magnetic fields which are sufficiently large to suppress these predicted phenomena. We have designed and built an apparatus to test a wide range of quantum magnetism parameters in the ultralow field regime. We use a two chamber ultra-high vacuum apparatus to precool atoms in a vapor cell magneto-optical trap and transfer the atoms into an ultra high vacuum magneto-optical trap, where they are efficiently loaded into a tightly confining, large volume optical dipole trap. The dipole trap will eventually be used to realize an all-optical BEC located inside a multi-layer mu-metal shield that reduces the background magnetic field to below 10^{-6} Gauss around the condensate region while still allowing good optical access for imaging and measurements. We have performed measurements of microwave-driven Rabi oscillations in the dipole trap to test the feasibility of ultralow magnetic field measurements inside the magnetic shield. We have also developed a method to dynamically control the aspect ratio of the BEC using a crossed beam dipole trap configuration we call a "scissor trap".

INDEX WORDS: Bose Einstein condensate, Laser cooling, Ultracold, Quantum simulation, Quantum magnetism, Magnetic shielding, Optical dipole trap, Scissor trap

An All Optical BEC Apparatus for the Investigation of Magnetic Order in Spinor Atomic Gases in Ultralow Magnetic Fields

by

Joshua Alan Hughes

B.A., University of Oklahoma, 2003

A Dissertation Submitted to the Graduate Faculty
of The University of Georgia in Partial Fulfillment
of the Requirements for the Degree

Doctor of Philosophy

Athens, Georgia

2013

©2013

Joshua Alan Hughes

All Rights Reserved

An All Optical BEC Apparatus for the Investigation of Magnetic Order in Spinor Atomic Gases in Ultralow Magnetic Fields

by

Joshua Alan Hughes

Major Professor: Chad D. Fertig

Committee: William M. Dennis
Michael R. Geller

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
August 2013

Dedication

For my parents.

Acknowledgments

I would like to extend my deepest gratitude to all my professors in the physics department who have shared their knowledge and stood by me throughout my entire time as a graduate student. I would especially like to thank my advisor, Dr. Chad Fertig, for introducing me to a field of cutting edge research and for all his assistance in building a lab from the ground up. Despite the frustration of numerous equipment failures along the way, he was always there to provide perspective and encouragement. I would like to thank my family, especially my parents, who have loved and supported me both emotionally and monetarily throughout my long career as a graduate student. They have always believed in me and never questioned my ability to succeed. Finally, I would also like to thank all my friends in the department, without whom I would have never lasted so long.

Contents

Dedication **iv**

Acknowledgments **v**

List of Figures **viii**

List of Tables **xi**

1 Introduction and Literature Review **1**

2 Vapor Cell and Ultra High Vacuum MOT Systems **6**

2.1 Vacuum chamber design and bakeout procedure 6

2.2 Laser frequency locking 18

2.3 VC MOT 26

2.4 Pusher/guide beam 33

2.5 UHV MOT 39

3 A widely tunable laser frequency offset lock with digital counting **43**

3.1 Abstract 44

3.2 Introduction 44

3.3 Structure of the paper 45

3.4 Comparison to other schemes 45

3.5 Description of the method 47

3.6 Quantization error limited stability of a digitally counted frequency servo loop 52

3.7 Experimental characterization 56

3.8 Conclusion 62

4 Crossed Beam Dipole Trap **63**

4.1 Overview of the apparatus 63

4.2 Scissor trap 64

4.3 Loading procedure from the UHV MOT 67

4.4 Dipole trap characterization 69

4.5 Evaporation and the "wings" problem 88

5 Magnetic Shield **100**

5.1 Design and simulation 100

5.2 Construction and initial measurements 105

5.3 Low-noise field canceling coils 113

5.4 Practical considerations for stand alone equipment inside the shield 113

6 Rabi Oscillations **115**

6.1 Rabi frequency and light shifts 115

6.2 Rabi frequency measurements 116

7 Conclusion **122**

Bibliography **123**

List of Figures

2.1 VC MOT vacuum chamber cell 7
2.2 Vacuum chamber ion pumps 9
2.3 UHV MOT vacuum chamber cell 10
2.4 Vacuum chamber in preparation for a bakeout 12
2.5 The oven used for vacuum chamber bakeouts 13
2.6 Bakeout temperature 15
2.7 Bakeout pressure 16
2.8 Bakeout current 17
2.9 Master laser saturated absorption spectroscopy beam diagram 19
2.10 Repumper laser saturated absorption spectroscopy beam diagram 20
2.11 Master laser and MOT laser energy level diagram 22
2.12 Master laser saturated absorption spectrum 23
2.13 Repumper laser energy level diagram 24
2.14 Repumper laser saturated absorption spectrum 25
2.15 MOT coils configuration 27
2.16 MOT Zeeman shift diagram 29
2.17 VC MOT loading curve 32
2.18 Pusher beam diagram 34
2.19 VC MOT loading curve in the presence of the pusher beam 39
2.20 UHV MOT loading curve 41
2.21 UHV MOT lifetime 42

3.1 Block diagram of circuit 48

3.2 Timing diagram of digital counting 50
3.3 Detecting frequency fluctuations by discretely sampled counting 52
3.4 Frequency stability 55
3.5 Long-term locking 57
3.6 Optical beat note linewidth 58
3.7 Tuning agility 59
3.8 Tuning range 61

4.1 Scissor trap optics 66
4.2 Crossed dipole trap that requires a double Gaussian fit 72
4.3 Horizontal cross section requiring a double Gaussian fit 73
4.4 Crossed dipole trap number 75
4.5 Expanded cloud image 77
4.6 Crossed trap frequencies in the tightly confined direction 80
4.7 Crossed trap frequencies in the loosely confined direction 81
4.8 Single trap frequencies for beam 1 82
4.9 Single trap frequencies for beam 2 83
4.10 Full power crossed trap lifetime 85
4.11 18/25 power crossed trap lifetime 86
4.12 12/25 power crossed trap lifetime 87
4.13 Number of atoms for a simulated evaporation ramp that reaches BEC 89
4.14 Trap temperature for a simulated evaporation ramp that reaches BEC 90
4.15 Phase space density for a simulated evaporation ramp that reaches BEC . . 91
4.16 Number of atoms for an experimental evaporation ramp 92
4.17 Trap temperature for an experimental evaporation ramp 93
4.18 Phase space density for an experimental evaporation ramp 93
4.19 Atoms migrating from the crossed dipole trap to the wings 94
4.20 Number of atoms for a simulation of evaporation taking into account the wings 96

4.21 Trap temperature for a simulation of evaporation taking into account the wings 97
4.22 Phase space density for a simulation of evaporation taking into account the wings 98
4.23 Crossed dipole trap with a tilted potential 99

5.1 Picture of the mu-metal shield 102
5.2 Magnetic shield simulation 104
5.3 Shielding factor measurements 107
5.4 Shielding factor experimental setup 108
5.5 Magnetic shield noise measurements 110
5.6 Magnetic shield noise Allan deviation 111
5.7 Imaging lens system diagram 114

6.1 Rabi oscillations 119
6.2 Simulated Rabi oscillations 121

List of Tables

5.1 Shielding factor 106

1 Introduction and Literature Review

Bose-Einstein condensates (BECs) of spinor atomic gases have become powerful tools in quantum simulation experiments. In particular, the study of magnetic order in quantum systems with internal degrees of freedom is of fundamental importance to condensed matter physics. Spinor BECs of atomic gases can potentially be used as analog systems to provide access to a range of quantum magnetism parameters inaccessible in the solid state [1, 2, 3]. A number of important quantum magnetism models with intractable analytical solutions or insufficient solid state examples exist that could potentially be simulated in a BEC, such as the 3D Bose Hubbard model, 1D Ising chain, 2D spin ladder, 2D XY magnet, and 2D Heisenberg magnet, among others [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27].

A spinor BEC in an optical lattice is a promising environment to study magnetic order in such models for several reasons. First, the microscopic interactions between atoms in an optical lattice are well known and can often be precisely tuned via well controlled applications of magnetic, electric, and laser fields. With such fine control, the atomic interactions can be adjusted to match the various couplings found in the different models. In addition, the properties of the optical lattice can be dynamically controlled by changing the laser polarization, intensity, and frequency. This allows a real time change in the tunneling strength between lattice sites and the on site interaction energies to provide a complete phase space map of the model Hamiltonian. Multiple lattice geometries are achievable, as well, which further increase the experimental possibilities to include studies of frustrated spin and other quantum magnetic phenomena related to crystal symmetry [28, 29, 30]. An optical lattice is also a defect-free environment, minimizing the external degrees of freedom acting on the sys-

tem. If desired, however, disorder can be controllably imposed by passing the optical lattice beams through a holographic diffuser to introduce an optical speckle field [31, 32, 33, 34].

The goal of this research project is to study new regimes of quantum magnetism in ultralow magnetic fields. In order for a spinor BEC in an optical lattice to be considered a good quantum simulator that can "solve" a classically intractable quantum magnetism many-body theory by simply measuring the incalculable parameters of the model, precise knowledge and control over the microscopic interactions as well as the internal and external degrees of freedom are critical. High precision measurements in an experimentally flexible apparatus will be required to fully characterize the atomic interactions. In order to achieve maximal control over the internal and external degrees of freedom, environmental perturbations must be minimized. This includes minimizing laser field intensity fluctuations and pointing instabilities by using stable laser sources and optics with low mechanical jitter. It also requires highly stable, narrow bandwidth laser frequencies that will not drift from known frequency values. For quantum simulations in the regime where the magnetic dipole-dipole interaction energy of a spinor condensate dominates, ultralow magnetic fields are predicted to induce phase transitions that break individual spin and space rotation symmetries [22, 23]. To reach such a regime, the ambient field and low frequency noise around the experiment must be cancelled with a passive mu-metal magnetic shield, with accurate manipulation of higher frequency low-level fields provided by actively controlled field-cancelling shim coils.

Ever since the first observations of the BEC phase transition in a dilute gas in 1995 [35, 36], new technological advances have made the process almost commonplace in the laboratory, now even available commercially in the form of an atom chip. The most common path to BEC involves precooling an atomic vapor to sub-mK temperatures using laser cooling techniques, and then transferring those atoms to a magnetic trap. The atoms are then further cooled below the BEC critical temperature via rf-induced evaporation of energy selective spin transitions. However, developing an experimental system that realizes a large BEC in volume and number inside a well characterized ultralow magnetic field environment remains

a challenging endeavor. As the term implies, a magnetic trap requires a magnetic field to provide the trapping potential for the atoms, making studies of BECs in ultralow magnetic fields impossible. Furthermore, atomic spins align along the direction of the magnetic field, freezing out the spin degree of freedom. Quantum magnetism experiments rely on the spin interactions between atoms, justifying the need for a spinor condensate. Another route to BEC is required.

The first all optical BEC was observed in 2001 without the presence of a magnetic trapping potential [37]. Atoms precooled with laser cooling techniques are loaded into an optical dipole trap and then evaporatively cooled by decreasing the trapping beam intensity. Evaporative cooling relies on removing high energy (i.e. "hot") atoms from the trap, while the remaining atoms rethermalize via elastic collisions to a lower average temperature. For successful evaporation in an optical dipole trap, several conditions must be met. The dipole trap must provide a tight confining potential in all dimensions to ensure high atomic densities, and thus fast rethermalization rates. The initial population in the trap must also be large since the evaporative process requires the removal of a large fraction of atoms from the trap. As an indication of the efficiency of evaporation, the ratio of elastic collisions, which are necessary for thermalization, to inelastic collisions, which heat the atoms and cause trap losses, must be large.

Unfortunately, achieving sufficient phase space densities to reach BEC in an optical dipole trap suffers from severe limitations not present in magnetic evaporation procedures. In a tightly confining dipole trap, a small trapping volume limits the number of atoms loaded into the trap but has high atomic densities. Optical dipole traps with large trapping volumes that load a large number of atoms are also possible, but at the expense of weak confinement and low elastic collision rates. Expanding on the method of [37] to overcome these issues, we use a crossed beam dipole trap to load a large number of atoms into a large trapping volume while also meeting the requirements of a tight confining potential. Unlike magnetic traps, optical dipole traps suffer from trap heating due to off resonant Rayleigh scattering of

trap photons. For each scattering event, the thermal energy in the trap increases due to the photon recoil. The evaporation rate in optical dipole traps must therefore proceed quickly to avoid significant trap loss due to heating.

To satisfy all necessary requirements, we have designed and built an apparatus that quickly loads a large number of cold atoms into a far off resonant crossed beam optical dipole trap. We use saturated absorption spectroscopy to frequency lock lasers with narrow laser linewidths necessary for high precision frequency stability required in laser cooling and measurement diagnostics. We have also developed a new digital electronic frequency offset lock to lock beams to frequencies not conveniently located near a ^{87}Rb resonance. To achieve the requisite low pressure environment for laser cooling atoms and subsequent transfer into the dipole trap, we have constructed a simple yet robust two chamber vacuum system.

Inside the vacuum, ^{87}Rb atoms are first loaded from a room temperature background vapor into a large volume vapor cell (VC) magneto-optical trap (MOT) to provide a cold source of atoms to the ultra high vacuum side of the chamber. A beam of atoms is subsequently pushed out of the VC MOT via a pusher laser beam and transferred to an ultra high vacuum (UHV) MOT with a high transfer efficiency. The UHV MOT exists in an extremely low pressure environment ($< 10^{-11}$ torr) to avoid collisions of cold trapped atoms with room temperature background atoms. A pair of tightly focused 1064 nm laser beams cross to form a dipole trap at the center of the UHV MOT. Their high power, narrow waists, long Rayleigh lengths, and shallow crossing angle all provide a large capture volume while maintaining tight trap confinement. We have also developed a method to dynamically change the crossing angle of the dipole trap beams without displacing the region of intersection. This will enable future experiments that test the predicted formation of spontaneous spin textures due to changing trap anisotropy in ultralow magnetic field environments [23]. Evaporation in the dipole trap was attempted, but with the current configuration of the dipole trapping beams it proceeds too slowly to reach BEC. We performed studies to determine the cause, and determined that atoms trapped in the non-intersecting regions of the beams play

a significant role in heating the atoms in the intersecting region. We propose that a simple reorientation of the dipole beam configuration to tilt the trapping potential will correct this problem and and allow evaporation to proceed quickly enough to reach BEC.

For experiments in ultra low magnetic field environments we have built and characterized a passive five layer cylindrical mu-metal shield that slides over the UHV vacuum chamber. The BEC will form in the center of the mu-metal shield where the background magnetic field reaches a minimum value of $\sim 1\,\mu$G. Effective canceling by the shield of low frequency field fluctuations have been observed. High frequency field fluctuations will be cancelled by an active low noise shim coil circuit to successfully realize the low field limit. Preliminary studies show the presence of Rabi oscillations in the dipole trap in the absence of the shield when driven by a microwave field of known frequency. This serves as a proof of principle experiment to show that our system is capable of driving and observing Rabi oscillations, which will be useful in measuring the Zeeman shift of the atoms due to a small magnetic field in a shielded environment.

2 Vapor Cell and Ultra High Vacuum MOT Systems

2.1 Vacuum chamber design and bakeout procedure

The vacuum chamber is one of the most critical elements in any laser cooling experiment. A chemically pure environment with good environmental isolation is a necessity, and trapped atom interactions with any background gas must be minimized as much as possible. To this end, a simple and robust vacuum chamber design that can reliably maintain pressures below 10^{-11} torr is required.

In our design, the vacuum chamber must operate under two different pressure regimes. For reasons discussed in Sections 2.3 and 2.5, the vapor cell (VC) magneto-optical trap (MOT) side of the chamber must operate at pressures on the order of 10^{-9} torr, while the ultra high vacuum (UHV) MOT side of the chamber must operate at pressures on the order of 10^{-11} torr or lower. The two orders of magnitude pressure difference between the two sides of the vacuum chamber is maintained by a differential pumping tube placed between the two sides of the chamber. In order to reach such low pressures, careful selection and preparation of all vacuum components is necessary. Because pump and chamber failure is sometimes inevitable, minimizing the impact of replacement and repair on the existing external experimental apparatus is quite useful. Therefore, a small, semi-portable vacuum chamber with as few valves and flanges as possible to minimize the opportunity for leaks is beneficial. The current experimental vacuum chamber design achieves these goals very well.

The VC MOT chamber, seen in Figure 2.1, consists of an all-glass cell with six 2 in diameter MOT beam windows and one 1 in diameter window for the guide and probe beams (the guide and probe beams are explained in greater detail in Sections 2.4 and 4.4, respectively). By using a glass cell instead of a steel cell with glass windows, we avoid the extra flange

connections required for the windows and the often troublesome process of creating good glass-to-metal seals. The glass cell has a 1 in diameter extension down towards the optical table that connects to a steel bellows to reduce stress on the the fragile glass. These bellows connect to a four-pronged storage chamber, where each prong holds a glass ampoule of solid Rb. The glass VC MOT cell also has a 1 in diameter extension along the line of sight of the probe and guide beam window that connects to another steel bellows to reduce stress that leads to the rest of the vacuum chamber.

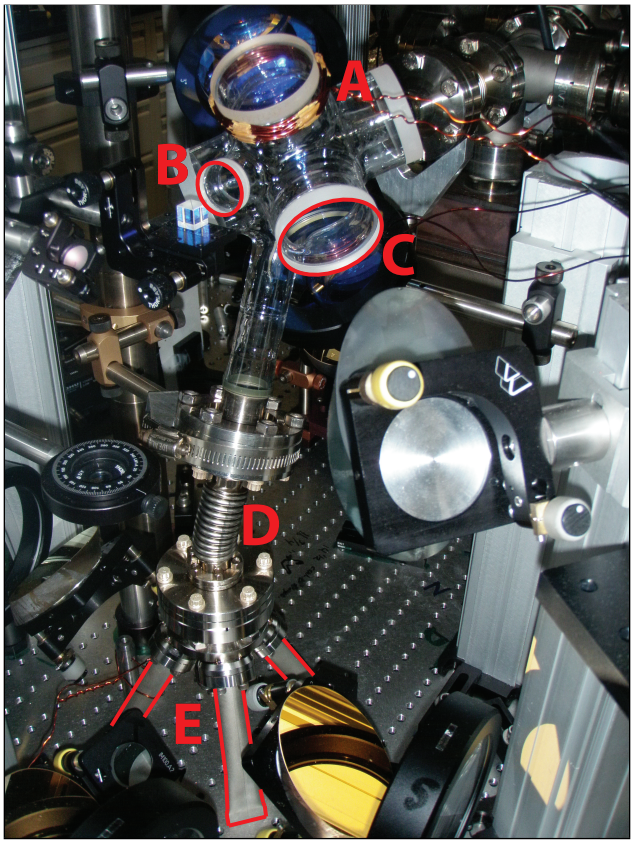

Figure 2.1: This is a labeled photo of the VC MOT vacuum chamber cell. The VC MOT forms in the glass cell (A). The guide and probe beams pass through the highlighted 1 in diameter window (B) with a direct line of sight all the way through to the end of the UHV chamber cell (not depicted). There are six 2 in diameter windows (C) for the VC MOT beams. The steel bellows (D) sits between the glass VC MOT cell and the four-pronged ^{87}Rb ampoule storage chamber (E) and serves to reduce the stress applied to the glass cell. Each prong stores one ampoule of ^{87}Rb. The remaining vacuum chamber components continue up and to the right, out of the photograph.

A V-clamp supports the steel bellows connecting the VC MOT glass cell to the stainless steel portion of the vacuum chamber. This part of the vacuum chamber, shown in Figure 2.2, consists of a Gamma Vacuum TiTan Ion Pump to provide pumping on the VC MOT side of the vacuum chamber, a differential pumping tube to maintain a two orders of magnitude pressure differential between the VC MOT and UHV MOT sides of the chamber, a Varian Vacuum Technologies StarCell Ion Pump to provide pumping on the UHV MOT side of the vacuum chamber, and a titanium sumblimation pump to provide as needed pumping of excessive background gas load buildup. The only valve present in the entire vacuum system extends off a flange connected to the Gamma ion pump. This valve is only opened during the bakeout process to connect roughing and turbo pumps until pressures are low enough in the system for the ion pumps to safely be the only pumping sources.

On the UHV MOT side of the vacuum chamber, a 1 in diameter, 18 in long glass tube extends along the line of sight of the guide and probe beams window of the VC MOT cell. To the end of the tube is attached a 1 in $\times$ 1 in $\times$ 1.5 in rectangular experiment cell. This tube is designed to slide through the center hole of the magnetic shield and extend to the center of the shield. The UHV MOT, dipole trap, and BEC are all made and held in the center of the experiment cell, shown in Figure 2.3.

In order to achieve the low pressures necessary for ultracold atom experiments, care must be taken in the assembly and bakeout of the vacuum chamber. Handling of all parts to be used in the assembly of the vacuum chamber should be done while wearing latex gloves to prevent body oils from contaminating any surfaces. Prior to assembly, all stainless steel parts must be thoroughly cleaned to remove any oils or residue left over from the manufacturing process or exposure to atmosphere. This usually involves placing each part in a powdered detergeant (such as Alconox) solution in a sonic bath for about an hour. The parts are then removed and rinsed with distilled water to remove any leftover soap. Each part is then submerged in methanol and placed in a sonic bath for another hour; this step is typically performed twice to ensure the surfaces are completely clean. Once all steel parts have been

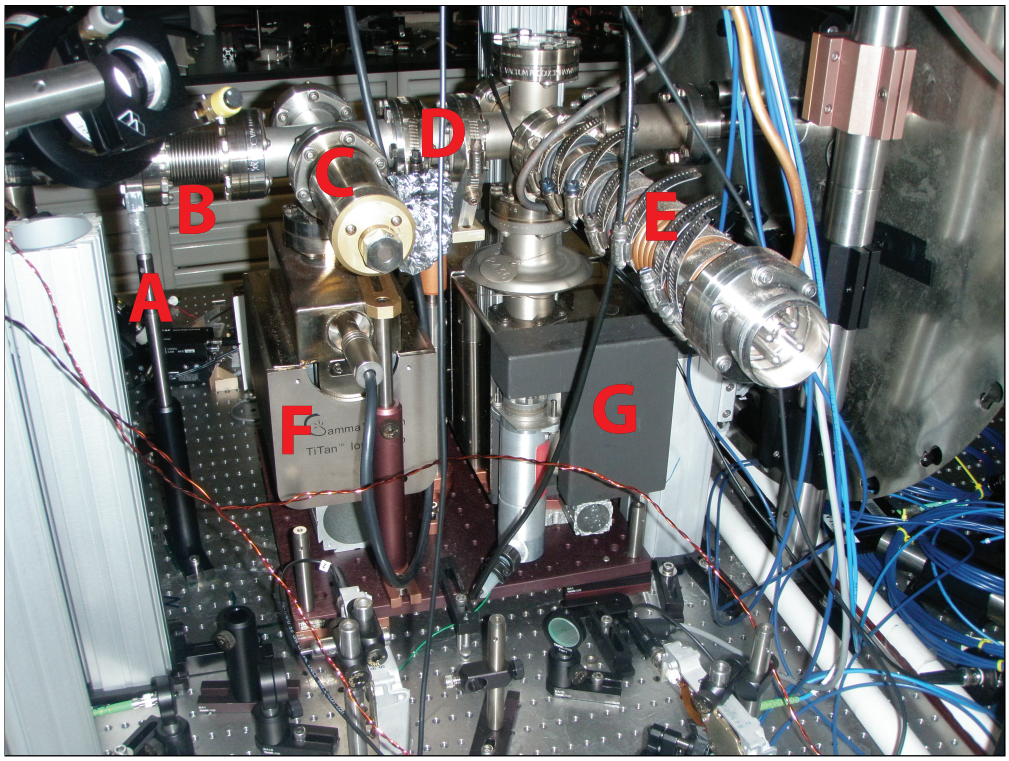

Figure 2.2: This is a labeled photo of the ion pump section of the vacuum chamber. The glass VC MOT cell, hidden from view in the upper left hand corner of the photo, is supported by a V-clamp (A) at the point where it connects to a steel bellows (B). The only valve (C) in the entire vacuum chamber is used during the bakeout process to connect roughing and turbo pumps. The differential pumping tube is located at point (D) in between the two ion pumps. The titanium sublimation pump electrodes (E) are only used when we believe we need to reduce the vacuum pressure due to increased background gas load over time. The Gamma ion pump (F) pumps on the VC MOT side of the vacuum chamber, while the Varian ion pump (G) pumps on the UHV MOT side (which extends out of the picture to the right).

cleaned in this way, they are wrapped in special UHV foil to prevent excess contamination from dust and other particles floating around in the air until the chamber is ready for assembly.

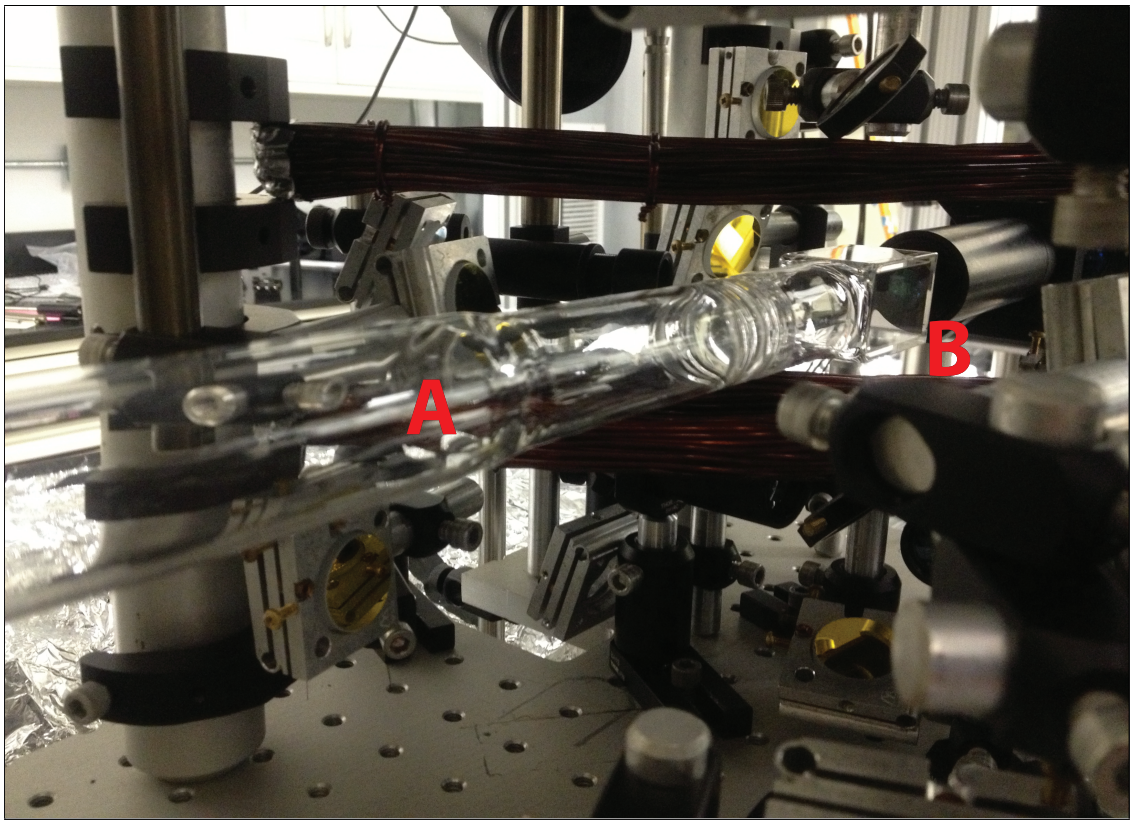

Figure 2.3: This is a labeled photo of the UHV MOT vacuum chamber cell. An 18 in long glass tube (A) extends along the line of sight of the pusher and probe beams window of the VC MOT cell. The 1 in $\times$ 1 in $\times$ 1.5 in rectangular glass UHV cell (B) is attached to the end of the tube and is located at the center of the magnetic shield.

When connecting flanges, each copper gasket should be removed from its wrapper and thoroughly wiped down with methanol just before inserting it into the flange housing. Before inserting the bolts into the flange assembly, each bolt should be lightly coated with an anti-seizing compound to prevent seizing of the nuts and bolts during the bakeout process. Each bolt should be tightened in small, incremental steps, going around the flange in a star pattern until a 0.005 in spacer cannot fit between the gap in the flange joints. This helps to create a uniform, airtight seal of the copper gasket located inside the flange.

Several options are available for baking out the vacuum chamber. One common method involves wrapping the entire chamber in silicon heater tape. This has the benefit of allowing

an in situ bakeout process with minimal displacement of the the external experimental apparatus. In practice, however, wrapping the heater tape to get sufficient thermal contact on all vacuum chamber surfaces often still requires the removal of much of the external experimental apparatus due to space constraints. In addition, putting the heater tape in contact with the glass surfaces may introduce scratches on the glass. Any surfaces not in uniform thermal contact with the heater tape may also introduce thermal gradients during the bakeout process which may cause regions of the chamber to expand at different rates, especially at glass to metal seals, which can cause failures.

Another option is to construct an oven in situ around the vacuum chamber. This, however, requires the removal of most of the external experimental apparatus because the optics cannot withstand the high temperatures of a bakeout. Unlike the heater tape method, though, the glass components of the vacuum chamber will not get scratched, and far greater temperature uniformity can be achieved. Constructing an oven in situ can be a very tedious process, too, because space constraints often require oddly shaped, inefficient oven designs.

The method that was ultimately used for this experiment involves constructing an oven away from the vacuum chamber, and then removing the vacuum chamber from the optics table and placing it in the oven in a different location. A picture of the full vacuum chamber removed from the experimental setup in preparation for a bakeout is shown in Figure 2.4, and Figure 2.5 shows the oven in use during a bakeout. This still requires the removal of some of the external experimental apparatus, but only a minimal amount to safely slide the vacuum chamber out of place. The oven is constructed of 3 in thick fiberglass insulating panels on all sides wrapped in about five layers of aluminum foil. Each wall and lid of the oven can be removed independently to allow easy transfer of the vacuum chamber into and out of the oven. Holes are cut into some of the walls to allow the parts of the vacuum chamber that might be damaged by the high temperatures of the bakeout to be exposed to room temperature air. These parts include the turbo pump valve, the titanium sublimation pump electrodes, and the ion pump high voltage connections. Heating inside the oven is provided

by one or more ceramic heaters, shielded from direct heating of the vacuum chamber by a thick shield of aluminum foil. Thermocouples placed in free space inside the oven and in contact with various parts of the vacuum chamber show that the temperature is fairly uniform throughout the volume, and temperature gradients are well within acceptable limits.

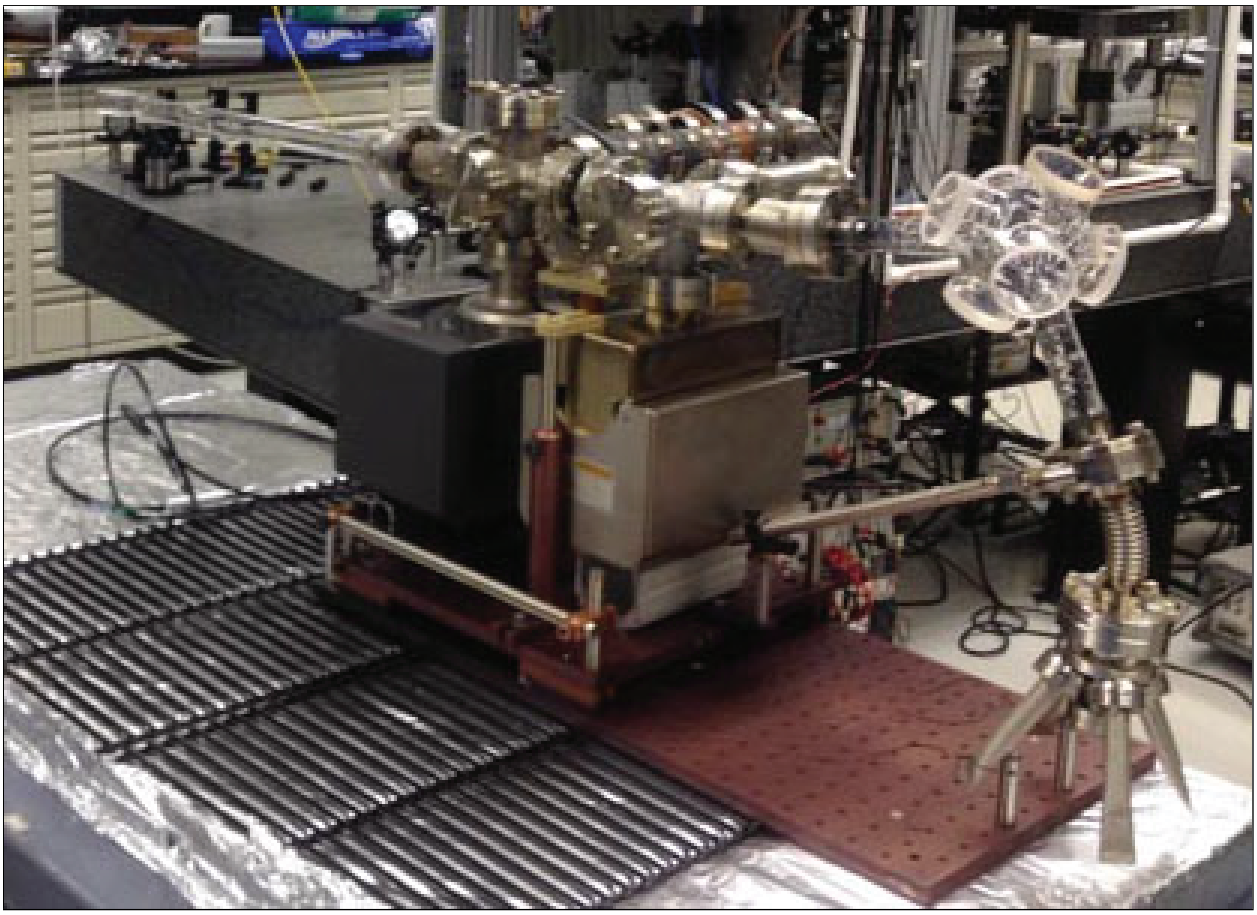

Figure 2.4: This is a photo of the full vacuum chamber removed from the experimental setup in preparation for a bakeout.

A vacuum chamber bakeout is necessary to achieve the lowest possible base pressures. At very low pressures, many substances outgas and increase the vacuum chamber pressure. Baking out the vacuum chamber removes these substances and is necessary to achieve the lowest possible base pressures. The bakeout process begins by turning on a roughing pump for about ten minutes to get the vacuum chamber down to a pressure where the turbo pump

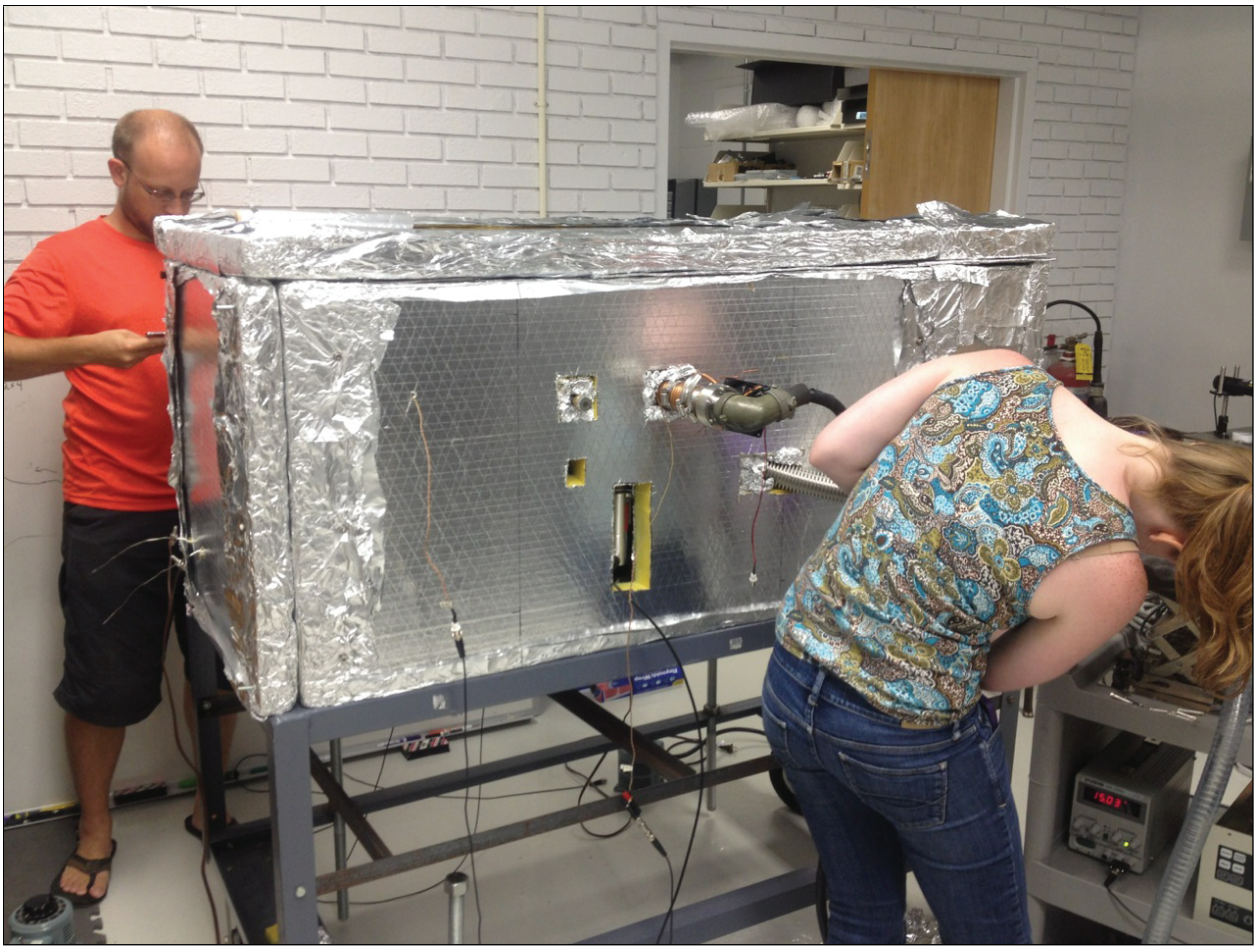

Figure 2.5: This is a photo of the oven in use during a bakeout. Holes cut into the walls can be seen that allow parts of the vacuum chamber apparatus sensitive to high temperatures to be exposed to room temperature air.

can take over. Once the turbo pump is turned on, it is allowed to pump at room temperature for at least four hours. After the turbo pump has had a chance to pump the system at room temperature for several hours, the heaters are turned on in the oven and monitored to increase the temperature at a rate not to exceed 1° C/min. At this point, one of the filaments of the titanium sublimation pump is also turned on with a continuous 25 A supply of current. Once the vacuum chamber has reached 250° C, pumping continues at this temperature for about 12 hours. After pumping for 12 hours at 250° C, the turbo pump should not be able to reduce the vacuum chamber pressure any more on its own. The ion pumps are turned

on, and if any leaks are present in the system, the pressure will be too high and they will turn off. If no leaks are present, however, the titanium sublimation pump may be degassed. For each filament, the current is increased to 40 A for one minute and then reduced back to 0 A. After completion of the degassing process, one of the electrodes is returned to 25 A while the heaters are still on. Once the ion pumps are turned on, their pressure is monitored while maintaining a temperature of 250° C. Their pressure should decrease and reach an asymptote, at which point the pressure in the system will not decrease any further at 250° C. At this point, it is safe to begin the cooling process, and the temperature again must be monitored to not exceed a rate of decrease of 1° C/min. The pressure in the vacuum system should continue to drop as the temperature drops. The turbo pump valve may be hand tightened if it does not show an increase on the ion pump pressure gauges. Once back at room temperature, the vacuum chamber is once again checked for leaks, and if none are present, the turbo pump valve is fully closed with a wrench. Using this procedure, a pressure below the threshold of the Varian ion pump gauge was achieved on the UHV MOT side of the vacuum chamber. It is estimated that this is less than 10^{-11} torr. The Gamma ion pump showed an estimated pressure of 10^{-10} torr on the VC MOT side of the vacuum chamber. An example of a bakeout procedure can be seen in Figures 2.6, 2.7, and 2.8, showing the temperature at various points in the oven and on the vacuum chamber, the pressure at each ion pump as determined by the pump current, and the current between the ion pump element plates, respectively.

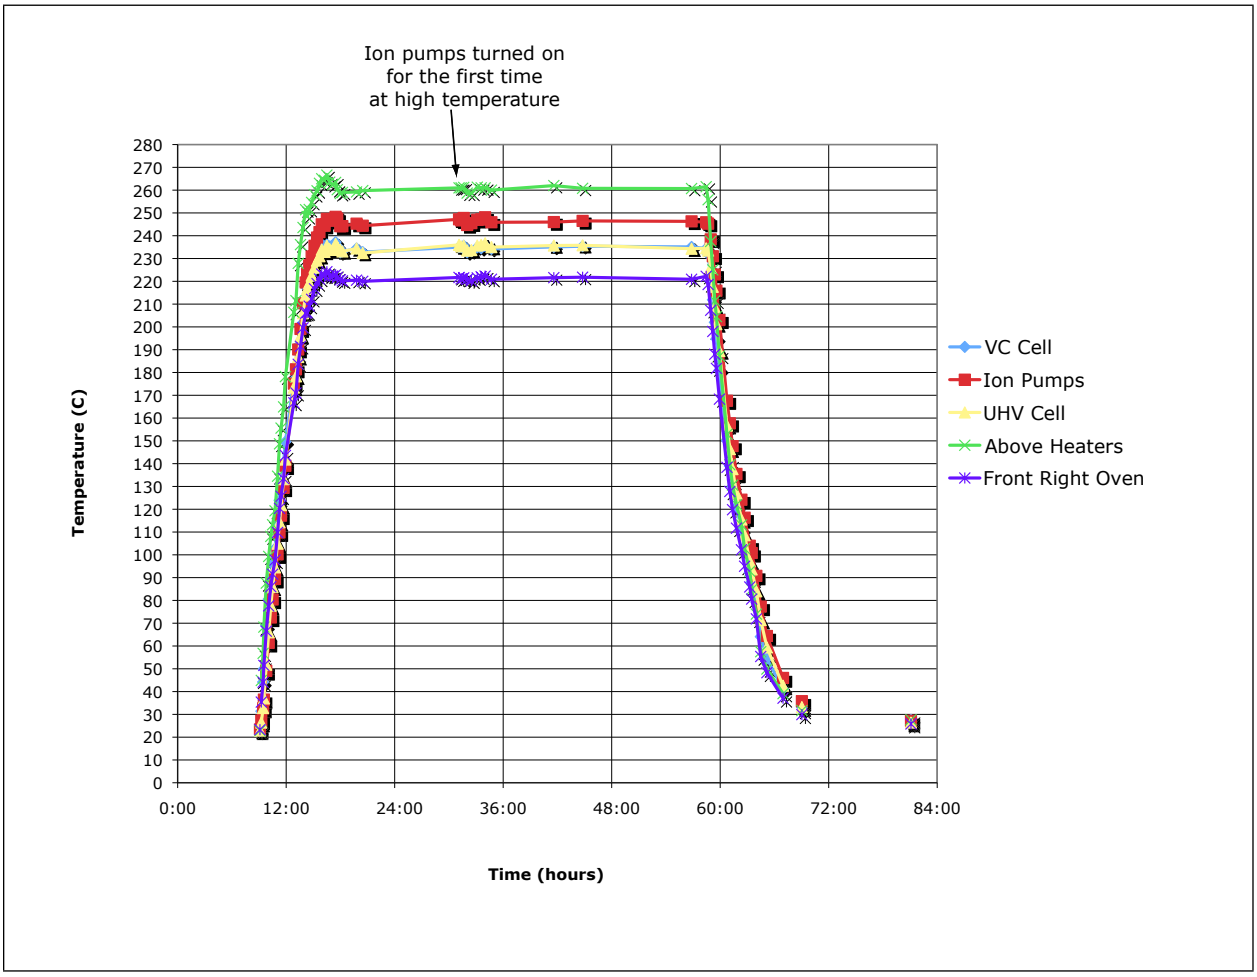

Figure 2.6: This plot shows the temperature at various points inside the oven during a bakeout of the vacuum chamber. Thermistors were placed in contact with the glass VC MOT cell, a flange connector between the ion pumps, and the glass UHV MOT cell. Thermistors were also placed in free space in the oven directly above the heaters and in a region in the oven far away from the heaters. The heating and cooling rates are $\leq 1\,^{\circ}\mathrm{C/min}$.

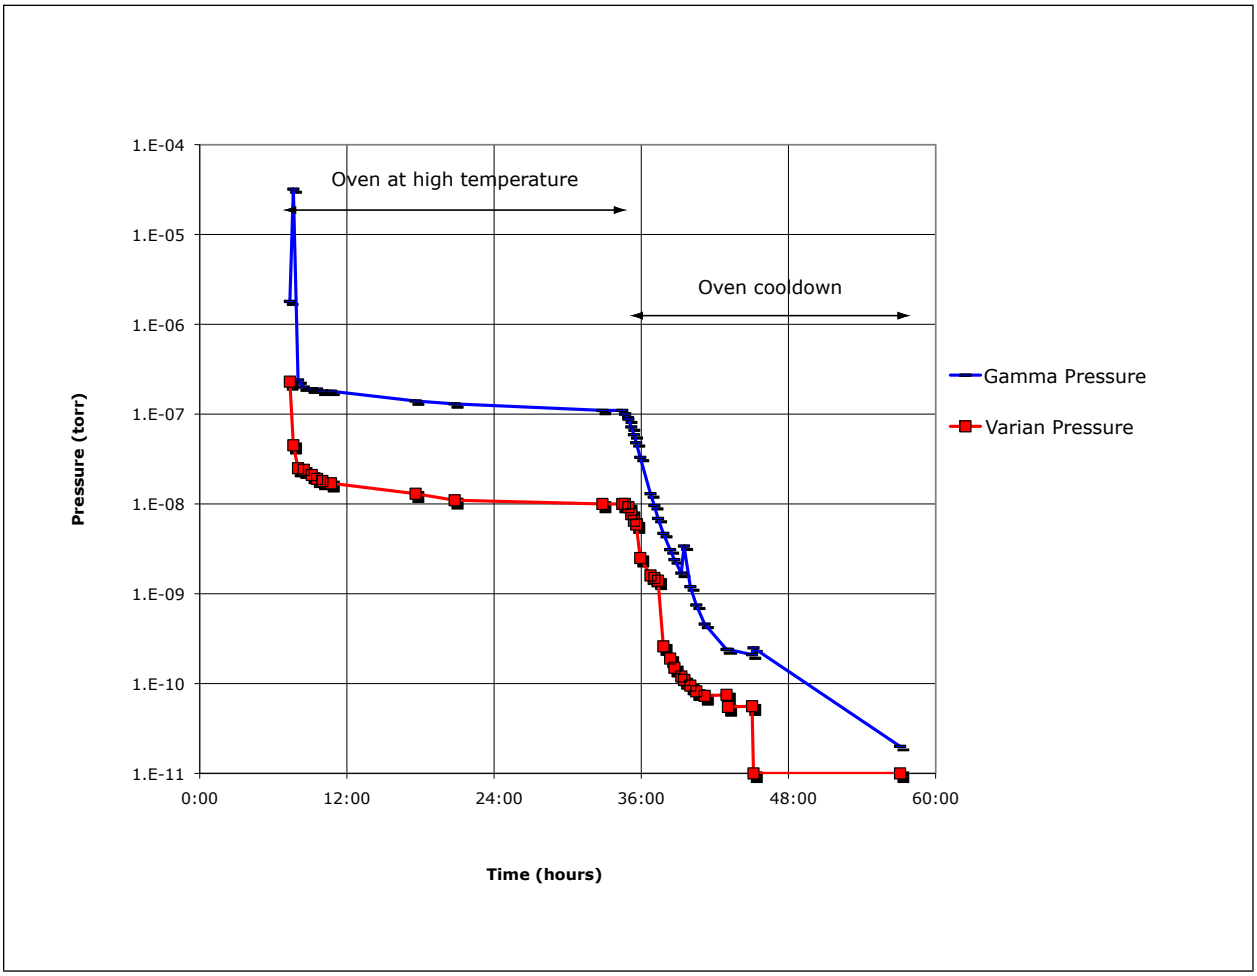

Figure 2.7: This plot shows the pressure at each ion pump as determined by the pump current during a bakeout of the vacuum chamber. The ion pumps do not directly measure the pressure in the vacuum chamber. Instead, the pump controller is calibrated to calculate the pressure based on the ion pump current according to manufacturer specifications. Therefore, pressures determined by the ion pumps can only be used as estimates of the pressure in the vacuum chamber in the immediate vicinity of the ion pumps. The final two data points for the Varian pressure indicate that the pump current is below the detection threshold of the ion pump controller and the pressure is below 10^{-11} torr.

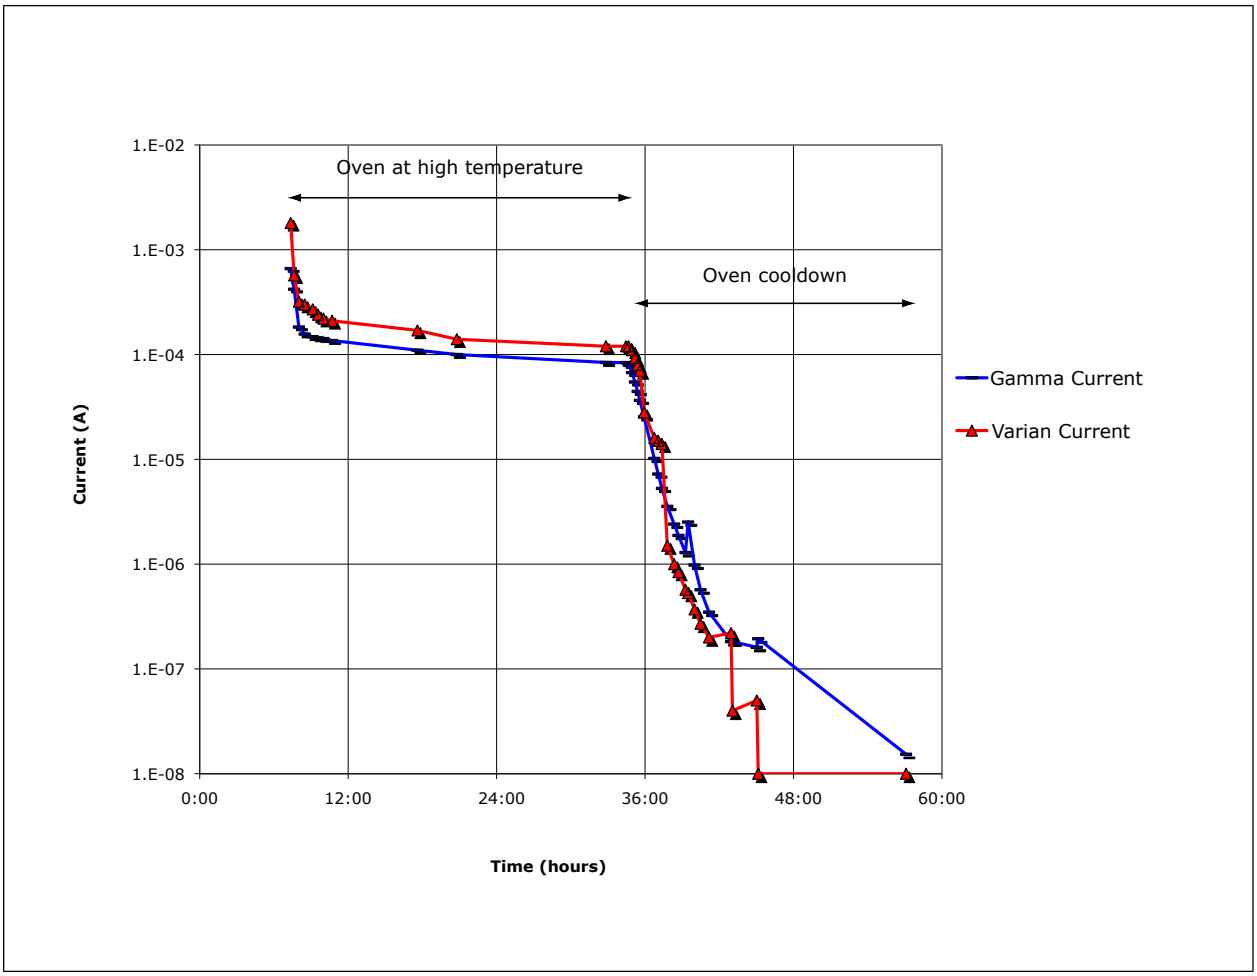

Figure 2.8: This plot shows the current between the ion pump element plates for each pump during a bakeout of the vacuum chamber. The final two data points for the Varian pump current are below the detection threshold of the ion pump controller, which is below 10 nA.

2.2 Laser frequency locking

Precise knowledge and control of laser frequencies relative to the atomic transitions of ^{87}Rb are required for the creation of a MOT. Because MOT light comes from multiple laser sources, this necessitates several methods for laser frequency locking and stabilization. The "master" and "repumper" lasers are locked directly to atomic transitions of ^{87}Rb using saturated absorption locks. The "MOT" laser uses a frequency offset lock described in Chapter 3 to lock a slave laser to an offset frequency relative to the frequency of the master laser.

A single beam passing through a reference vapor cell at room temperature and falling onto a photodiode will exhibit Doppler-broadened absorption peaks corresponding to the frequency separation of the ground and excited states. A laser frequency can be locked to this absorption peak, but it provides insufficient knowledge and control of the laser frequency. It would be better to resolve the hyperfine structure of these transitions and lock the laser frequency to one of these sharper peaks. To fix this, a dual pump-probe beam technique is used called saturated absorption spectroscopy.

The master laser is a Toptica DL Pro grating stabilized tunable single mode diode laser. The current and temperature of the diode are set to generate 780.24 nm light in the center of its modehop-free tuning range. A small fraction of the total output power of the laser is diverted into the two beams used for the saturated absorption lock. The relatively strong pump beam does a double pass through an AOM tuned for 80 MHz, frequency modulated with a small 70 kHz dither signal. The resultant beam, which is shifted 160 MHz to the red, then passes through a Rb vapor cell. The relatively weak probe beam passes through the Rb vapor cell in the opposite direction before falling onto a photodiode. The two beams cross inside the vapor cell to form a large intersection region but do not perfectly copropagate, as seen in Figure 2.9.

For atoms in resonance with the strong pump beam, the populations of the excited states are saturated. As a result, the weak probe beam experiences a Doppler-broadened absorption

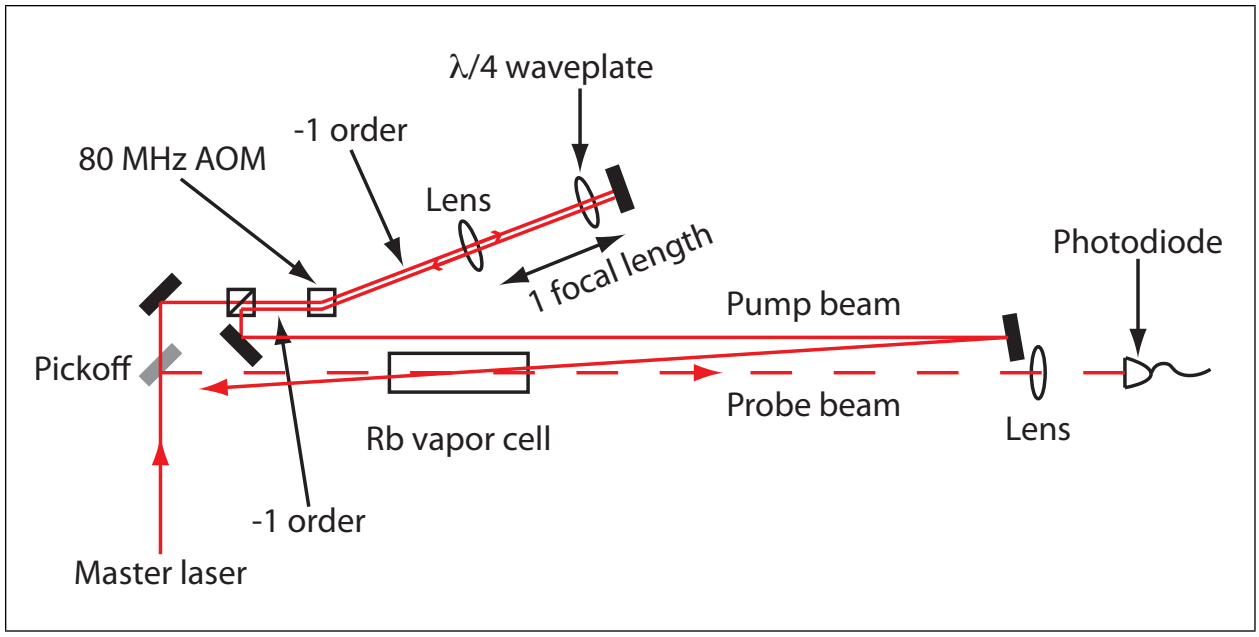

Figure 2.9: This is a diagram of the master laser saturated absorption spectroscopy setup.

peak with sharp "holes" in the absorption signal. These Doppler-free holes correspond to resonances between various hyperfine transitions in Rb, as well as crossover resonances between hyperfine sublevels. They arise because the two beams interact with atoms that have been Doppler shifted in opposite directions. Only those atoms that have a doppler shift of -80 MHz with respect to the bluer probe beam and $+80$ MHz with respect to the redder pump beam along their intersection axis will be resonant with both beams at the same time. Probe photons resonant with atoms that have been saturated by the pump beam will be less absorbed because the population is already bleached, resulting in an increase in transmitted light at that frequency. Because the pump beam is 160 MHz to the red of the probe beam, the saturated absorption spectrum will actually be red-detuned by 80 MHz from the $5^2S_{1/2,F=2} \rightarrow 5^2P_{3/2}$ hyperfine transitions, halfway between the frequencies of the two beams. The dither on the pump beam creates a derivative signal of the saturated absorption spectrum with a steep zero crossing that can be used for an error signal input into a Proportional-Integral-Derivative (PID) servo controller. Any deviation from the lock point

of the PID servo controller will be fed back to the laser current to correct for any drifts in laser frequency. Laser cooling requires frequency stabilities below a natural linewidth, Γ ($\Gamma = 5.75$ MHz for the $5^2S_{1/2} \rightarrow 5^2P_{1/2}$ transition of ^{87}Rb and $\Gamma = 6.07$ MHz for the $5^2S_{1/2} \rightarrow 5^2P_{1/2}$ transition). As our PID controlled saturated absorption spectroscopy locks exhibit frequency stabilities below 1 MHz, they demonstrate suitable stability for laser cooling.

The repumper laser is a Photodigm laser diode that has been injection locked by a grating stabilized single mode diode laser. The grating stabilized laser is frequency stabilized using a saturation absorption lock with a similar setup as the master laser. However, instead of the pump beam doing a double pass through an 80 MHz AOM before passing through the vapor cell, it only does a single pass, and the AOM is oriented so that its frequency is shifted in the opposite direction. As a result, the saturated absorption spectrum for the repumper laser shows the $5^2S_{1/2,F=1} \rightarrow 5^2P_{3/2}$ hyperfine transitions, blue-detuned by 40 MHz. A diagram of the repumper laser saturated absorption spectroscopy setup can be seen in Figure 2.10.

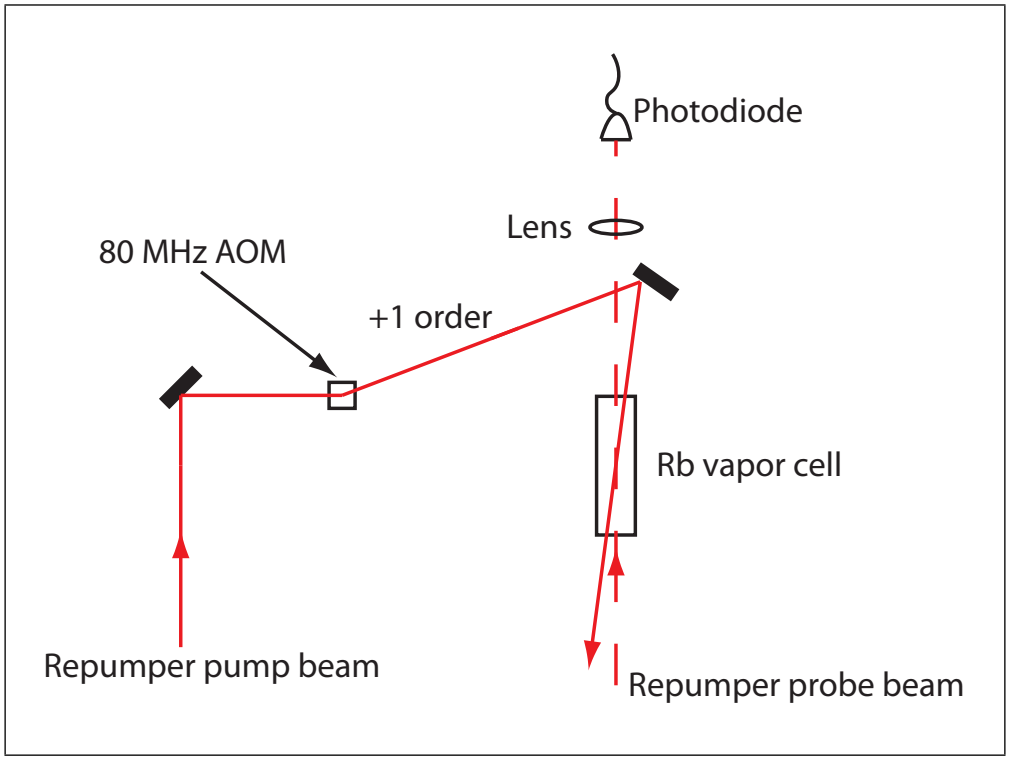

Figure 2.10: This is a diagram of the repumper laser saturated absorption spectroscopy setup.

Figure 2.11 shows the energy level diagram for the master laser and MOT laser frequency locks. Figure 2.12 shows a typical saturated absorption spectrum for the master laser. The frequency of the master laser is locked 80 MHz to the red of the $F = 2 \rightarrow F' = 2 \times 3$ crossover resonance. The MOT laser is frequency offset locked 120 MHz to the blue of the master laser frequency and subsequently upshifted 80 MHz by an AOM before entering the MOT chambers. The repumper laser is locked 40 MHz to the blue of the $F = 1 \rightarrow F' = 0 \times 2$ crossover resonance and then upshifted 80 MHz by an AOM before entering the MOT chambers, which can be seen in Figure 2.13. Figure 2.14 shows a typical saturated absorption spectrum for the repumper laser.

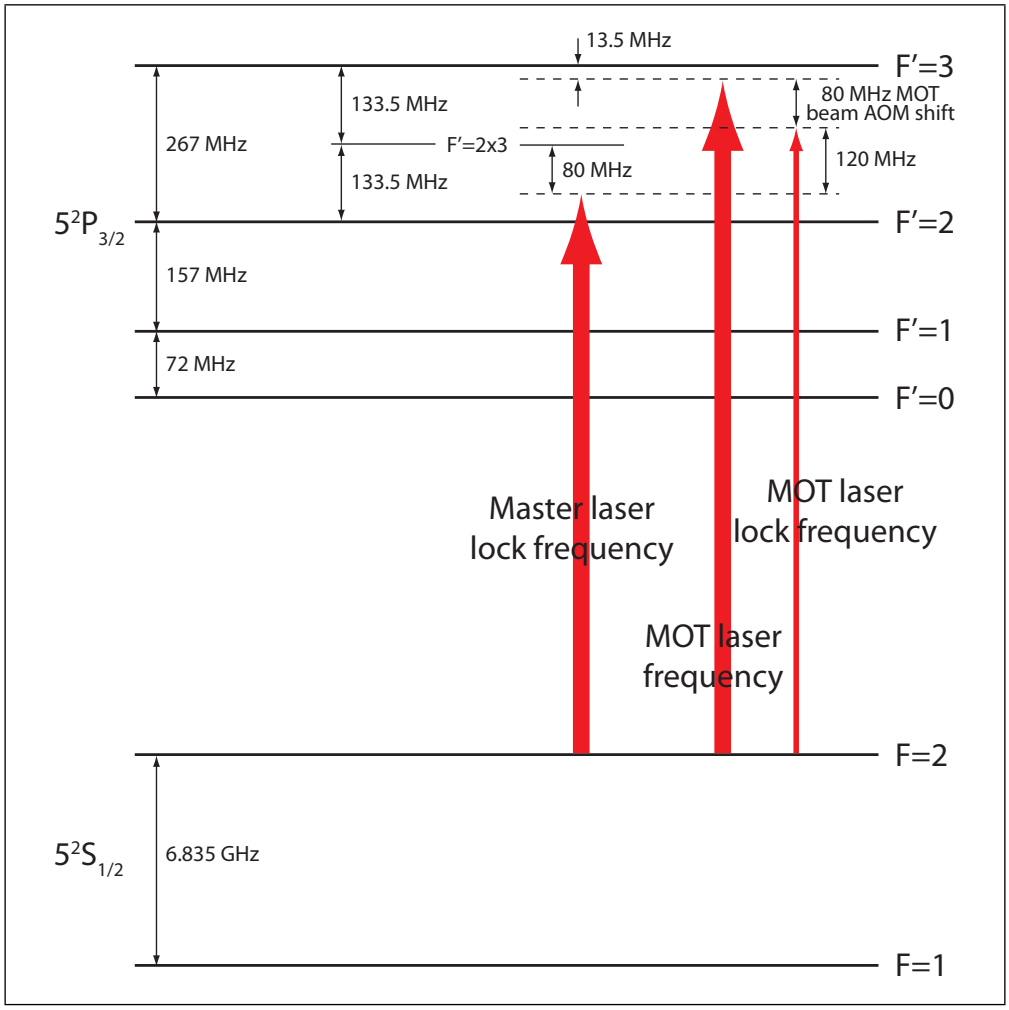

Figure 2.11: The energy level diagram for the ^{87}Rb D_2 line at 780 nm, showing the frequencies of the master and MOT laser locks. The master laser is locked 80 MHz to the red of the $F = 2 \rightarrow F' = 2 \times 3$ crossover resonance using a saturated absorption lock. The MOT laser is locked 120 MHz to the blue of the master laser frequency using a frequency offset lock. The MOT beam frequency is then upshifted by 80 MHz with an AOM. The detuning of the MOT beams during MOT loading is 13.5 MHz to the red of the $F = 2 \rightarrow F' = 3$ transition.

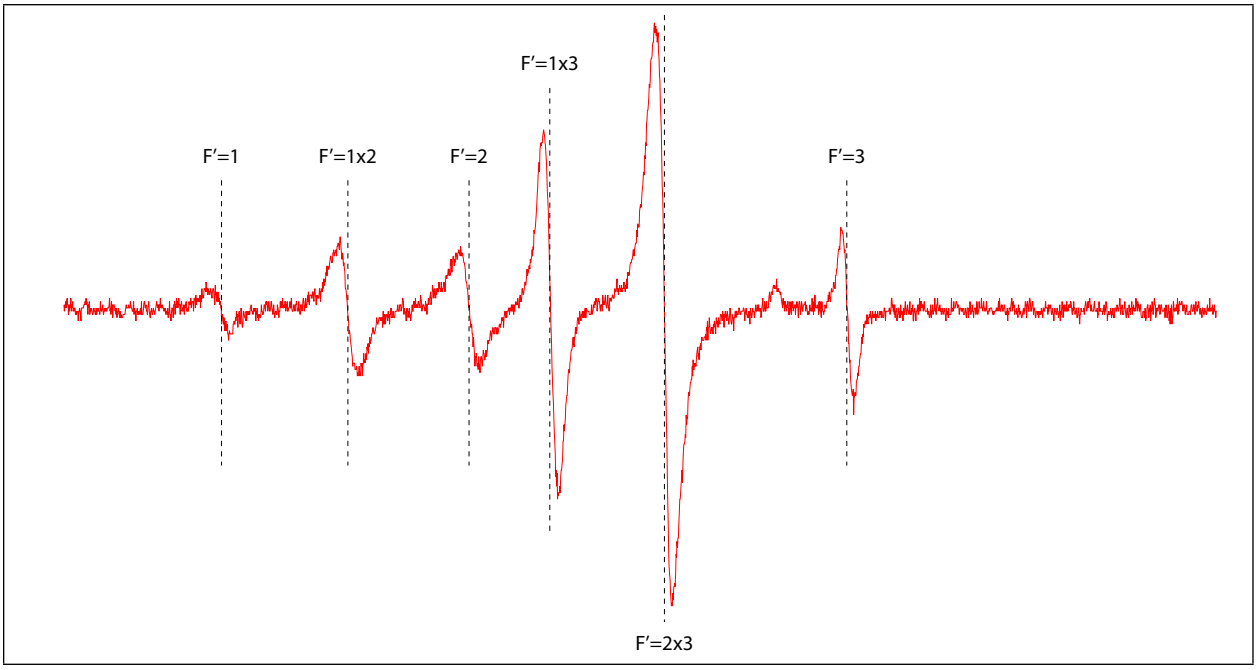

Figure 2.12: Saturated absorption spectrum for the master laser. This shows the transitions in the $F = 2 \rightarrow F'$ manifold. The master laser is locked to the zero crossing of the $F = 2 \rightarrow F' = 2 \times 3$ crossover, shifted to the red by 80 MHz.

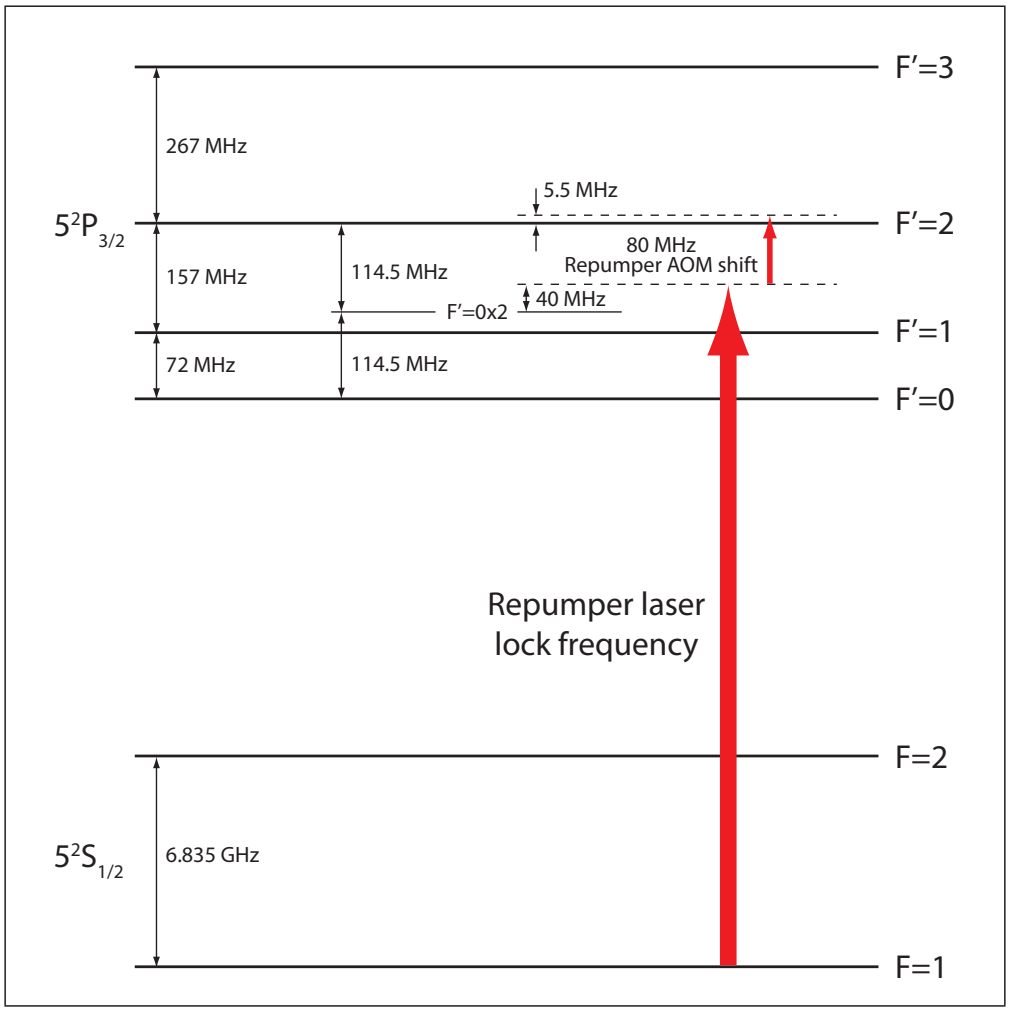

Figure 2.13: The frequency level diagram for the repumper laser lock. The repumper laser is locked 40 MHz to the red of the $F = 1 \rightarrow F' = 0 \times 2$ crossover resonance using a saturated absorption lock. The repumper beam frequency is then upshifted by 80 MHz with an AOM. The final detuning of the repumper beams is 5.5 MHz to the blue of the $F = 1 \rightarrow F' = 2$ transition.

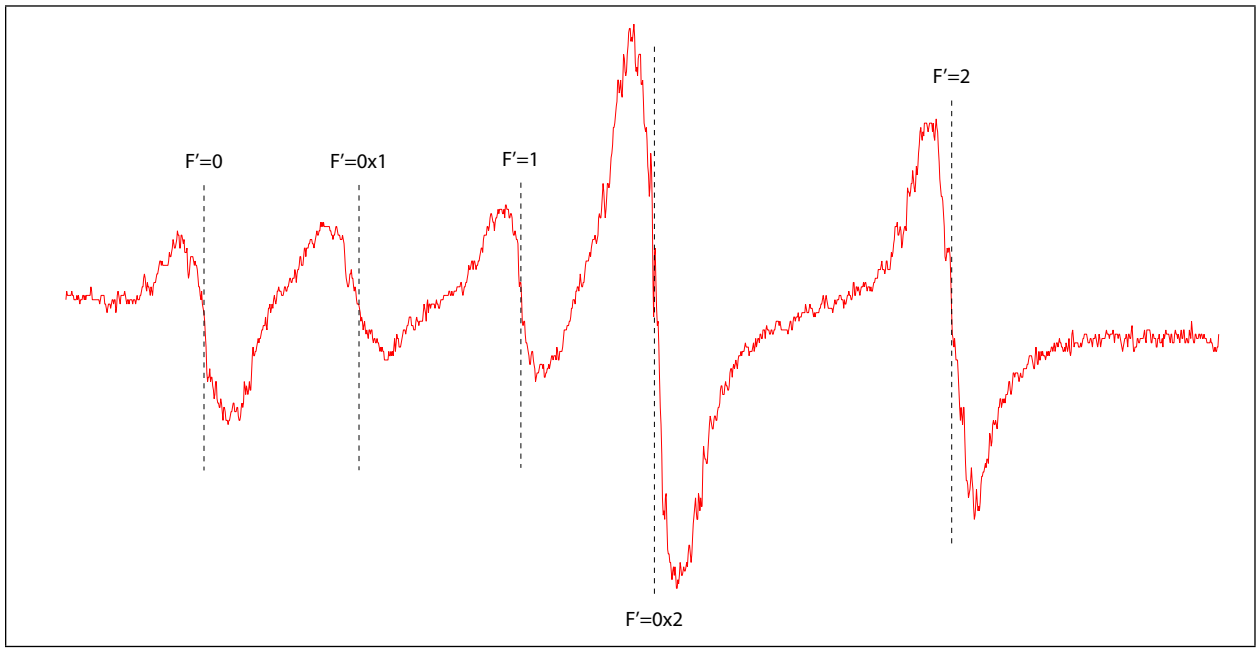

Figure 2.14: Saturated absorption spectrum for the repumper laser. This shows the transitions in the $F = 1 \rightarrow F'$ manifold. The $F = 1 \rightarrow F' = 1 \times 2$ crossover transition is not visible in this spectrum. The repumper laser is locked to the zero crossing of the $F = 1 \rightarrow F' = 0 \times 2$ crossover, shifted to the blue by 40 MHz.

2.3 VC MOT

Overview of the apparatus

A MOT in a UHV environment cannot be loaded directly from a background vapor and instead requires a cold atom beam source of rubidium. One common mechanism for loading a UHV MOT with a slow atomic beam is the Zeeman slower [38]. However, due to the large residual magnetic fields this apparatus can produce, it is an unattractive candidate for an experiment that relies on ultralow magnetic fields. To generate a cold atom beam source for our UHV MOT, a VC MOT loaded from a background vapor of room temperature rubidium precools the atoms before transfer into the UHV MOT. Efficient loading of a VC MOT requires a high rubidium vapor pressure to provide a source of atoms for the MOT. The rubidium source is a glass ampoule of solid rubidium located below the vapor cell chamber. Upon completion of the vacuum chamber bakeout, the steel finger holding the ampoule is squeezed with a pair of pliers until the glass cracks and the rubidium is exposed to the chamber. At room temperature, this will generate a vapor pressure of $\approx 1 \times 10^{-8}$ torr. If a higher vapor pressure is desired, the steel finger containing the rubidium can be heated slightly.

The VC MOT consists of three orthogonal retroreflected beams to form a six beam 3D trap. A pair of electromagnetic coils lies along the axis of one pair of retroflected beams in an anti-Helmholtz configuration as in Figure 2.15. These coils provide a magnetic quadrupole field in the region between them and are oriented so the MOT beams cross where the magnetic field strength is zero. The MOT beams are all circularly polarized, and upon retroreflection, each beam passes twice through a $\lambda/4$-waveplate to reverse its helicity.

A MOT utilizes a twofold approach to cool neutral atoms. First, a laser cooling process called optical molasses slows the atoms in the MOT via Doppler cooling, although near the zero field point in the center of the trap, additional sub-Doppler cooling mechanisms also exist [39]. To facilitate optical molasses, MOT beams are detuned slightly to the red

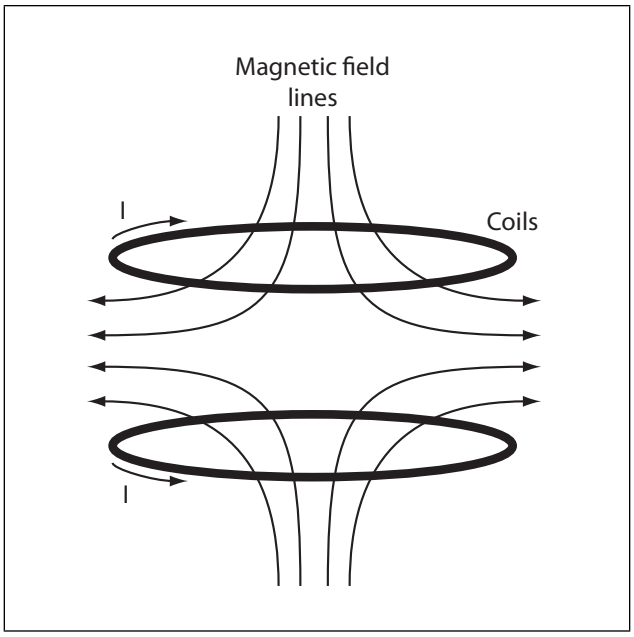

Figure 2.15: The MOT coils are arranged in an anti-Helmholtz configuration, each having a current I running in opposite directions. This produces a spherical quadrupole magnetic field in the region between the coils. The spherical quadrupolar field also has lines coming out of the page.

of the $F = 2 \rightarrow F' = 3$ cooling transition (typically -13.5 MHz during loading in our experiment). Atoms with a nonzero velocity along a beam path will therefore see a Doppler shifted frequency in the light field dependent on the direction of propagation of the beams. Atoms moving away from the center of the trap will absorb many more photons from the counterpropagating beam than from the beam traveling in the same direction as the atoms. These atoms will then spontaneously emit photons in a random direction, which slows the atoms down on average, cooling the atoms. Atoms at rest between two beams will absorb the same number of photons from each beam, which causes them to experience equal forces in each direction. By including three pairs of orthogonal beams, atoms can be cooled and slowed in every direction. In the absence of the MOT's magnetic gradient, optical molasses actually produces atom clouds with much lower temperatures than those achieved in the MOT due to sub-Doppler cooling processes (i.e. polarization gradient cooling) [39]. The

presence of the magnetic field spoils these processes and reduces their effects, limiting the minimum temperature achievable in a MOT.

Despite MOT beam frequencies being near the $F = 2 \rightarrow F' = 3$ resonant transition, there is a small probability an atom can be excited into the $F' = 2$ state instead of $F' = 3$. When an atom in the $F' = 2$ excited state spontaneously emits a photon, it has a 50% chance to decay into the $F = 1$ ground state instead of the $F = 2$ ground state [40]. This is called a "dark" state because the atom can be excited to the cooling transition with MOT beam frequency photons with negligible probability. Atoms in a dark state will no longer be confined in the trap because they will no longer scatter photons on the cycling transition, and thus the spontaneous scattering force vanishes. By including a small amount of repumper light in the MOT beams, the few atoms that fall into the $F = 1$ ground state will be excited to the $F' = 2$ excited state. These atoms will then have the opportunity to spontaneously emit a photon and decay back to the $F = 2$ ground state where they can once again absorb MOT beam photons.

While optical molasses cools the atoms and operates in velocity space, an overlapping magnetic quadrupole field confines the atoms and operates in position space. The field-producing coils are arranged so the magnetic field strength increases linearly along the direction of each pair of beams, and the field strength is zero at the point where all three beams cross. Consider a simplified atom with total angular momentum of the ground state $J_g = 0$ and total angular momentum of the excited state $J_e = 1$. In the presence of a magnetic field, this transition splits into three Zeeman components, with each component excited by a different beam polarization. The $m_J = \pm 1$ components are excited by circularly polarized $\sigma^{\pm}$ light, and the $m_J = 0$ component is excited by linearly polarized π light. Due to the Zeeman effect, the frequency of each component varies with the magnetic field strength, and therefore the position in the trap, as seen in Figure 2.16. On one side of the trap, the magnetic field will shift the $m_J = +1$ state up in frequency and the $m_J = -1$ state down by the same frequency. On this side of the trap, because the MOT beams are detuned below the

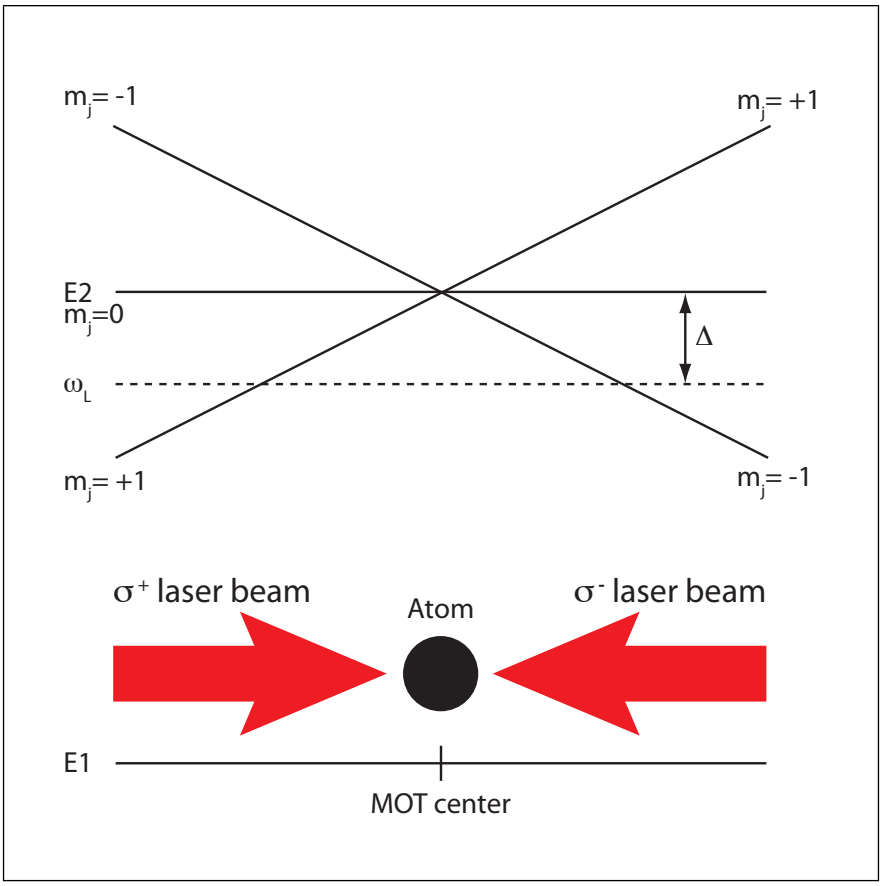

Figure 2.16: In the presence of a magnetic field of linearly varying strength, the excited atomic state is split into three sublevels with $m_J = 0, \pm 1$. A laser field with frequency ω_L detuned by Δ to the red of the zero field resonance selectively interacts with the $m_J = \pm 1$ sublevels depending on its polarization and the position of the atom in the trap.

zero field atomic resonance, they should be arranged so that the beam with σ^- polarization points to the center of the trap. In this way, more light will be scattered from the σ^- beam than the σ^+ beam, pushing the atoms to the center of the trap. On the other side, the $m_J = \pm 1$ states are reversed, and the MOT beam with σ^+ polarization points to the center of the trap. This results in a net force on the atoms, confining them in all directions at the point in space where the magnetic field is zero. While this description is for a simplified atomic transition of $J_g = 0 \rightarrow J_e = 1$, it works the same for the more complicated level structure of ^{87}Rb, which uses the transition $J_g = 2 \rightarrow J_e = 3$.

Each VC MOT beam is 200 mW clipped to a 1.5 in diameter when it enters the VC MOT chamber. Due to small beam power losses at each optics interface, the beams are designed to converge very slightly so each retroreflected beam matches the intensity of the original beam at the point where all six cross. A total of 10 mW of repumper light is combined with the MOT light in a fiber optic cable before entering the VC MOT chamber. The MOT light is detuned 13.5 MHz to the red of the $F = 2 \rightarrow F' = 3$ transition as shown in Figure 2.11. This value was chosen experimentally to load the largest number of atoms possible into the UHV MOT and does not necessarily yield the VC MOT with the largest number of atoms. The electromagnetic coils, which have a 2 in diameter and are spaced 4 in apart, are 20 turns each with 12 A of current, which generates of field gradient of 7 G/cm.

VC MOT characterization

The number of atoms in the VC MOT is determined using a simple laser stimulated fluorescence technique. A photodiode is set up to collect photons emitted by the fluorescing MOT, and the resulting voltage that is displayed on an oscilloscope is used to calculate the number of atoms in the trap. First, a photodiode calibration parameter is measured by shining a beam with a known amount of power onto the photodiode and measuring the voltage it produces. Because the MOT emits fluorescence photons in every direction, the photodiode only sees a fraction of the total emitted light. This fraction is determined by the solid angle seen by the photodiode. If r_{pd} is the radius of the photodiode, and d is the distance from the photodiode to the MOT, this fraction is

$$A = \frac{\pi r_{pd}^2}{4\pi d^2}. \tag{2.1}$$

Therefore, if F_0 is the fluorescence seen by the photodiode, the total fluorescence of the MOT, F, given in photons/s is

$$F = \frac{F_0}{A}. \tag{2.2}$$

Fluorescence is the spontaneous emission of photons from atoms in the excited state, so to calculate the total number of atoms in the MOT, the fraction of the atomic population in the excited state, ρ_{ee}, must be found using [39]

$$\rho_{ee} = \frac{s_0/2}{1 + s_0 + (2\Delta/\Gamma)^2}. \tag{2.3}$$

Here, Δ is the MOT beam detuning in MHz, $\Gamma = 6.1$MHz is the natural line width of the transition, and $s_0 = I/I_s$ is the on resonance saturation parameter, where I is the measured intensity of all the beams at the center of the MOT, and $I_s = 1.6$ mW/cm^2 is the saturation intensity for $\sigma^{\pm}$-polarized light for the $F = 2 \rightarrow F' = 3$ cycling transition [41]. Finally, the number of atoms in the MOT, N, is given by

$$N = F\frac{\tau}{\rho_{ee}}, \tag{2.4}$$

where $\tau = 26.24 \times 10^{-9}$s is the excited state natural lifetime. Using the photodiode voltage in Figure 2.17 as a measure of the fluorescence, the number of atoms in the VC MOT calculated this way is $N = 5 \times 10^{10}$ atoms.

The loading time of the VC MOT depends on the background vapor pressure in the VC MOT chamber [42]. We do not directly measure the background vapor pressure, but a rough estimate based on the temperature of the rubidium source gives a value of $\approx 3 \times 10^{-7}$ torr. The loading time, however, can be measured directly by observing the MOT fluorescence on a photodiode. In the absence of density-dependent losses, the rate equation governing the number of atoms in the MOT is given by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = L - N\gamma, \tag{2.5}$$

where L is the loading rate of the MOT, and γ is the loss rate due to background collisions. Because the vapor pressure in the VC MOT chamber is kept two orders of magnitude higher

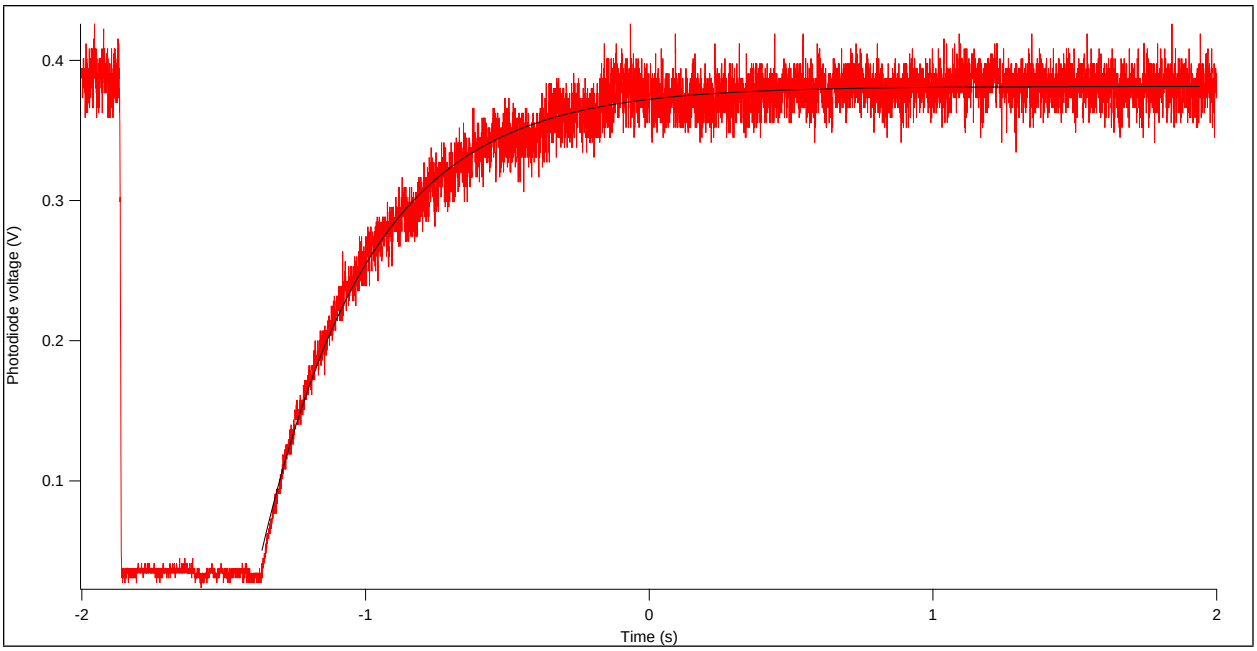

Figure 2.17: This plot shows the photodiode voltage for a typical VC MOT fluorescence loading curve. The minimum and maximum photodiode voltage levels give the background and fully loaded VC MOT fluorescence levels, respectively, which can then be used to calculate the number of atoms in the trap. The black curve is an exponential fit to the loading portion of the VC MOT fluorescence to find the loading time. For this curve, $N = 5 \times 10^{10}$ atoms and $\tau = 0.35$ s.

than the vacuum pressure, all collisional losses are assumed occur due to collisions with room temperature rubidium atoms. The number of atoms in the MOT at any point in time is therefore

$$N(t) = L\tau \left(1 - e^{-t/\tau}\right), \tag{2.6}$$

where $\tau = 1/\gamma$ is the MOT lifetime. Upon examination of Equation (2.6), it is evident that τ also represents the MOT loading time. Therefore, by measuring the $1/e$ loading time of the VC MOT, the loss rate due to background collissions for a wide range of vapor pressures can be inferred [42, 43]. Solving Equation (2.5) for the steady state, the loading rate of the VC MOT is simply

$$L_{VC} = N\gamma. \tag{2.7}$$

The loading time can be found from Figure 2.17. By measuring a loading time of $\tau = 0.35$ s, a loading rate of $L_{VC} = 1.43 \times 10^{11}$ atoms/s can be inferred.

For a VC MOT with losses determined primarily by collisions with a room temperature rubidium background vapor, the loss rate can be expressed as

$$\gamma = \frac{1.05 p \sigma}{k_B T} \sqrt{\frac{2 k_B T}{m}}, \tag{2.8}$$

where p is the background vapor pressure, $\sigma = 2.5\,\mathrm{nm}^2$ is the total collision cross section for a Van der Waals interaction, $T = 300$ K is the temperature of background gas atoms, and $m = 1.443 \times 10^{-25}$ kg is the atomic mass of rubidium [44]. Given the loss rate measured in our VC MOT, this implies a background vapor pressure of $p = 1.4 \times 10^{-8}$ torr.

2.4 Pusher/guide beam

Overview of the apparatus

Once the atoms have been precooled by the VC MOT, a mechanism is required to transfer these atoms efficiently to the UHV MOT trapping region without additional heating. Several methods are available to transfer atoms from the VC MOT via a cold atom beam into the UHV MOT that do not require the use of magnetic guiding fields. A pyramidal MOT [45, 46], conical mirror funnel [47], or low velocity intense source (LVIS) [48] scheme can provide a sufficiently slow atomic beam for transfer from a VC MOT and subsequent recapture in a UHV MOT, however they both require optics inside the vacuum chamber. Likewise, a two-dimensional MOT scheme has been shown to generate slow atomic beams sufficient for recapture in a UHV MOT but at the expense of added complexity in the electromagnetic coils configuration surrounding the VC MOT [49, 50, 51]. As a common theme, these methods all rely on radiation pressure imbalances to create an extraction column of atoms in the VC MOT pointing in the direction of the UHV MOT.

A hybrid pusher/guide beam is used in this experiment to transfer atoms from the VC MOT to the UHV MOT, similar to the setup used in [43]. A single 250 mW linearly polarized laser beam coming from a Toptica DLX 110 tunable high power single-mode diode laser is focused to a waist with $1/e^2$ radius $w_0 = 100\,\mu$m at a point 2 cm in front of the VC MOT. The beam diverges to $w_1 = 110\,\mu$m inside the VC MOT and $w_2 = 5$ mm at the UHV MOT. The beam points slightly above the location of the UHV MOT, which, along with the larger size of the beam at this spot, serves to minimize perturbations of the UHV MOT trapping mechanisms. Figure 2.18 shows a diagram of the pusher beam as it passes through the VC and UHV MOT chambers.

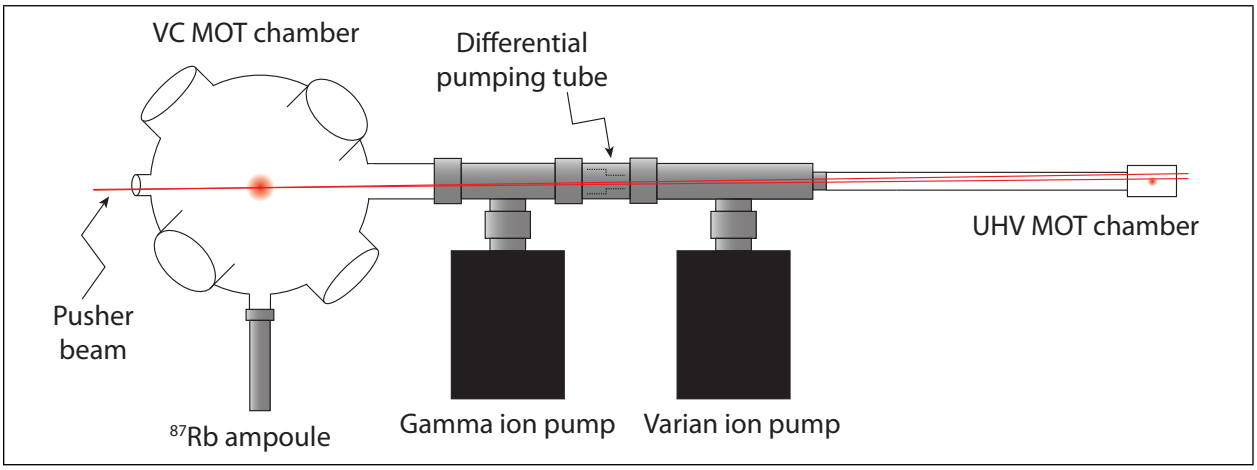

Figure 2.18: The pusher beam enters the VC MOT chamber through a 1 in diameter window on the glass cell. It focuses to a waist 2 cm before it passes through the VC MOT, passes through the differential pumping tube, and diverges to a 5 mm waist at the UHV MOT. The pusher beam is tilted slightly up in order to minimize perturbations on the UHV MOT.

At low powers and small detunings, the pusher beam creates a region of radiation pressure imbalance in the VC MOT. As atoms trapped by the VC MOT are cooled and move toward the center of the trap, they enter the extraction column created by the pusher beam and are accelerated out of the trap in the direction of the UHV MOT. As the pusher beam contains no repumper light, the atoms will only scatter pusher beam photons in the region of the VC MOT. While the atoms in the pusher beam will be heated slightly in the longitudinal direction as they are accelerated out of the VC MOT, the cooling provided by the VC MOT

limits the transverse velocities, and therefore the divergence of the atomic beam. While it has not been verified in our experiment, further transverse cooling may also occur during extraction due to an on-axis light shift of the atomic levels caused by the pusher beam.

Despite the low divergence of the atomic beam as it leaves the VC MOT, without the presence of an additional guiding potential very few atoms reach the UHV MOT trapping region. At low pusher beam powers ($< 200\,\mu$W) and small detunings ($\Delta \approx -20$ MHz from the $F = 2 \rightarrow F' = 3$ transition), atoms have been successfully loaded in the recapture UHV MOT, but only at 5% the maximum fluorescence level observed when larger powers and larger red detunings are used in the guide beam regime (see Section 2.5). Atoms in the VC MOT experience a force due to the pusher beam of

$$F = \hbar k \frac{\Gamma}{2} \frac{I/I_s}{1 + 4\Delta^2/\Gamma^2 + I/I_s}, \tag{2.9}$$

where $k = 2\pi/\lambda$ is the wave vector, $\Gamma = 2\pi \times 6.07$ MHz is the natural linewidth of the D_2 line, $I_s = 1.6\,\mathrm{mW/cm^2}$ is the saturation intensity, and Δ is the laser detuning [52, 43]. This leads to a velocity along the axis of the pusher beam of

$$v(z) = \left(\frac{\hbar k \Gamma z}{m} \frac{I/I_s}{1 + 4\Delta^2/\Gamma^2 + I/I_s} \right)^{1/2}, \tag{2.10}$$

where m is the mass of a ^{87}Rb atom [43]. An estimate of the final velocity of the atomic beam can be made for a pusher beam in the regime of low power and small detuning by calculating the velocity as it exits the VC MOT, assuming that outside the region of the VC MOT all the atoms are in the $F = 1$ ground state and do not interact with the pusher beam. In this case the velocity of the atomic beam will be 30 m/s. The distance between the VC MOT and the UHV MOT is 1 m, so it will take $t = 0.033$ s for atoms to travel to the UHV MOT. The capture velocity of the UHV MOT is

$$v_{capture} = \sqrt{\frac{r\Gamma\hbar k}{m}}, \tag{2.11}$$

where $r = 0.75$ cm is the radius of the UHV MOT beams [43]. For our UHV MOT, $v_{capture} = 30$ m/s, which is just high enough to capture atoms from the pusher beam. However, without the presence of a guiding potential outside the region of the VC MOT, the effects of gravity are significant. In $t = 0.033$ s of travel time, the atoms fall 1 cm, which could be enough to keep them from reaching the UHV MOT capture region.

At higher powers and larger detunings, the atoms may experience a confining optical dipole potential that helps them maintain their trajectory until they reach the UHV MOT. In this regime, it is beneficial to refer to the pusher beam as a guide beam in order to differentiate between the applicable situations. In the electric field of the guide beam, the atoms experience an induced dipole potential which leads to a dipole force

$$\vec{F}(\vec{r}) = -\frac{3\pi c^2}{2\omega_0^3}\frac{\Gamma}{\Delta}\nabla I(\vec{r}) \tag{2.12}$$

that scales like $1/\Delta$, where ω_0 is the transition frequency of the D_2 line. The photon scattering rate is

$$\Gamma_{sc} = \frac{3\pi c^2}{2\hbar\omega_0^3}\left(\frac{\Gamma}{\Delta}\right)^2 I(\vec{r}), \tag{2.13}$$

which scales like $1/\Delta^2$ [53]. Therefore, large detunings lead to strong confining forces with minimal scattering. Once again using Equation (2.10) to estimate the velocity of the atoms as they leave the VC MOT, a guide beam with power $P = 280$ mW and detuning $\Delta = -6$ GHz gives the same velocity $v = 30$ m/s of the atomic beam.

The optical guiding potential confines the atoms in 2D, allowing them to remain in the path of the guide beam without falling to the boundaries of the UHV cell before they reach the UHV MOT. As the guide beam diverges, the beam intensity decreases, and the depth of the dipole potential, which is well approximated by

$$U(r) = \frac{3\pi c^2}{2\omega_0^3}\frac{\Gamma}{\Delta}I(r), \tag{2.14}$$

also decreases [53]. This leads to a weaker confinement of the atomic beam, but also provides a means of adiabatic cooling in the transverse direction. As the depth of the potential decreases, atoms with the highest transverse temperatures escape over the potential boundaries, allowing the remaining atoms to rethermalize to a lower average temperature. The condition for adiabaticity is given by $|\mathrm{d}\omega_g/\mathrm{d}t| \ll \omega_g^2$, where

$$\omega_g(z) = \omega_g(0)\frac{w_0^2}{w^2(z)} = \omega_g(0)\frac{z_R^2}{z^2 + z_R^2} \tag{2.15}$$

is the transverse oscillation frequency of the guide beam [43]. To check the condition for adiabaticity, the rate of change of the transverse oscillation frequency of the guide beam is

$$\frac{\mathrm{d}\omega_g}{\mathrm{d}t} = -2\omega_g(0)\frac{\dot{z}z(z_R^2)}{(z^2 + z_R^2)^2}. \tag{2.16}$$

To compare $|\mathrm{d}\omega_g/\mathrm{d}t|$ with ω_g^2, $\omega_g(0)$ can be found using

$$\omega_g(0) = \sqrt{\frac{4U(0)}{mw_0^2}}. \tag{2.17}$$

Assuming a constant velocity for simplicity, under our experimental conditions the ratio $|\mathrm{d}\omega_g/\mathrm{d}t|/\omega_p$ is a minimum at the VC MOT and continues to increase until it reaches its maximum value of 0.026 when the atoms reach the UHV MOT. The guide beam meets the condition for adiabaticity and likely provides cooling of the atomic beam in the transverse direction.

Pusher/guide beam characterization

The optimal pusher beam power and detuning were determined experimentally by optimizing the number of atoms loaded into the UHV MOT. For larger pusher beam powers, the peak in the number of atoms in the UHV MOT occurred at larger detunings. Furthermore, the maximum number of atoms attainable grew every time the power in the pusher beams was

increased up to the maximum available power of 280 mW incident on the VC MOT. At this power, the optimal detuning is estimated to be ≈ 6 GHz to the red of the $F = 2 \rightarrow F' = 3$ transition. This detuning was estimated using a Fabry-Perot interferometer with a 1 GHz free spectral range as the laser was continuously tuned from resonance to its optimal detuning. The performance of the pusher beam is insensitive to frequency drift on the order of ± 100 MHz, so no frequency lock was used.

In the presence of the pusher beam, the number of atoms in the VC MOT, N_p, depends on an additional loss rate γ_p due to the extraction of atoms from the MOT into the pusher beam. The loading rate for the VC MOT in the presence of the pusher beam now becomes

$$L_{VC} = N_p \left(\gamma + \gamma_p\right). \tag{2.18}$$

Using the value of L_{VC} measured for the VC MOT in the absence of the pusher beam, the flux of atoms extracted out of the VC MOT into the pusher beam can be found using

$$L_{out} = \gamma_p N_p = L_{VC} - \gamma N_p. \tag{2.19}$$

N_p is calculated experimentally using the same method as described in Section 2.3.2. Figure 2.19 shows a loading curve with $N_p = 4 \times 10^{10}$ atoms. This leads to a calculated flux of $L_{out} = 2.87 \times 10^{10}$ atoms/s.

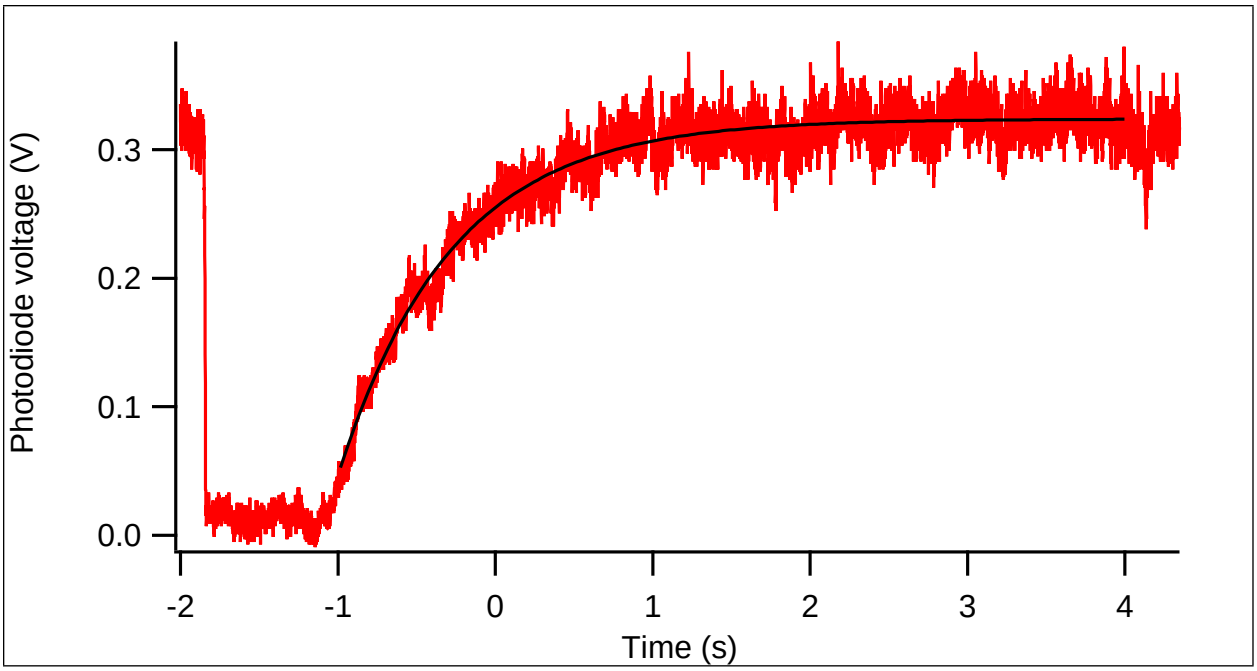

Figure 2.19: This plot shows the photodiode voltage for a typical VC MOT fluorescence loading curve in the presence of the pusher beam. The minimum and maximum photodiode voltage levels give the background and fully loaded VC MOT fluorescence levels, respectively, which can then be used to calculate the number of atoms in the trap. The black curve is an exponential fit to the loading portion of the VC MOT fluorescence to find the loading time. For this curve, $N_p = 4 \times 10^{10}$ atoms.

2.5 UHV MOT

Overview of the apparatus

The ultra high vacuum (UHV) MOT cools and traps atoms in much the same way as the VC MOT. The primary difference is that the UHV MOT exists in a region with significantly lower background vapor pressure. Whereas the source of atoms for loading into the VC MOT comes from the background vapor, the UHV MOT uses the atomic beam created by the presence of the pusher laser as its source of atoms. Because the atoms loaded into the UHV MOT have been precooled to a temperature much lower than the temperature of the background vapor, the capture velocity of the UHV MOT may be significantly lower than

the capture velocity of the VC MOT. This allows us to use smaller MOT beams. In addition, the dipole trap must be loaded in a region of ultra high vacuum with a background pressure $\sim 10^{-11}$ torr. The dipole trap must have a sufficiently long lifetime for evaporative cooling to be effective, which takes of order seconds to accomplish. The background collision limited lifetime of a dipole trap loaded in the region of the VC MOT would be on the order of 350 ms, which is too short to allow evaporative cooling. If the vapor pressure in the VC MOT were reduced so the dipole trap lifetime was 10 s, the VC MOT would simply never load. This is why loading atoms into a UHV MOT is necessary.

All the UHV MOT optics are designed to fit inside the volume of the magnetic shield, which is a cylinder with inner dimensions of 18 inches in diameter and 16 inches in length. Six fiber optic cables, each containing a combination of MOT light and repumper light, are brought into the shield. The light from each cable is brought into free space to form the six MOT beams. At their point of intersection, each beam has a radius $r = 0.75$ cm, and the total power in all the beams is 180 mW of MOT light with 14 mW of repumper light. The UHV MOT light is detuned 13.5 MHz to the red of the $F = 2 \rightarrow F' = 3$ transition during the initial loading of the UHV MOT. The UHV MOT coils are 25 turns each with 10 A of current, which forms a field gradient of 12 G/cm.

UHV MOT characterization

The number of atoms in the UHV MOT, as well as the loading time, can be measured experimentally using fluorescence detection in the same manner described in Section 2.3. Figure 2.20 shows a typical loading curve for a UHV MOT with $N = 1.76 \times 10^{10}$ atoms and $\tau = 1.28$ s. The loading rate for the UHV MOT is therefore $L_{UHV} = 1.375 \times 10^{10}$ atoms/s. The transfer efficiency between the VC MOT and UHV MOT is given by the ratio L_{UHV}/L_{out} and is found to be 50%.

A measurement of the UHV MOT lifetime can tell us how low the pressure is in the UHV vacuum chamber using Equation (2.8). The lifetime of the UHV MOT depends mainly on

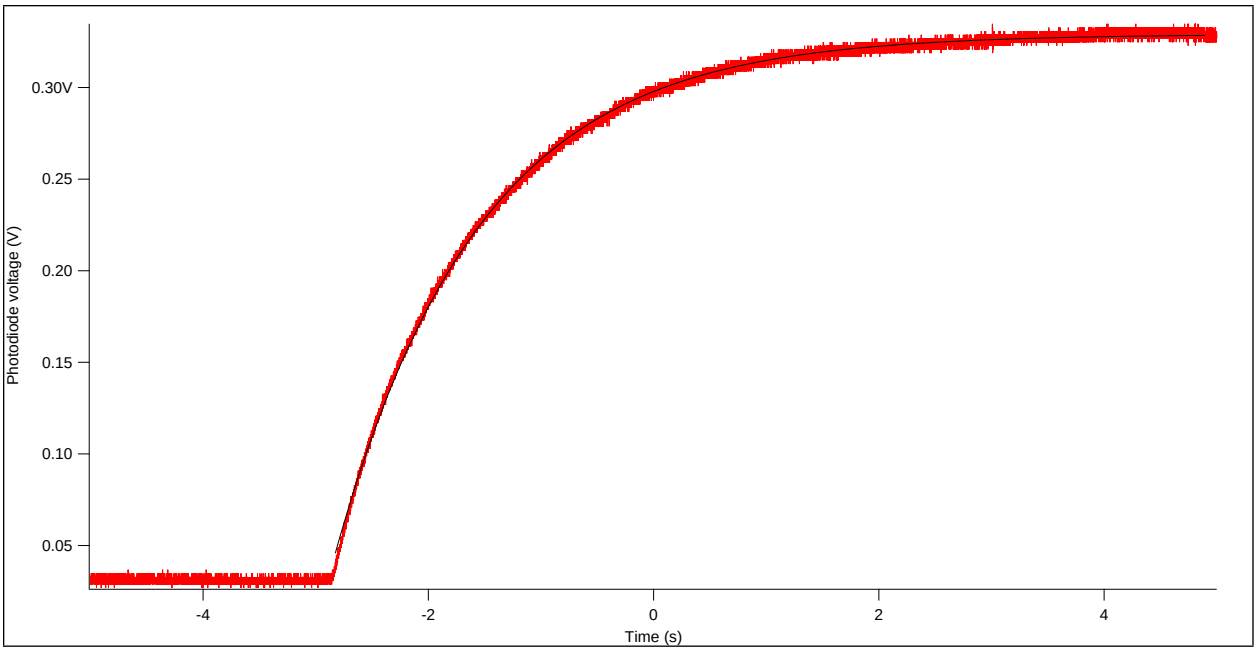

Figure 2.20: This plot shows the photodiode voltage for a typical UHV MOT fluorescence loading curve. The minimum and maximum photodiode voltage levels give the background and fully loaded UHV MOT fluorescence levels, respectively, which can then be used to calculate the number of atoms in the trap. The black curve is an exponential fit to the loading portion of the UHV MOT fluorescence to find the loading time. For this curve, $N = 1.76 \times 10^{10}$ atoms and $\tau = 1.28$ s.

collisions with room temperature background atoms, so long lifetimes indicate low background pressures. To measure the UHV MOT lifetime, the atom source for MOT loading (the pusher beam) is turned off while the remaining atoms are held in the trap. The MOT fluorescence is observed over time to determine the lifetime. Figure 2.21 shows a lifetime curve for the UHV MOT that indicates a long $\tau_c = 56$ s lifetime, which implies a background vapor pressure $p = 8.8 \times 10^{-}11$ torr. The VC MOT lifetime cannot be measured in this way because the loading source cannot be turned off. According to Equation (2.6), the loading time of the UHV MOT should be the same as the UHV MOT lifetime. We see here that the UHV MOT lifetime is much longer than the loading time, however. During the MOT loading, the pusher beam is on, while the lifetime measurement was taken in the absence of the pusher beam. This indicates that, while providing a source of atoms for the UHV MOT,

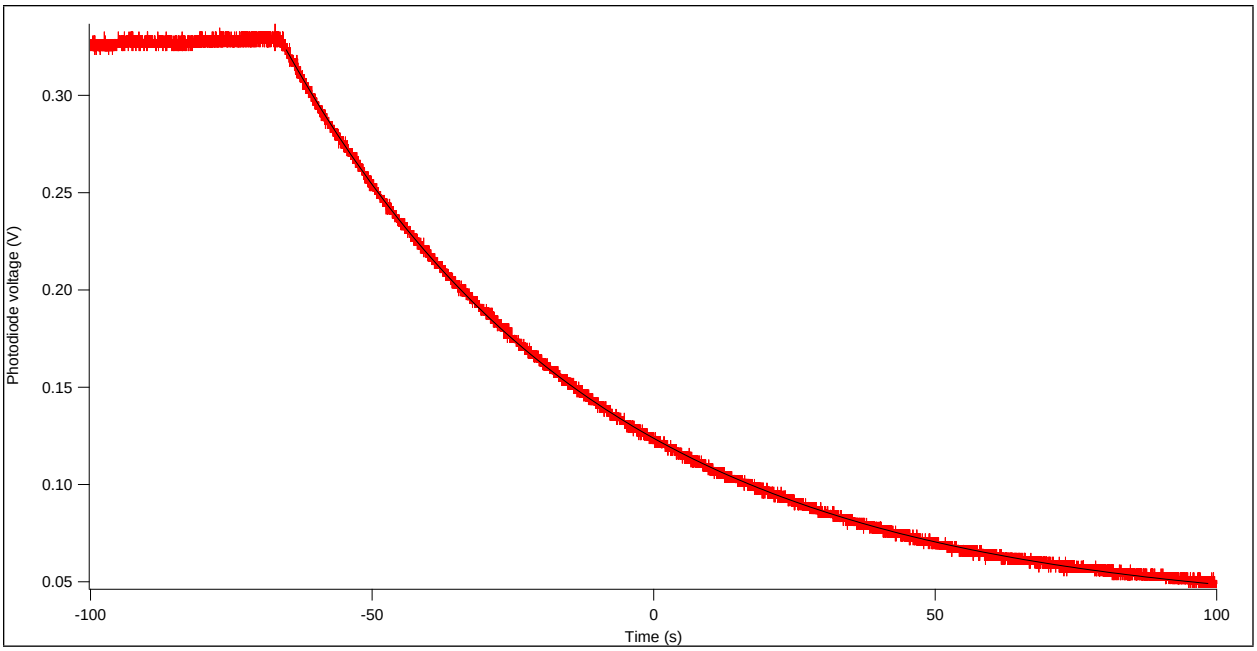

Figure 2.21: This plot shows the photodiode voltage for a typical UHV MOT fluorescence lifetime curve. Once the pusher beam is turned off, the UHV MOT begins to decay with a lifetime dependent on the background pressure in the system. The black curve is an exponential fit to the decay portion of the UHV MOT fluorescence to find the lifetime. For this curve, $\tau = 56$ s.

the pusher beam also introduces a loss term $\gamma_{p,UHV}$ in addition to losses caused by collisions with background atoms γ_c. In this case, Equation (2.5) becomes

$$\frac{\mathrm{d}N}{\mathrm{d}t} = L - N(\gamma_c + \gamma_{p,UHV}). \tag{2.20}$$

Equation (2.6) is still valid, now with

$$\tau = \frac{\tau_c \tau_{p,UHV}}{\tau_c + \tau_{p,UHV}}. \tag{2.21}$$

3 A widely tunable laser frequency offset lock with digital counting[1]

[1] J. Hughes and C. Fertig. 2008, *Rev. Sci. Instrum.*, 79(10):103104, Reprinted here with permission of publisher. Copyright 2008, American Institute of Physics.

3.1 Abstract

We demonstrate a hybrid analog+digital electronic lock to stabilize a dynamically tunable RF frequency offset between two lasers. Our method features an 80 MHz capture range, ± 7 GHz tuning range, frequency agility of 1 MHz/μs, and low (< 30 ppm) drift in the absolute optical frequency difference after $\sim$1000 s. With this scheme, multiple slave lasers can easily be referenced to one stable master laser, while each remains rapidly and accurately tunable over the wide frequency ranges encountered in typical laser cooling and trapping experiments.

3.2 Introduction

Atomic laser cooling and trapping experiments typically require several independently tuned laser beams distributed in frequency across atomic ground and excited state manifolds. For example, in ^{87}Rb the $5S_{1/2}$ ground state and $5P_{3/2}$ excited state hyperfine manifolds (connected by the 780 nm D2 line) have widths of 6.8 GHz and 500 MHz, respectively. In the course of a typical experiment it is often necessary to rapidly and accurately scan or jump the frequency of one or more lasers over these ranges, while maintaining absolute frequency reference to a particular transition.

To meet these goals, we have developed a scheme to electronically lock a rapidly, widely and accurately tunable frequency difference between two lasers based on digitally counting their optical beat frequency. Here we report on two particular implementations of the scheme: a "small-Δ" lock, in which a slave laser is locked 120 MHz above a stable master laser; and a "large-Δ" lock, in which the slave is locked 6.7 GHz above the master, using nearly the same electronics.

3.3 Structure of the paper

The remainder of the paper is structured as follows. In Section 3.4 we review a number of published alternative schemes for locking the frequency difference between two lasers. In Section 3.5 we present a general description of our method, specializing to the two configurations realized in our laboratory (viz., the small- and large-Δ locks described above). In Section 3.6 we analyze theoretically the performance limits set by quantization error for a frequency servo loop based on sampled digital counting. In Section 3.7 we present the results of experiments assessing the stability, accuracy and agility of the locks. We conclude in Section 3.8 with a brief discussion of potential improvements.

3.4 Comparison to other schemes

A large number of techniques to stabilize the frequency difference between two lasers have been developed over the long history of laser spectroscopy. Here we review two broad categories that are most closely related to our method, and (not coincidentally) are often employed in laser cooling and trapping laboratories.

In *optical injection locks* a weak "seed" beam derived from a master laser is launched into the cavity of a slave laser, forcing the slave to lase on a longitudinal mode having frequency equal to that of the seed. Frequency offsets between master and slave of 100's of MHz can be achieved by frequency shifting the seed beam with an acousto-optic modulator (AOM). [54] In this case, not only is the frequency difference between the lasers locked, but their optical fields are in fact phase coherent. Disadvantages are that careful optical mode matching of the seed beam to the slave laser is required; that the dynamical tuning range is limited by the modulation bandwidth of the AOM (typically less than 50 MHz); and that the range of possible offset frequencies is limited by the small selection of commercially available AOMs.

Replacing the AOM with an electro-optic modulator [55] (EOM) can extend the dynamical tuning range beyond 100 MHz. EOMs are available over a wide range of operating

frequencies up to 10's of GHz. However, these benefits come at considerably higher cost and additional complexity, to the extent that optical injection locking with EOM modulation is typically used only when phase-coherence is essential.

It is also possible to use FM sidebands imposed on the master laser as seed beams [56, 57]. While a large range of offsets can be achieved with this technique, it is only suitable for diode master lasers, and when only one or possibly two slaves are to be locked to one master, as it is challenging to isolate the effects of the spatially overlapping and co-propagating sidebands.

Electronic offset locks, the family to which ours belongs, work by electronically detecting the optical beat note between beams from the master and slave lasers superimposed on a fast photodiode, then processing this signal to steer the frequency of the slave laser to maintain the desired offset. An electronic offset lock designed around a commercial frequency-to-voltage (F-to-V) IC is described in Ref. [58]. Unfortunately, commercial F-to-V chips are limited in operating frequency to at most 1 MHz, necessitating a deleteriously large prescaling factor to realize our small-Δ lock, and rendering unfeasible our large-Δ lock. Furthermore, the charge pump output stage of commercial F-to-V chips typically exhibits a large (20%) ripple that leads to undesirable dithering of the slave laser frequency even with aggressive low-pass post-filtering. The hybrid analog+digital locking scheme we demonstrate here realizes high performance F-to-V conversion by using fast AC CMOS electronics to digitally sample and count the beat note frequency and a 10-bit DAC to generate an analog error signal suitable for processing by analog servo electronics.

Ref. [59] describes an electronic offset lock based on a RF Mach-Zehnder interferometer. The RF beat note is split, transmitted through two unequal lengths of coaxial cables and recombined in an RF power combiner. The length of one cable is manually adjusted to produce, through total destructive interference, a node in the total transmitted power when the difference frequency between the two lasers is the desired value. This scheme has the advantages of large capture range and a continuous selection of possible offset frequencies. Unfortunately, since the offset frequency is tuned by a manual adjustment of the cable

length, real-time, rapid tuning of the offset is impossible without additional frequency shifting components (e.g., AOMs). Another serious drawback is the lack of absolute frequency reference for the offset, which is determined by the physical dimensions of cables.

Finally, electronic optical phase-locked-loops (OPLL) [60, 61, 62, 63] have been used to lock two lasers with a tunable offset. Our lock is well suited to experiments that do not warrant the complexity of an OPLL, or that involve broad-line lasers for which the wideband control demands of on OPLL would be difficult to meet. (Our control loop has a bandwidth of $\sim$ 100 kHz, but there is no critical level that must be met for stable operation, unlike for an OPLL.) Also, our scheme allows for multi-GHz tunable offsets without the need for a tunable local oscillator in the microwave range, as is typically employed in OPLLs.

Our hybrid analog+digital electronic frequency offset lock is appropriate when a single, low-cost, flexible solution is required for multiple locks spanning a wide range of offsets from one stable master laser.

3.5 Description of the method

Fig. 3.1 shows the essential optical and electronic components of our scheme. A few milliwatts of light from the master laser (optical frequency ν_m) and slave laser (optical frequency ν_s) are combined on a fast photodetector. The induced photocurrent comprises a DC component proportional to the total incident optical power and an AC component oscillating at the optical beat note frequency $\Delta\nu_{\mathrm{obn}} = \nu_m - \nu_s$. The AC signal is amplified, prescaled in the case of the large-Δ lock, shifted via heterodyne mixing, and ultimately counted with digital electronics. The digital frequency counting starts with a fast zero-crossing detector generating a train of logic pulses synchronous with its analog input. The number of logic pulses (typically 200) in each sample window of duration $\mathcal{T} \approx 5$ μs is counted by a gated 12-bit binary counter. At the conclusion of each sample window the counter's data is latched into a digital-to-analog converter (DAC) that generates an output voltage V_{DAC} proportional to the average frequency difference between the two lasers during that window. Subtracting

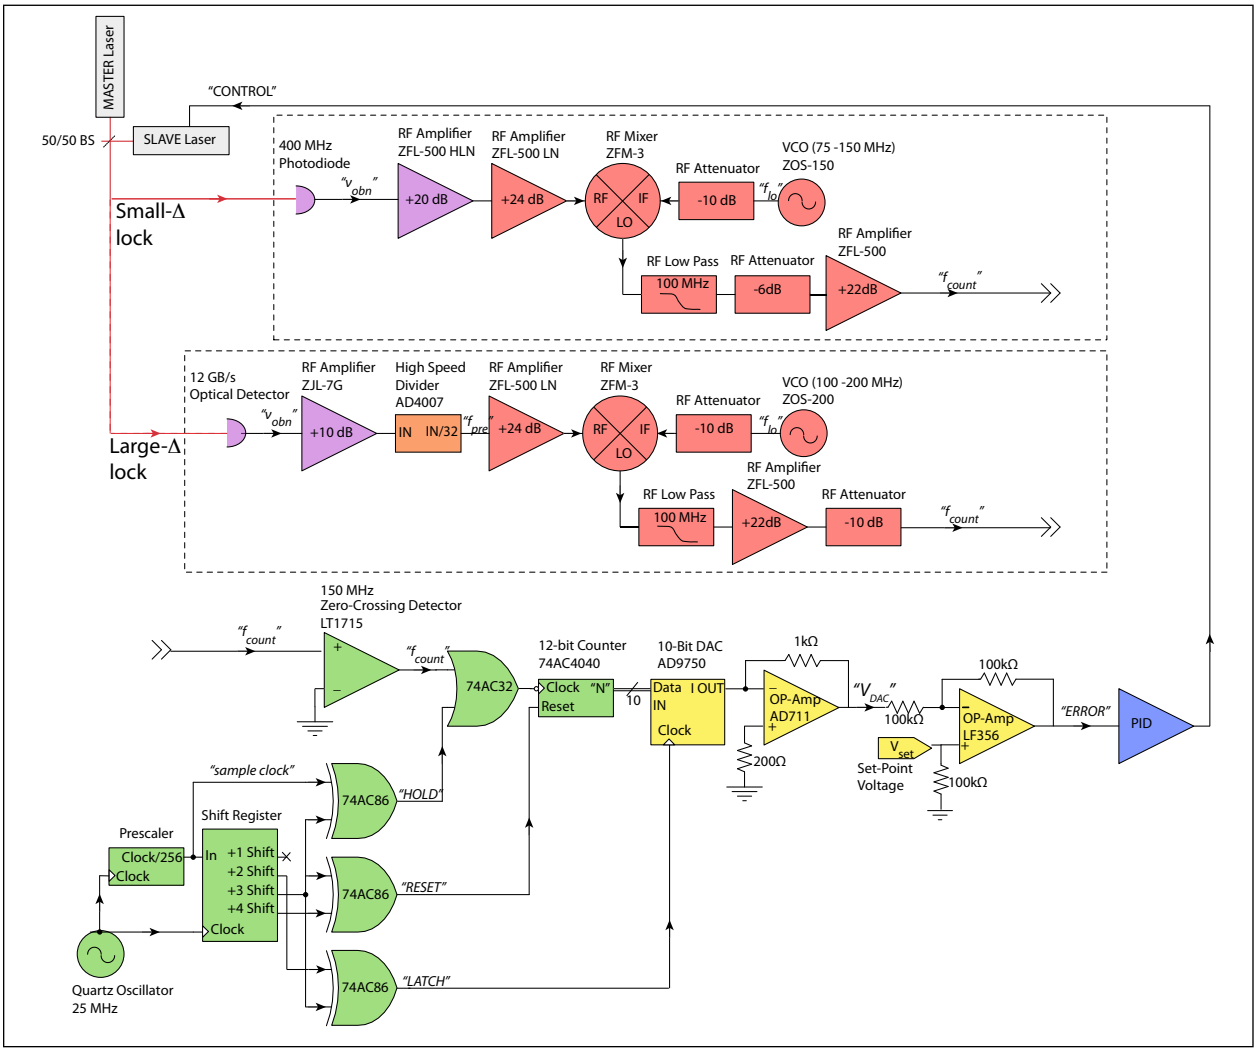

Figure 3.1: (Color online) Block diagram of circuit. Red arrows represent laser light, black arrows electronic signals. Violet symbols: detection and amplification of the optical beat note. Red symbols: heterodyning stage. Orange symbols: fast prescaling stage for large-Δ lock. Green symbols: digital counting. Yellow symbols: D-to-A and generation of error signal. Blue symbols: PID feedback control.

$V_{\rm DAC}$ from a computer-controlled set-point voltage $V_{\rm set}$ in a differencing amplifier produces an error signal suitable for subsequent processing by a PID (proportional-integral-differential) servo controller. The servo controller generates an electronic control signal that steers the slave laser frequency to maintain the desired offset with respect to the master.

We implement the digital electronics using AC CMOS logic, which has a theoretical maximum operating frequency of ~ 120 MHz. Since both our small- and large-Δ locks are intended to maintain values of $\Delta\nu_{\rm obn}$ at or exceeding this limit, we must take steps to translate the frequency of the optical beat note down into the working range of our counting electronics. For example, to count the $\Delta\nu_{\rm obn} = 120$ MHz beat note of our small-Δ lock, we heterodyne the optical beat note with a local oscillator of frequency $f_{\rm LO} = 80$ MHz in an RF mixer, then digitally count the *down*-shifted copy of the beat note at $f_{\rm count} = \Delta\nu_{\rm obn} - f_{LO} = 40$ MHz. Besides shifting the beat note into the working range of the counting electronics, the heterodyning step has other significant benefits. First, it provides an attractive method to tune the master-slave frequency difference by adjusting f_{LO} while leaving the reference voltage $V_{\rm set}$ centered in the circuit's capture range.[2] Second, by counting the *up*-shifted copy of the beat note from the RF mixer, we can lock a master-slave offset frequency for which the optical beat note would otherwise be too low in frequency to digitally count with the fidelity required for a high-bandwidth servo loop.

To count the $\Delta\nu_{\rm obn} = 6.7$ GHz beat note of our large-Δ lock we use a high bandwidth photoreceiver (Discovery Semiconductor DSC-R402-89) and divide the frequency of the optical beat note by 32 using a high frequency divider IC, producing a signal at the prescaled frequency $f_{\rm pre} = \Delta\nu_{\rm obn}/32 = 209$ MHz. The prescaled signal is then RF-heterodyned as in the small-Δ lock, though with $f_{LO} = 174$ MHz, to produce a (down-shifted) signal at $f_{\rm count} = f_{\rm pre} - f_{\rm LO} = 35$ MHz, suitable for digital counting.

[2]While the upper end of the capture range is straightforwardly set by the maximum operating frequency of the digital electronics, the lower end is (less obviously) set by the frequency at which the master and slave lasers cross (i.e., $\Delta\nu_{\rm obn}$ changes sign), or at which $f_{\rm count}$ crosses $\Delta\nu_{\rm obn}$, whichever is higher. It is at this point that the "lock" becomes an "anti–lock" (i.e., the servo loop feedback becomes positive, rather than negative).

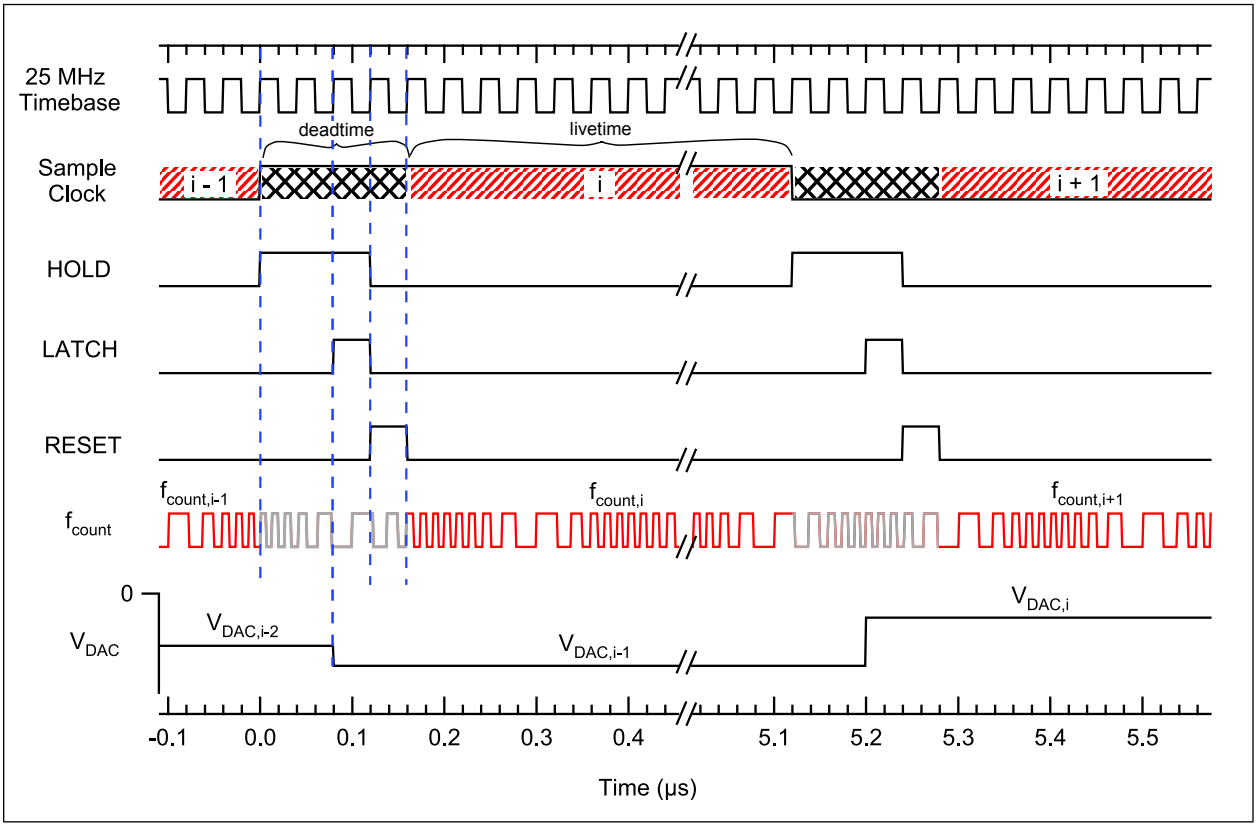

Figure 3.2: (Color online.) Timing diagram of digital counting. Vertical blue dashed lines are guides to the eye. Counting occurs only in red shaded regions of the sample clock waveform, yielding a livetime fraction of ∼97% (note the broken time axis). V_{DAC} is updated at the conclusion of a given counting window; for example, $V_{\mathrm{DAC,i\text{-}1}}$, proportional to the average counted frequency in the $(i-1)$ sample window, is output by the DAC while the binary counter is accumulating counts in the i-th window. This ~ 5 μs latency adversely effects the stability of the servo loop at frequencies near the Nyquist frequency $1/\mathcal{T} \approx 200$ kHz, and thus imposes a limit on the loop gain.

The post-heterodyne signal at f_{count} is input to a high speed zero-crossing detector that generates a synchronous AC CMOS digital signal.[3] This signal is binned and counted in sample windows of duration $\mathcal{T} = 4.96$ μs at a rate of 195.3 kSamples/s. The sample clock is derived from a 25.000 MHz crystal oscillator divided by 256 with an 8-bit binary counter. As shown in Figs. 3.1 and 3.2, a multiple-output shift register generates three auxiliary logic signals from the 195.3 kHz sample clock, shifted by +2, +3, and +4 cycles of the 25 MHz crystal oscillator (i.e., by 80, 120 and 160 ns, respectively). The unshifted and "+3" signals are XOR'd to generate the "HOLD" signal; "+2" and "+3" are XOR'd to generate the "LATCH" signal; and "+3" and "+4" are XOR'd to generate the "RESET" signal. The rising edge of RESET asynchronously clears the 12-bit counter; counting commences 40 ns later, after the fall of RESET, and continues for $\mathcal{T} = 4.96$ μs (the "sample window"), until HOLD toggles high. (A logical-high HOLD suspends counting by preventing falling edges from the zero-crossing detector from arriving at the 12-bit counter.) As long as RESET and HOLD are both low, the 12-bit counter will count logic pulses from the zero-crossing detector. The 80 ns delay between the close of the counting window and the start of data conversion by the DAC (on the rising edge of LATCH) is included to satisfy a minimum data setup time specification of our DAC. The result of this gating logic is a livetime fraction of approximately 97%.

The lowest 10 bits of the 12-bit counter serve as the inputs to a 10-bit DAC. The DAC is operated in buffered, single-ended output configuration with an output voltage range of $V_{\text{DAC}} = 0$ to -11.6 V. To generate the "ERROR" signal for the analog PID servo loop, V_{DAC} is subtracted from the computer generated set-point voltage V_{set} in a differencing amplifier. We choose $V_{\text{set}} \approx -2$V to -3V separately for the small- and large-Δ configuration so that f_{count} is approximately centered in the capture range when the servo loop is locked.

[3]For reasons that are not fully understood, our zero-crossing detector was fairly sensitive to input level, which was fine-tuned to optimize performance.

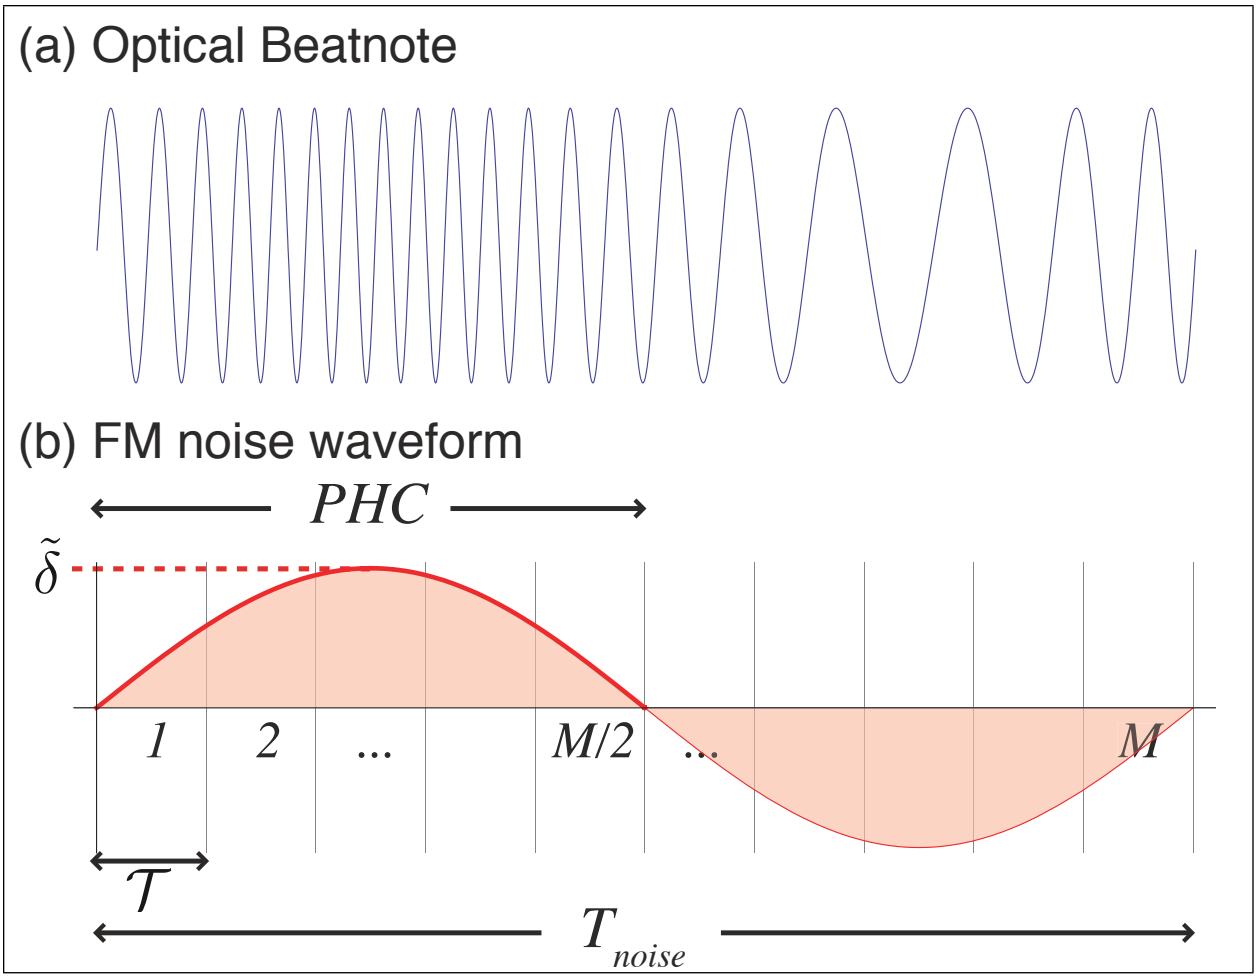

Figure 3.3: (Color online.) Detecting frequency fluctuations by discretely sampled counting. (a) Pictorial representation of the optical beat note frequency modulated by the noise waveform of (b). (b) Representative FM noise waveform, identifying various symbols used in the text.

3.6 Quantization error limited stability of a digitally counted frequency servo loop

There is a fundamental limit to the frequency stability of a lock such as ours that employs sampled digital frequency counting. Suppose that $N = M\mathcal{T}f$ counts (i.e., zero crossings) of a noiseless carrier at frequency f are detected in M sample windows of duration $\mathcal{T}$ (neglecting deadtime). Introduce monochromatic FM noise $\delta f(t) = \tilde{\delta}\sin(2\pi f_{\rm noise}t)$ on the carrier. For simplicity we restrict our calculation to noise frequencies with period an integer number

of sample windows: $T_{\rm noise} = M\mathcal{T}, M = 2, 3, 4, \ldots$. This FM noise on the carrier causes *additional* counts[4]

$$\Delta = \int_{\rm PHC} \delta f(t)\ dt = \frac{T_{\rm noise}\tilde{\delta}}{\pi}$$

during the noise waveform's positive half-cycle ("NWPHC," from $t = 0$ to $t = T_{\rm noise}/2$ (see Fig. 3.3), where $T_{\rm noise} = 1/f_{\rm noise}$. We define the *minimum detectable noise amplitude* $\tilde{\delta} = \tilde{\delta}_{\rm min}$ as that fluctuation of the carrier frequency that causes exactly $\Delta = 1$ to accumulate over the NWPHC. With the perspective that the "signal" we wish our circuit to detect (and subsequently to correct via feedback) is the instantaneous frequency fluctuation $\delta f(t)$, and that the "noise" that disturbs this effort is the unavoidable ± 0.5 count uncertainty in the least significant bit of the binary counter, we say that the fluctuation $\tilde{\delta}_{\rm min}$ will be sensed by the circuit with a quantization error limited ("QEL") signal-to-noise ratio $S/N \sim 1$, neglecting all other sources of noise. A fluctuation smaller than $\tilde{\delta}_{\rm min}$ would go unnoticed (and therefore uncorrected) by the circuit, while larger fluctuations would be sensed (and corrected) with $S/N > 1$.

Using relations developed in the preceding discussion, we find

$$\tilde{\delta}_{\rm min} = \frac{\pi}{T_{\rm noise}}.$$

In a sampled digital frequency counter such as ours, the minimum detectable frequency fluctuation of the carrier decreases with increasing FM noise period (decreasing FM noise frequency). This imposes a fundamental stability limit to a carrier whose frequency is controlled by a feedback loop that employs such a counter. In our lock, the "carrier" is the optical beat note at $\Delta\nu_{\rm obn}$. The QEL implies that fluctuations of the average optical beat note frequency between successive time-averages of duration τ cannot be smaller than $\tilde{\delta}_{\rm min}(\tau) = \pi/\tau$, and

[4]We neglect issues related to the relative phases of the counted signal, the noise waveform and the sample clock; such considerations affect the results of this analysis at the fractional level of $1/MN$, which are never greater than 1% for us.

thus will exhibit a (QEL) relative rms deviation of no less than

$$\sigma_{\mathrm{y}}^{\mathrm{QEL}}(\tau) = \frac{\pi}{\sqrt{2}\tau\Delta\nu_{\mathrm{obn}}}.$$

The division of the optical beat note by 32 in the large-Δ lock increases $\tilde{\delta}_{\mathrm{min}}$, and therefore the *absolute* frequency instability, by the same factor. However, since the large-Δ lock prescaled frequency $f_{\mathrm{pre}} = 209$ MHz is 1.7$\times$ larger than the unscaled optical beat note $\Delta\nu_{\mathrm{obn}} = 120$ MHz of the small-Δ lock, the QEL *relative* instability limit is only a factor of $32/1.7 \approx 19$ higher (i.e., worse) for the large-Δ lock. Fig. 3.4 shows lines demarcating the QEL instability limits for both locks, in both relative and absolute units.

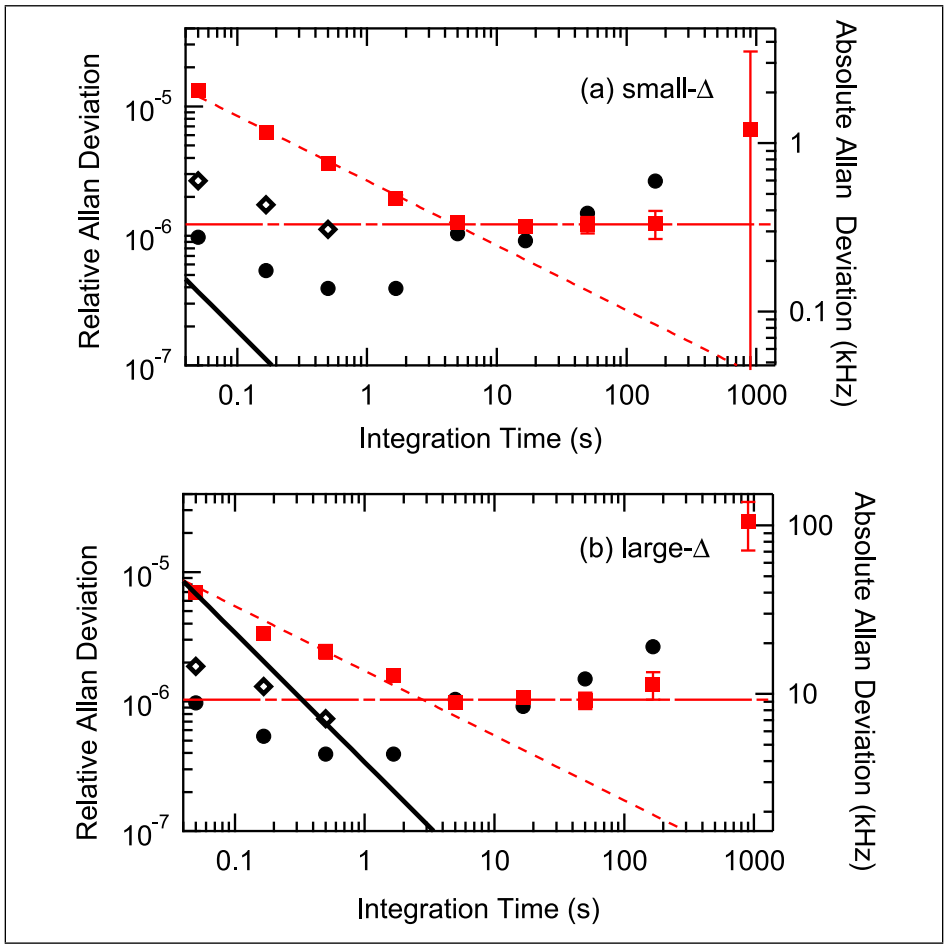

Figure 3.4: (Color online.) Frequency stability. Absolute and relative frequency Allan deviation data are shown as red solid squares for the small-Δ (a) and large-Δ (b) locks, for integrations times 0.05 s $\leq \tau \leq$ 167 s.. Points at 900 s are not true Allan deviations but rather the average deviation of 13 discrete measurements of the beat note taken at 15 minute intervals over 3 hours. The dashed and dashed-dot lines are fits to the measured Allan deviation points, as described in the text. Heavy black lines in both figures show the quantization error limited (QEL) relative instability $\sigma_y^{\mathrm{QEL}}(\tau)$ discussed in Sec. 3.6. Solid black circles are the measured relative Allan deviation of the computer-controlled reference voltage. Open black diamonds show the counting noise floor of the counting electronics as measured in a separate experiment in which the optical beat note was replaced by the output of a high-quality electronic synthesizer set to the same frequency, with the high speed divider either removed (a), or set to divide by 16 (b). There is no contradiction in these measurements lying below the QEL line, as they were made in open loop, whereas the QEL applies only to the closed loop stability of the *locked* laser beat note. Error bars are 1-standard deviation statistical uncertainties.

3.7 Experimental characterization

Here we report the results of several experiments to characterize the performance of both small- and large-Δ locks. To facilitate comparison of the two configurations, the same master laser, slave laser and lock electronics were used in all experiments, save for the explicit differences indicated in Fig. 3.1.

The master and slave lasers are both grating stabilized, external-cavity diode lasers with $\sim$ 1 MHz intrinsic "fast" linewidths. Fig. 3.5 shows a 5 hr time series of the optical beat note frequency for both locks under continuous operation, measured at 10 s intervals with an RF spectrum analyzer.

We evaluated the short time stability of the optical beat note by directly measuring its linewidth over a 10 ms integration time using an RF spectrum analyzer (see Fig. 3.6). The beat note linewidth was measured to have FWHM of 1.5 MHz (4.0 MHz) for the small- (large-)Δ locks, respectively. The broadening of the large-Δ linewidth is qualitatively consistent with quantization error limited performance (see Sec. 3.6).

To characterize frequency stability over longer times, a second, independent copy of the lock electronics was used to make an "out of loop" measurement of the frequency Allan deviation $\sigma_y(\tau)$ of the beat note for integration times 0.05 s $\leq \tau \leq$ 167 s; these data are presented in Fig. 3.4. To make these measurements, fluctuations of the error signal voltage of the out-of-loop counting box were recorded using a digital multimeter, and the measured voltage noise subsequently converted to frequency fluctuations. The deadtime fraction for the measurements is $< 12\%$ and is ignored in the analysis, except for the points at $\tau = 900$ s (see Fig. 3.4 caption). For the large-Δ lock, the data is fit to a white-noise-limited model for 0.05 s $\leq \tau \leq$ 5 s with a result $\sigma_y(\tau) = 1.6 \times 10^{-6}\tau^{-1/2}$. A 1/f noise dominated "flicker floor" model is fit to the data for 5 s $\leq \tau \leq$ 167 s with the result $\sigma_y(\tau) = 1.0 \times 10^{-6}$. Similar fits are made to the small-Δ lock data, yielding $\sigma_y(\tau) = 2.6 \times 10^{-6}\tau^{-1/2}$ for 0.05 s $\leq \tau \leq$ 5 s, and $\sigma_y(\tau) = 1.2 \times 10^{-6}$ for 5 s $\leq \tau \leq$ 167 s. These fits are shown in Fig. 3.4. We could not

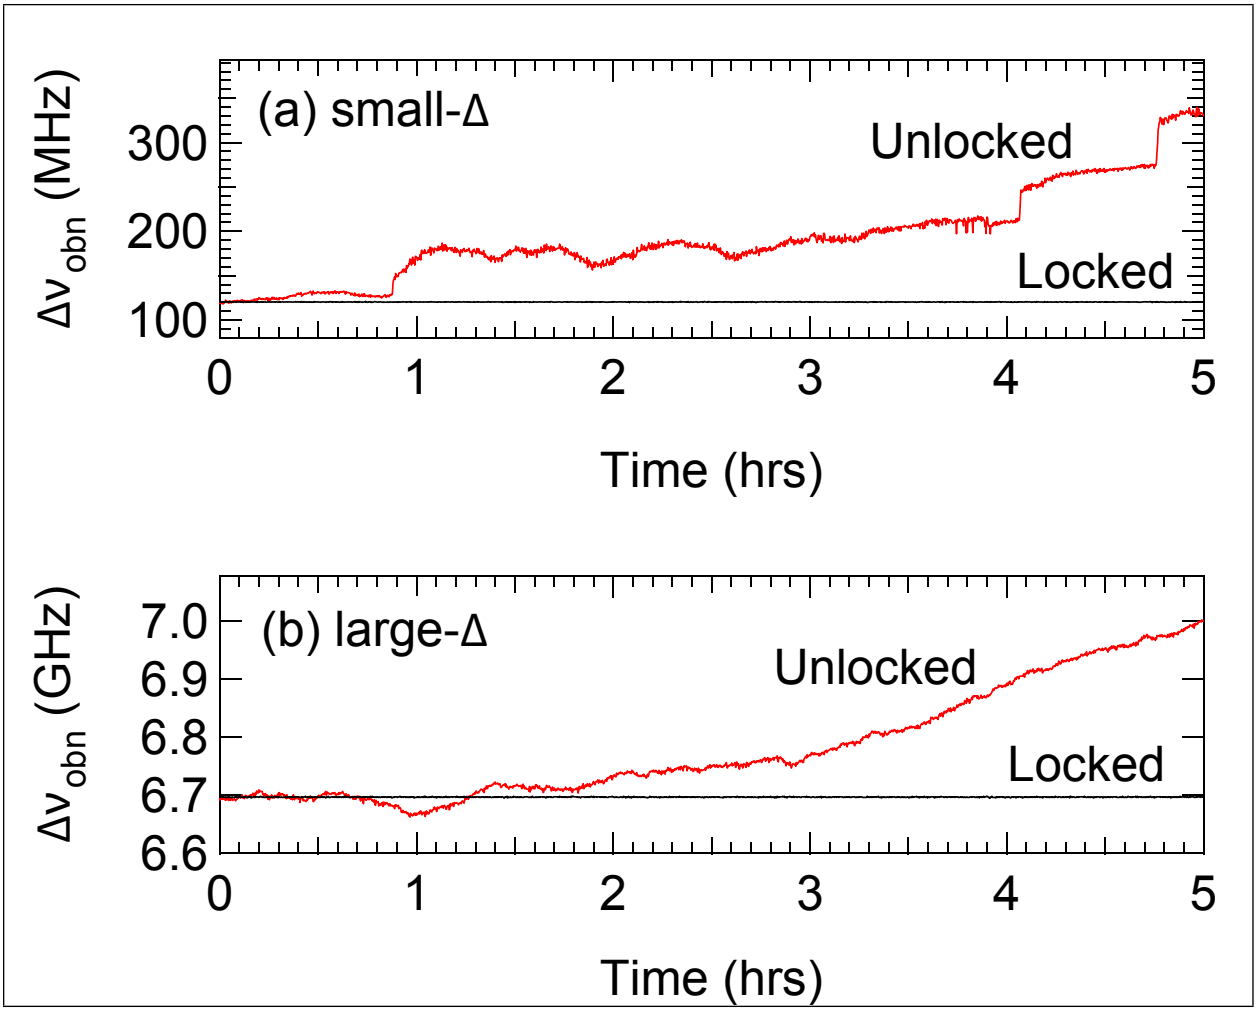

Figure 3.5: (Color online.) Long-term locking. Time series of the optical beat note frequency $\Delta\nu_{\mathrm{obn}}$, acquired over a 5 hour period, for the small-Δ (a) and large-Δ locks (b). For comparison, time series are also shown for the slave laser unlocked.

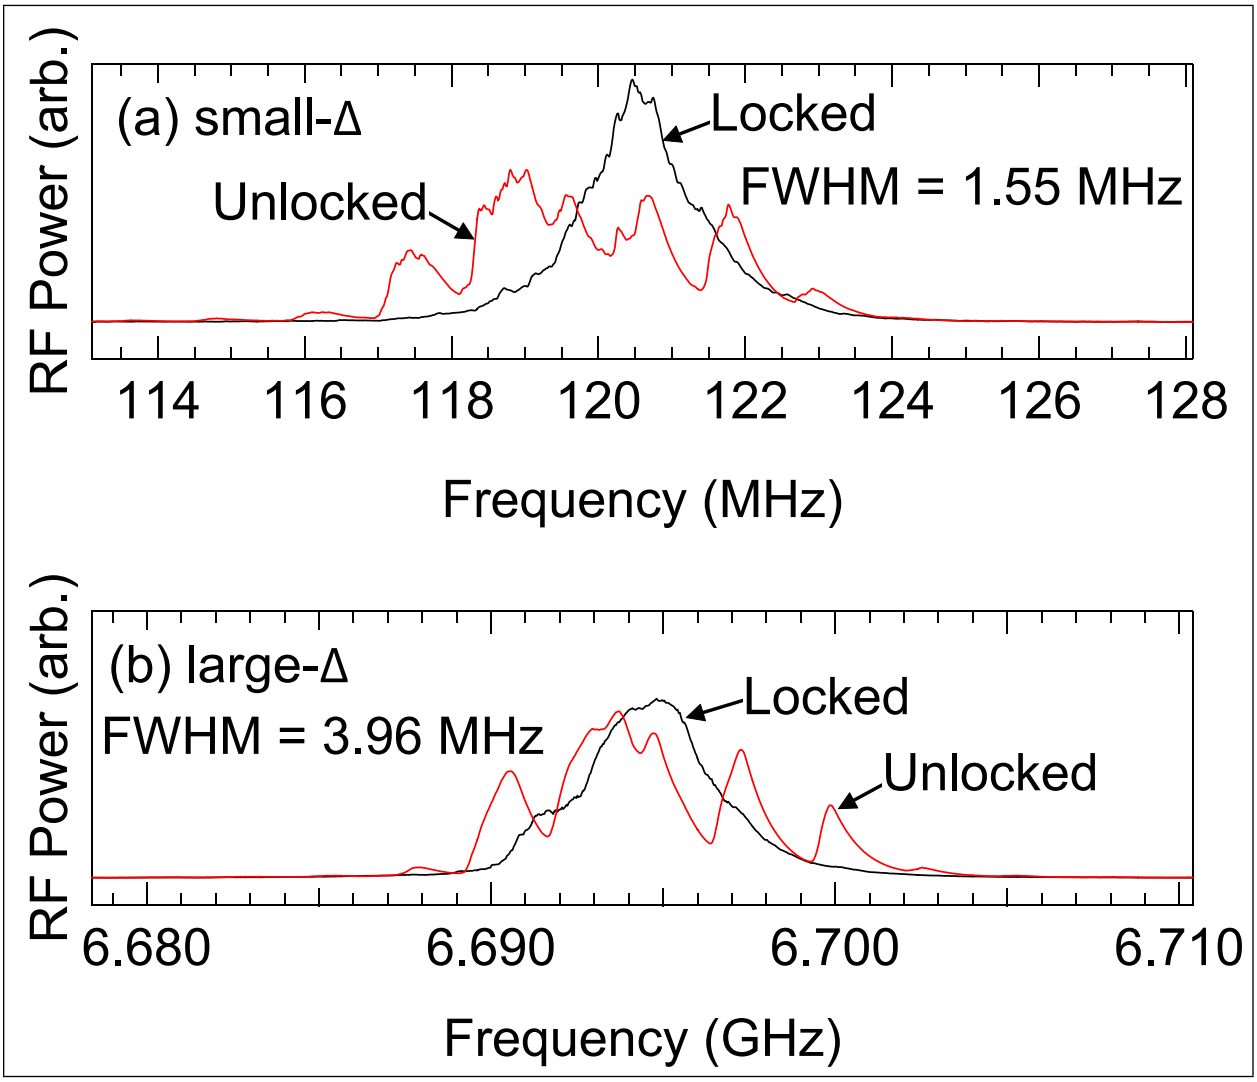

Figure 3.6: (Color online.) Optical beat note linewidth. RF spectrum analyzer traces of the optical beat note for small-Δ (a) and large-Δ (b) locks for the servo loops both locked and unlocked. FWHM linewidths are given for the locked data. For these measurements, the resolution bandwidth of the spectrum analyzer was 10 kHz (100 kHz) for the small- (large-)Δ locks, respectively. The full-screen sweep time was approximately 20 ms for all traces.

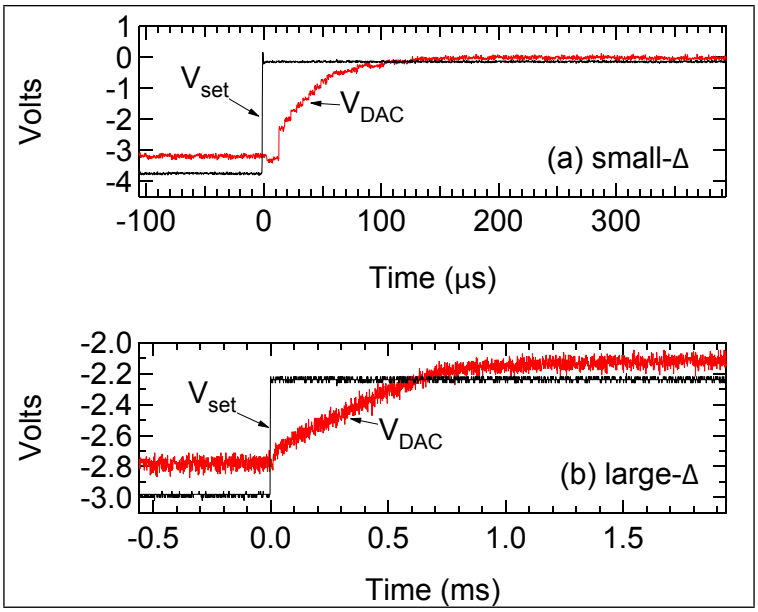

Figure 3.7: (Color online.) Tuning agility. Oscilloscope traces of V_{DAC}, proportional to $\Delta\nu_{\mathrm{obn}}$, show the response of the slave laser to an abrupt step of V_{set}, corresponding to jumps of 70 MHz in the small-Δ lock (a) and 416 MHz in the large-Δ lock (b). The stair-case pattern visible in V_{DAC} in (a) directly reflects the discrete sampling and counting of the beat note. The steps are not visible in (b) due to its compressed time axis.

measure Allan deviations for integration times sufficiently short to probe the onset of QEL performance for either lock. By extrapolating the data we estimate that the small- (large-)Δ locks would reach QEL at $\tau = 5 \times 10^{-5}$ s ($\tau = 0.045$ s), respectively.

Fig. 3.4 also shows measurements of the Allan deviation of the reference voltage V_{set}, which in our system is supplied by a PC-based multifunction card, which are consistent with the observed level of the "flicker floor" for both locks. Inspection of the raw time-series data show that the increase in instability of V_{set} at $\tau = 1$ s is the result of rare, discrete jumps in V_{set}. We do not have an explanation for the jumps, which are much smaller than can be attributed to LSB noise in the PC card's output DAC. In the future we plan to replace the PC card with a higher-performance model, upgrade electrical connections and install opto-isolators to break ground loops, in an attempt to suppress these jumps.

We quantify the tuning agility of the lock as the average frequency slewing rate over the time it takes the beat note to track from 10% to 90% of its asymptotic value ("10-90" time), following a large step in $V_{\rm set}$. Fig. 3.7 shows traces of $V_{\rm DAC}$, proportional to $\Delta\nu_{\rm obn}$, from which we extract a 10-90 time of 71 μs (690 μs) after a 70 MHz (416 MHz) step, for an average slew rate of 1.0 MHz/μs (1.7 MHz/μs) for the small- (large-)Δ locks, respectively. These results are consistent with the designed 100 kHz bandwidth of the servo loop.

The optical beat note can be accurately tuned over a wide range by adjusting the heterodyne local oscillator frequency $f_{\rm LO}$. Fig. 3.8 shows data demonstrating the range and accuracy of closed-loop tuning of the slave laser in which $f_{\rm LO}$ is changed in 20 MHz steps for both the small- and large-Δ locks. (For the large-Δ lock, because of the division-by-32 prescaling of the optical beat note, the tuning steps are 32×20 MHz = 640 MHz.) While the 80 MHz capture range of the servo loop limits the size of an individual abrupt step, smooth adjustments to $f_{\rm LO}$ permit tuning limited on the high-end only by the bandwidth of the RF photodetector and on the low-end by the steepness of the post-heterodyne low-pass filter.[5] As discussed in Section 3.5, we reach the lowest lockable beat notes by selecting the up-shifted heterodyne component, while choosing the polarity of the servo loop appropriate to the sign of $\Delta\nu_{\rm obn}$ so as to yield overall negative feedback.

[5]The lock range is also (obviously) limited by the mode-hop free tuning range of the slave laser.

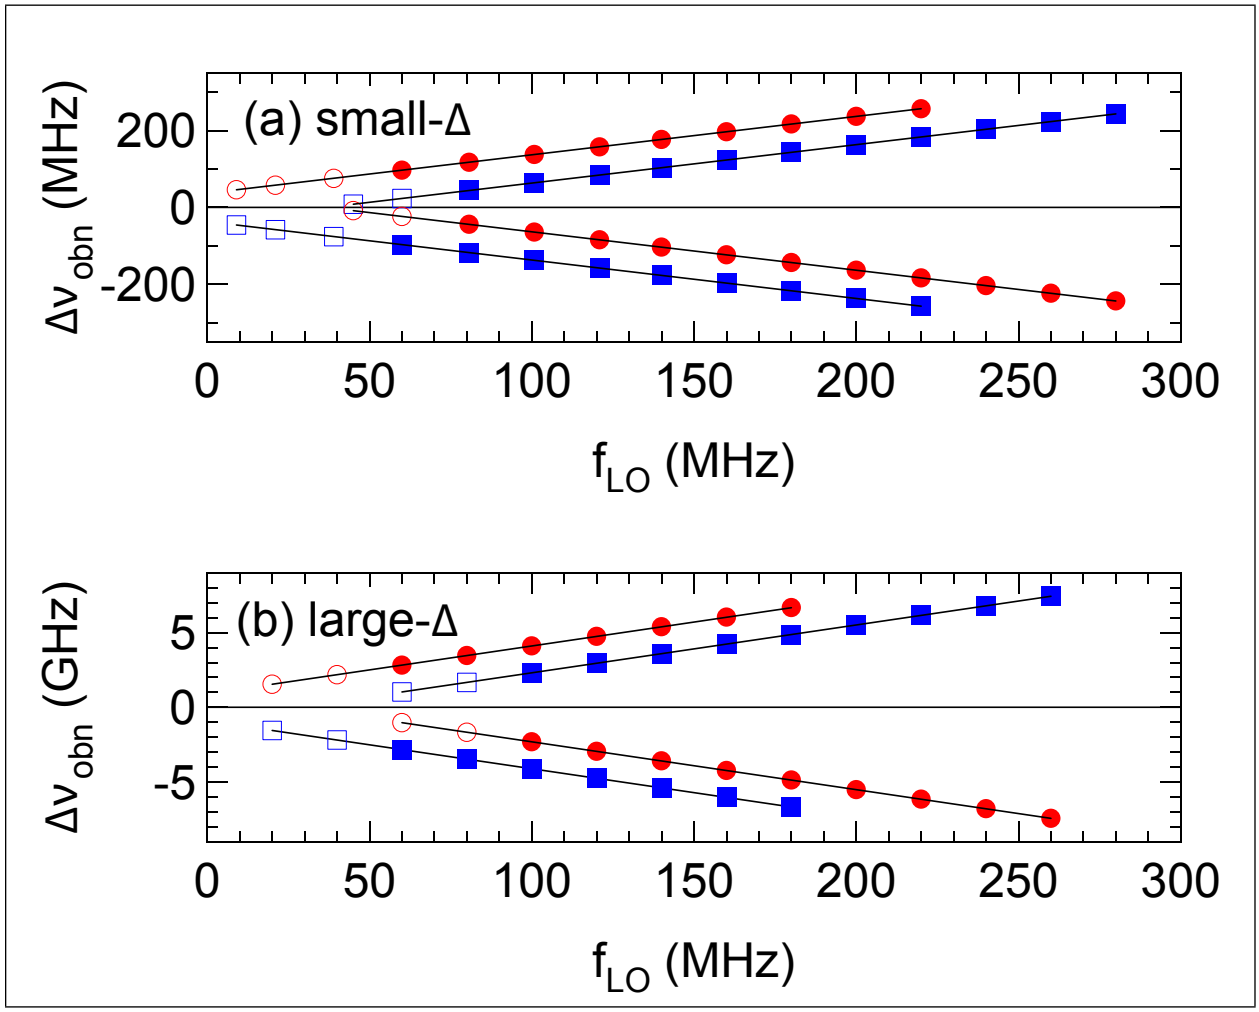

Figure 3.8: (Color online.) Tuning range. The frequency of the locked optical beat note $\Delta\nu_{\rm obn}$, for the small-Δ lock (a) and large-Δ lock (b) versus heterodyne local oscillator frequency $f_{\rm LO}$, counted using an RF spectrum analyzer. A positive (negative) beat note corresponds to the slave laser frequency higher (lower) than the master laser. Solid red circles (solid blue squares) correspond to locking to the down- (up-)shifted copy of the beat note from the heterodyne stage. Open symbols represent measurements made with a different value of the post-heterodyne RF low-pass filter so as to lock the lowest optical beat notes for a given configuration. Curves through the data are lines of slope=1 (=32) through the lowest measured beat note of each series, for the small- (large-)Δ locks, respectively.

3.8 Conclusion

We have designed and implemented a hybrid analog+digital scheme to lock the difference frequency between two lasers. The long term drift, short term stability, tuning range, tuning accuracy and overall flexibility of the scheme make it well suited to applications where the frequencies of several independent lasers need to be stabilized to an absolute reference *and* rapidly and accurately tuned over several GHz. Improvements to the digital electronics, including optimized component layout, better ground plane practices, and the use of stripline leads for signal transport should allow the present design to reach the maximum clock rate for AC CMOS logic of 120 MHz, for a modest increase in performance. A large jump in performance could be realized by switching to a higher speed logic family, such as ECL, which would enable count rates of 500 MHz or more, improving the signal-to-noise, capture range, and closed-loop bandwidth of the lock. It would then be feasible to implement fully digital-domain processing of the error signal, which could greatly improve both the long-term stability and absolute accuracy of the lock.

In conclusion, we have demonstrated a method to lock the offset frequency between two lasers based on the direct digital counting of the optical beat note, which is well suited to atomic laser cooling and trapping experiments that call for many widely tunable laser beams, each with absolute frequency reference.

This work was sponsored in part by grants from the ARO, the DURIP, the ORAU, and the University of Georgia.

4 Crossed Beam Dipole Trap

4.1 Overview of the apparatus

Once the atoms have been loaded into the UHV MOT, further cooling is still required to reach the phase transition to a BEC. This is typically done using evaporative cooling techniques in time-orbiting potential (TOP) or Ioffe-Pritchard magnetic traps [35, 36, 64]. Unfortunately, in an experiment where background magnetic fields in excess of $\sim 1\,\mu$G are potentially devastating, the large fields generated by magnetic trapping schemes are unacceptable. Therefore, this experiment utilizes an all optical dipole trapping scheme, a far off-resonant trap (FORT), similar to the one first implemented in [37].

In our FORT, two tightly focused high power far off resonant laser beams cross at their waists inside the region of the UHV MOT to trap cold atoms. The electric field of the beams induces a dipole moment in the atoms. The interaction potential of the induced dipole moments in the electric field leads to a dipole force, which can be expressed as [53]

$$\vec{F}(\vec{r}) = -\frac{3\pi c^2}{2\omega_0^3}\frac{\Gamma}{\Delta}\nabla I(\vec{r}), \tag{4.1}$$

where Δ is the laser detuning from the D_2 line, Γ is the natural linewidth of the D_2 line, and $I(\vec{r})$ is the intensity of the laser beam. The scattering rate can be expressed as

$$\Gamma_{sc} = \frac{3\pi c^2}{2\hbar\omega_0^3}\left(\frac{\Gamma}{\Delta}\right)^2 I(\vec{r}). \tag{4.2}$$

Because the scattering rate scales as $1/\Delta^2$ while the dipole force only scales as $1/\Delta$, it is critical to use far off resonant light to minimize heating in the trap while still maintaining strong confinement forces.

The dipole trapping beams are generated by an IPG Photonics 1064 nm 50 W linearly polarized, single mode YAG fiber laser with a NP Photonics 1064 nm 30 mW pump. Upon exiting the fiber, the beam passes through a 50/50 beam splitter to form the two beams of the crossed dipole trap. The two beams enter the volume of the magnetic shield in free space through one of the holes on the radial surface of the magnetic shield. Once inside the shield, both beams pass through a single lens located 1 focal length away from the center of the UHV MOT, which focuses the beams onto the atoms. After the beams exit the UHV cell, they are steered out of the magnetic shield in free space through another hole on its radial surface and into a beam dump to prevent stray 1064 nm light from bouncing around the lab.

4.2 Scissor trap

The scissor trap allows the crossing angle of the dipole trapping beams to be dynamically swept over a limited range. By doing so, the aspect ratio of the BEC can be dynamically changed in situ, opening up possibilities for new experimental regimes. In particular, [23] predicts the formation of spontaneous spin textures in a spin-1 BEC when the trap shape dynamically changes from a pancake shape to a cigar shape in magnetic fields below 100 μG. In the future, we would like to use this new scissor trap technique to test these predictions. In our experiment, the crossing angle can range from 0° (perfectly overlapped beams) to a maximum crossing angle of 30°, limited only by the diameter of the focusing lens just before the UHV cell.

Special care must be taken in the design and placement of the optics for the scissor trap. The goal is to achieve an axis of rotation for each beam that overlaps with both waists at all crossing angles. This is a difficult task, but possible with careful alignment. Beams from two 1064 nm YAG fiber lasers with similar divergence properties coming out of the fibers are allowed to expand in free space until both beams have the same $1/e^2$ diameter. At that point, both beams pass through identical focusing lenses to match their spatial profiles. From here, as long as the pathlengths between the corresponding optics for each beam are

the same, the beam diameters and expansion properties of each beam should match once they are recombined onto shared optics and sent into the magnetic shield.

The design and layout of the dipole trap optical system can be seen in Figure 4.1. Coming out of the fiber, the beam diverges very slightly. After exiting the fibers, each beam is allowed to expand in free space until they each reach a diameter of 1 mm. At this point, each beam passes through a high power YAG anti-reflection coated lens, L_1, with focal length $f_1 = 75$ cm. An AOM is centered at the waist of each focused beam with an RF input frequency that can range from 60 MHz to 90 MHz. At the output of the AOMs, the $+1$ diffracted order of one beam is used, while the -1 diffracted order of the other beam is used. Because the dipole trapping beams are far off-resonant, this small change in frequency by the AOMs does not change the shape of the dipole trapping potential. By shifting each beam's frequency in opposite directions, however, standing wave interference patterns are prevented in the trap. Another lens, L_2, with the same focal length as L_1, is located slightly more than $2f_1$ away from L_1. This almost collimates the beam, but still allows it to converge slightly to a shallow waist. After each beam passes through L_2, they each pass through a $\lambda/2$-waveplate where one beam is rotated to S-polarization while the other is rotated to P-polarization. These two orthogonally linearly polarized beams are then recombined on a single polarizing beam splitting cube, PBS_1, and share optics for the rest of their beampaths.

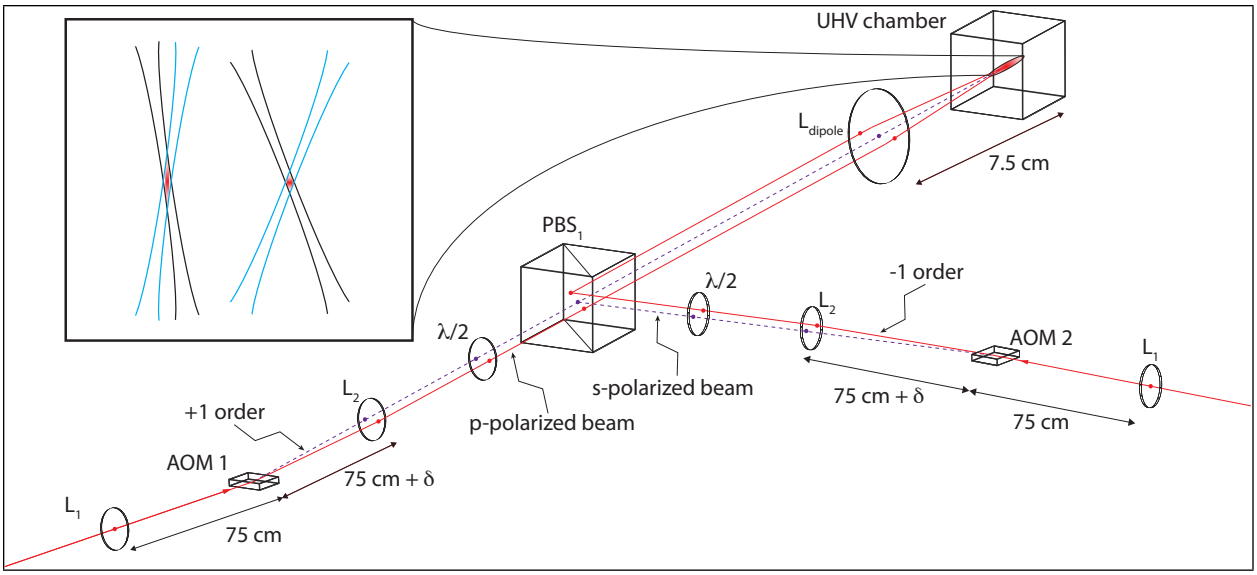

Figure 4.1: This is a diagram of the optics used to make the scissor trap. The solid red lines represent a maximum crossing angle of 30°, which occur at an AOM rf input frequency of 90 MHz. The dashed purple lines represent perfectly overlapped dipole trapping beams, which occur at an AOM rf input frequency of 60 MHz. The inset in the diagram shows how different crossing angles change the aspect ratio of the trap. Note: not to scale.

Inside the magnetic shield, the dipole trap focusing lens, L_{dipole}, with focal length $f_{dipole} = 7.5$cm, is located f_{dipole} away from the center of the UHV MOT and aligned so each dipole trapping beam passes directly through the center of the UHV MOT. The path-length between L_2 and L_{dipole} is arranged so the shallow waist occurs at a distance f_{dipole} before L_{dipole}. This arrangement ensures that the intersection between the beams occurs at the beam waists for all crossing angles.

The "scissoring" behavior of this configuration comes from the ability to change the RF input frequency of the AOMs, thereby deflecting the diffracted beams by different angles. Typically, this would be done using a voltage controlled oscillator with a computer controlled input voltage for maximum experimental flexibility and timing control. The two beams are first independently aligned using a RF input frequency of 60 MHz for both AOMs. When the beams are recombined onto PBS_2 at this frequency, they are aligned to be perfectly

copropogating as they enter the magnetic shield. The best way to check the beam overlap is to place a beam profiler at various points along the beam path and steer mirrors until the beams are perfectly overlapped at all points. With perfectly overlapped beams, they can then be steered using their shared optics into the shield and carefully aligned through the center of L_{dipole}. As the RF input frequency of the AOMs is increased, the deflection angle of the beams increases, and the beams shift away from their 60 MHz overlapped paths on the same plane by equal distances. As long as L_{dipole} is aligned perpendicular to the beams, all beams that pass through any point of the lens will pass through a point in space f_{dipole} after the lens, which is designed to be in the center of the UHV MOT. If all optics in the system have been properly placed, this should also be where both beam waists occur, allowing for dynamic control of the intersection angle of the beams at the waists by smoothly ramping the RF input frequency of the AOMs.

4.3 Loading procedure from the UHV MOT

In order to efficiently load atoms into a dipole trap, the atoms must first be precooled and loaded from the UHV MOT. The procedure is to first load as many atoms as possible into the UHV MOT for transfer into the dipole trap, but at this point the conditions are still not optimal for the transfer of the largest number of atoms. The UHV MOT has a large volume ($1.77\,\mathrm{cm}^3$) with a low atomic density ($9.9 \times 10^9\,\mathrm{atoms/cm}^3$) and very little overlap with the dipole trapping beams. The trapping volume for a crossed dipole trap at the maximum scissor angle is $8.75 \times 10^{-7}\mathrm{cm}^3$, so only 0.0000494% of the UHV MOT overlaps with the crossed region of a dipole trap in this configuration. Therefore, in order to transfer the largest number of atoms into the dipole trap, the UHV MOT must be further compressed to increase the density in the region of the dipole trapping beams.

Once the UHV MOT is fully loaded, the first step in the process is to turn off the source of atoms into the UHV chamber by turning off the pusher beam. This incoming flux of atoms can potentially scatter loaded atoms out of the dipole trap, as evidenced by the impact the

pusher beam had on the UHV MOT lifetime when it was on. After the pusher beam has been fully extinguished by a shutter, additional cooling of the MOT atoms may begin. This is called the "dark detuned MOT" phase. The MOT beams are stepped down in power from fully on (180 mW) to 60% power (108 mW), and the repumper beams are stepped down in power from fully on (14 mW) to 0.8% power (100 μW), creating what is called a dark MOT. Under normal MOT conditions, collisions between ground and excited state atoms can lead to trap loss, limiting the maximum number of atoms that can be trapped by the MOT. In a dark MOT, by reducing the power in the repumper beams, most of the atoms are pumped into the lower $F = 1$ hyperfine "dark" ground state, which is only weakly coupled to the trapping light field, thereby reducing the number of collisions between atoms in the ground and excited states. MOT density is also limited by the reabsorption of scattered photons, which causes an outward radiation pressure that balances the confinement of the MOT light. By pumping atoms into a dark state that is weakly coupled to the trapping beams, these rescattering processes are diminished [65, 66, 67]. Therefore, by entering a dark MOT state, the MOT density increases, giving the dipole trap access to a larger number of atoms.

Additionally, the MOT beam detuning is linearly ramped from -2.2Γ to -9.8Γ over 190ms. It is possible that a larger final detuning would result in improved loading of the dipole trap, but in our apparatus the amount of detuning is limited by the capture range of the frequency offset lock described in Chapter 3. It is also possible that a sudden step would improve performance over a ramped change of the detuning, but this was limited by the response time of the MOT beam laser to its own reference cavity stabilization lock. The decrease in MOT beam power and increase in detuning serves to further decrease the coupling of the atoms to the light field, thus also reducing the effects of multiple scattering processes and increasing MOT density.

The next step to achieve efficient loading of the dipole trap is compression of the MOT into a dense sample around the region of the dipole trapping beams. With the repumper beams still at 100 μW, the MOT beams are returned to 180 mW. At this point, the current

in the UHV MOT coils is stepped from 10 A to 20 A, which corresponds to an increase in the field gradient from 12 G/cm to 24 G/cm. This sudden increase in the field gradient compresses the UHV MOT into a dense region around the intersection of the two dipole trapping beams [68]. The current in the coils is then linearly ramped down over 60ms until the coils are completely turned off. The repumper beams are completely extinguished 1ms before the end of the magnetic field ramp while the MOT beams are still on so all the atoms in the trap fall into the lowest F=1 hyperfine state. The MOT beams are extinguished at the same time the magnetic field turns completely off.

Both dipole trapping beams remain on at full power (37 W total in the crossing region) the entire time the UHV MOT loads from the VC MOT as well as during the cooling and compression stages of the UHV MOT. We found that waiting to turn on the dipole trapping beams until the onset of cooling and compression of the UHV MOT appears to have no effect on the final loading characteristics of the dipole trap. The dipole trapping beams are positioned so they cross at their beam waists, and so the crossing point occurs at the center of the UHV MOT. Alignment of the dipole trapping beams in other arrangements leads to decreased efficiency of loading into the dipole trap.

4.4 Dipole trap characterization

Image analysis

The dipole trap is observed and measured with a 1:1 imaging lens system onto a CCD camera using absorption imaging techniques. Four pictures are taken to subtract any background light and normalize the images. A probe beam tuned near the $F = 2 \rightarrow F' = 3$ transition passes through the dipole trap and the shadow cast by the atoms is imaged onto the camera. Because all the atoms in the trap have been prepared in the $F = 1$ hyperfine state, repumper light must be present for atoms to absorb any of the probe light. The first picture, I_1, is taken in the presence of the probe beam but with no repumper light so that no atoms are observed. Because no repumper light is present, the atoms do not interact with the probe

light and remain in the dipole trap. The second picture, I_2, is taken in the presence of both probe and repumper light so atoms may be observed. At this point, the atoms have been expelled from the dipole trap due to interaction with the near-resonant probe beam. The next picture, I_3, contains no probe or repumper light and serves as an image of the background light to be subtracted from the first picture. The final picture, I_4, is taken in the presence of repumper light but with no probe light. This picture serves as an image of the background light to be subtracted from the second picture, since the repumper light is considered part of the background. The final image used for analysis is a normalized image

$$\frac{I}{I_0} = \frac{I_2 - I_4}{I_1 - I_3}. \tag{4.3}$$

Images I_1 and I_3 have exposure times of 30 μs each, while, due to technical limitations of the camera, images I_2 and I_4 have exposure times of 150 ms each. However, the probe beam is on for the same amount of time in both I_1 and I_2 (10 μs), and the repumper beam is on for the same amount of time in both I_2 and I_4 (10 μs).

The dipole trapping beams are aligned through the UHV MOT on a plane parallel to the optical table. Likewise, the beams cross to form a dipole trap that is aligned on an orthogonal plane to the camera. These precautions allow us to define a convenient and consistent frame of reference when performing analysis of the images. Technical graphing and data analysis is all performed using the IGOR Pro software from Wavemetrics. When discussing images analyzed in IGOR Pro, I will use the following convention: the horizontal axis is the x-axis, the vertical axis is the y-axis, and the z-axis refers to the axis that points out of the page. The result is that the long axis of the crossed dipole trap is aligned along the x-axis, and the short axes of the trap are aligned along the y- and z-axes. Due to the quasi-rotational symmetry of the crossed dipole trap, all quantities either measured or calculated along the y-axis are taken to be the same along the z-axis. In the discussion of trap frequencies in Section 4.4 for the crossed dipole trap, we explicitly take into account the lack of perfect

rotational symmetry, and it is shown that using the same values for quantities along the z-axis as for quantities along the y-axis is reasonable under many circumstances. In Section 4.5, however, this assumption breaks down when discussing the effects the uncrossed regions of the dipole trapping beams have on the crossed region, and the lack of rotational symmetry must be taken into account.

The dipole trap is a 3D harmonic potential, $U(x, y, z)$, of the form

$$U(x, y, z) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2). \tag{4.4}$$

As a result of the Boltzmann equation, the density distribution, $n(x, y, z)$, of the thermal cloud is

$$n(x, y, z) = n_0 \exp\left(-\frac{U(x, y, z)}{k_B T}\right), \tag{4.5}$$

which leads to a Gaussian spatial profile in all directions of the trap:

$$n(x, y, z) = n_0 \exp\left(-\frac{x^2}{2(\Delta x)^2}\right) \exp\left(-\frac{y^2}{2(\Delta y)^2}\right) \exp\left(-\frac{z^2}{2(\Delta z)^2}\right), \tag{4.6}$$

where ω_i are the trap frequencies, and (Δi) are the widths of the Gaussian peaks in each direction [53]. The simplest analysis of the crossed dipole trap can be performed by taking horizontal and vertical cross-sections at the center of the trap. Each camera pixel is a point in the plot, and performing a Gaussian fit to each cross-section provides all the information necessary to find all relevant quantities. Along the y- and z-axes, a single Gaussian fit is sufficient. Along the x-axis, however, two regions of interest exist: a central region with a large optical depth and an outer region with a low but non-zero optical depth. The central region with a large optical depth corresponds to the region where the crossed dipole beams overlap. The outer region corresponds to the region where the beams do not overlap but still behave as single beam dipole traps. A single Gaussian profile along the x-axis is often insufficient to get accurate measurements of the dipole trap if the outer region contains any

number of atoms. In this case, a double Gaussian profile is required to isolate the relevant properties of the central, high optical depth region from the irrelevant properties of the outer, low optical depth region as in Figures 4.2 and 4.3.

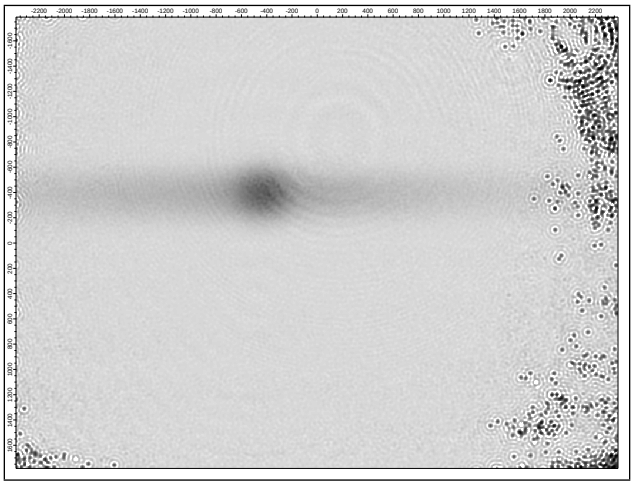

Figure 4.2: This is a normalized image of a crossed dipole trap with atoms in the single beam horizontal extensions that requires a double Gaussian fit.

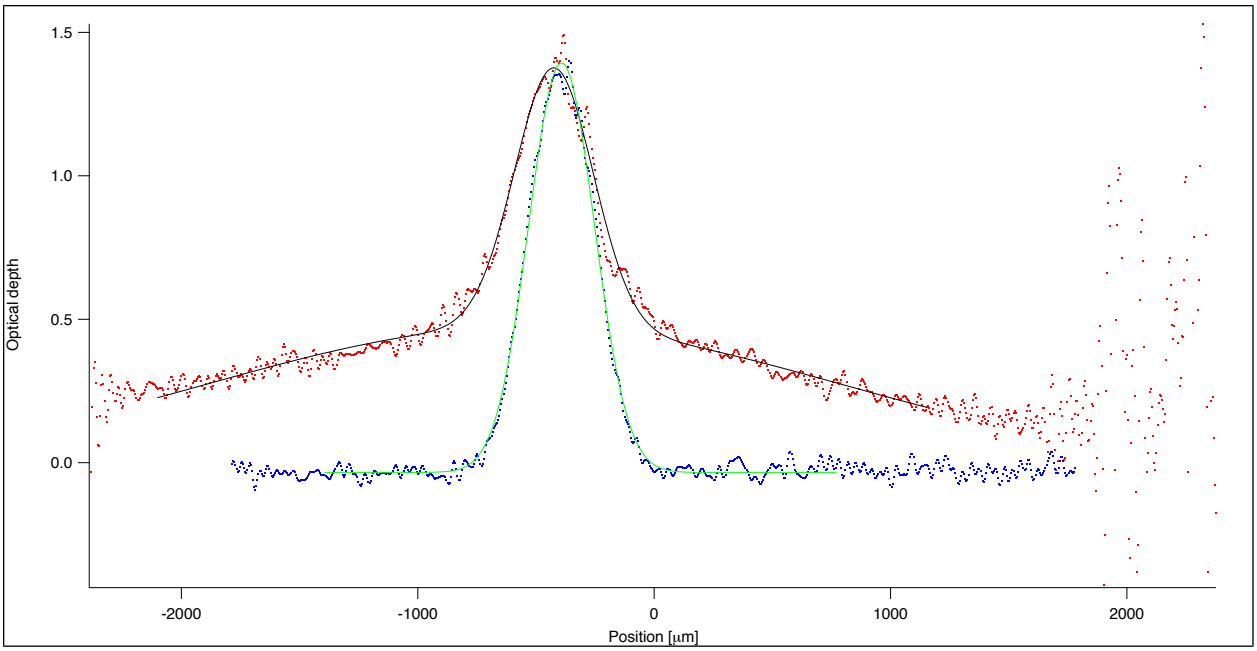

Figure 4.3: These are vertical and horizontal cross sections of the image in Figure 4.2. The horizontal cross section does not behave as a single Gaussian due to atoms that are trapped outside the crossed region of the beams. The vertical cross section is fit with a single Gaussian, but the horizontal cross section is fit with a double Gaussian of the form $f(x) = a + be^{-[(x-c)/d]^2} + ue^{-[(x-v)/w]^2}$.

Number of atoms

The quantity of interest for all calculations is the optical depth of the image, OD, given by

$$I = I_0 e^{-OD} \tag{4.7}$$

according to Beer's Law. Alternatively, one can express equation (4.7) in terms of the atom column density, $n(x, y)$, and the absorption cross section, σ:

$$I = I_0 e^{-n(x,y)\sigma}. \tag{4.8}$$

The absorption cross section is given by

$$\sigma = \frac{\sigma_0}{1 + 4(\Delta/\Gamma)^2 + (I/I_{sat})}, \tag{4.9}$$

where

$$\sigma_0 = \frac{\hbar\omega\Gamma}{2I_{sat}} \tag{4.10}$$

is the on resonance saturation parameter, Δ is the detuning of the probe laser from the atomic resonance, Γ is the decay rate of the transition, and I_{sat} is the saturation intensity. Combining equations (4.7) and (4.8) yields $OD = n(x, y)\sigma$.

The column density $n(x, y)$ is defined as the 3D density $n(x, y, z)$ integrated over the path of the probe beam through the trap so that

$$n(x, y) = \int n(x, y, z)\,\mathrm{d}z = n(x, y, z)\pi^{1/2}(\Delta z). \tag{4.11}$$

Integrating the density over the the volume of the dipole trap gives the number of atoms, N, in the trap:

$$N = \iiint n(x, y, z)\,\mathrm{d}x\mathrm{d}y\mathrm{d}z = n(x, y, z)\pi^{3/2}(\Delta x)(\Delta y)(\Delta z). \tag{4.12}$$

Up to 5.6×10^6 atoms have been observed in the fully loaded dipole trap, as seen in Figure 4.4.

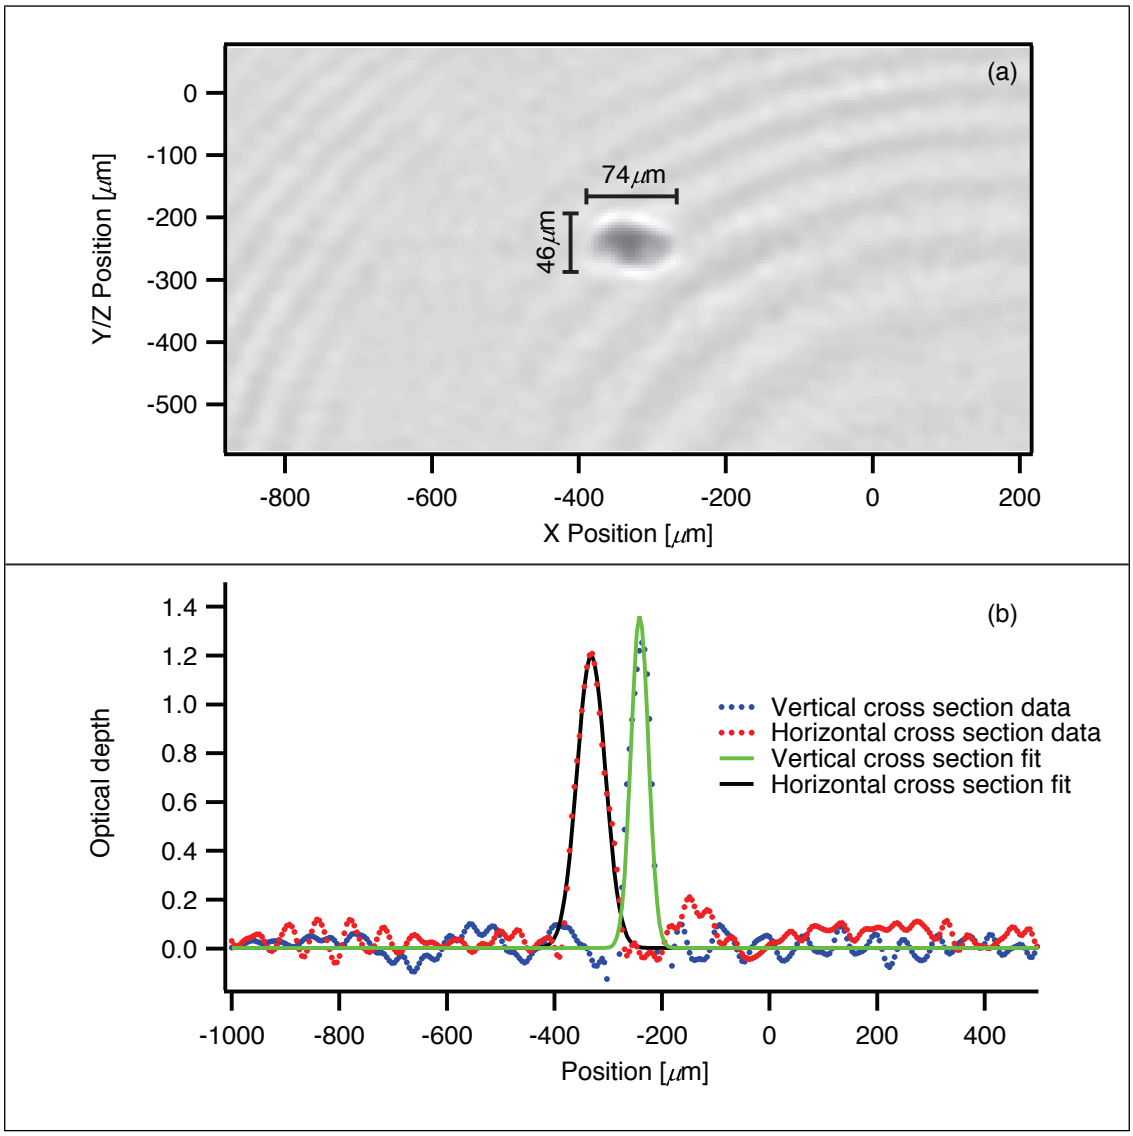

Figure 4.4: (a) CCD image of the fully loaded crossed dipole trap, with background subtracted and normalized. (b) Horizontal and vertical cross sections of the center of the dipole trap along with Gaussian fits. This image shows a trap containing 5.6×10^6 atoms with a density $n(x, y, z) = 52$ atoms/μm^3.

Temperature

In equation (4.6), the widths of the Gaussian peaks in each direction are given by

$$(\Delta i) = \omega_i^{-1}\sqrt{k_B T/m}. \tag{4.13}$$

Therefore, the temperature of the cloud is related to its spatial properties by

$$T = \frac{m}{k_B}(\Delta i)^2\omega_i^2. \tag{4.14}$$

Our imaging system does not have the resolving power to take an absorption image of the in situ dipole trap. Instead, the dipole trapping beams are extinguished, and the cloud is permitted to expand ballistically for some finite amount of time before an absorption image is taken. The resulting cloud after a thermal expansion provides information on the velocity distribution of atoms in the cloud, and thus the temperature.

With knowledge of the trap frequencies, it is not necessary to take multiple pictures of the trap after different expansion times to calculate the temperature. A simple scaling law exists between the size of the expanded cloud, (Δi), and the original size of the in situ cloud, $(\Delta i)_0$, dependent only on the in situ trap frequencies and the time allowed for expansion:

$$\frac{(\Delta i)}{(\Delta i)_0} = \sqrt{1 + (\omega_i t)^2}. \tag{4.15}$$

Now that the initial size of the cloud is known, the width of the cloud with temperature T at time t after all trapping potentials have been turned off can be expressed as

$$(\Delta i) = \sqrt{\frac{k_B T}{m}t^2 + (\Delta i)_0^2}. \tag{4.16}$$

Solving for T gives

$$T = \frac{m}{k_B}\left(\frac{(\Delta i)^2 - (\Delta i)_0^2}{t^2}\right). \tag{4.17}$$

The average temperature of the atoms in the fully loaded crossed dipole trap is 290 μK. An image of an expanded cloud used to calculate this temperature can be seen in Figure 4.5.

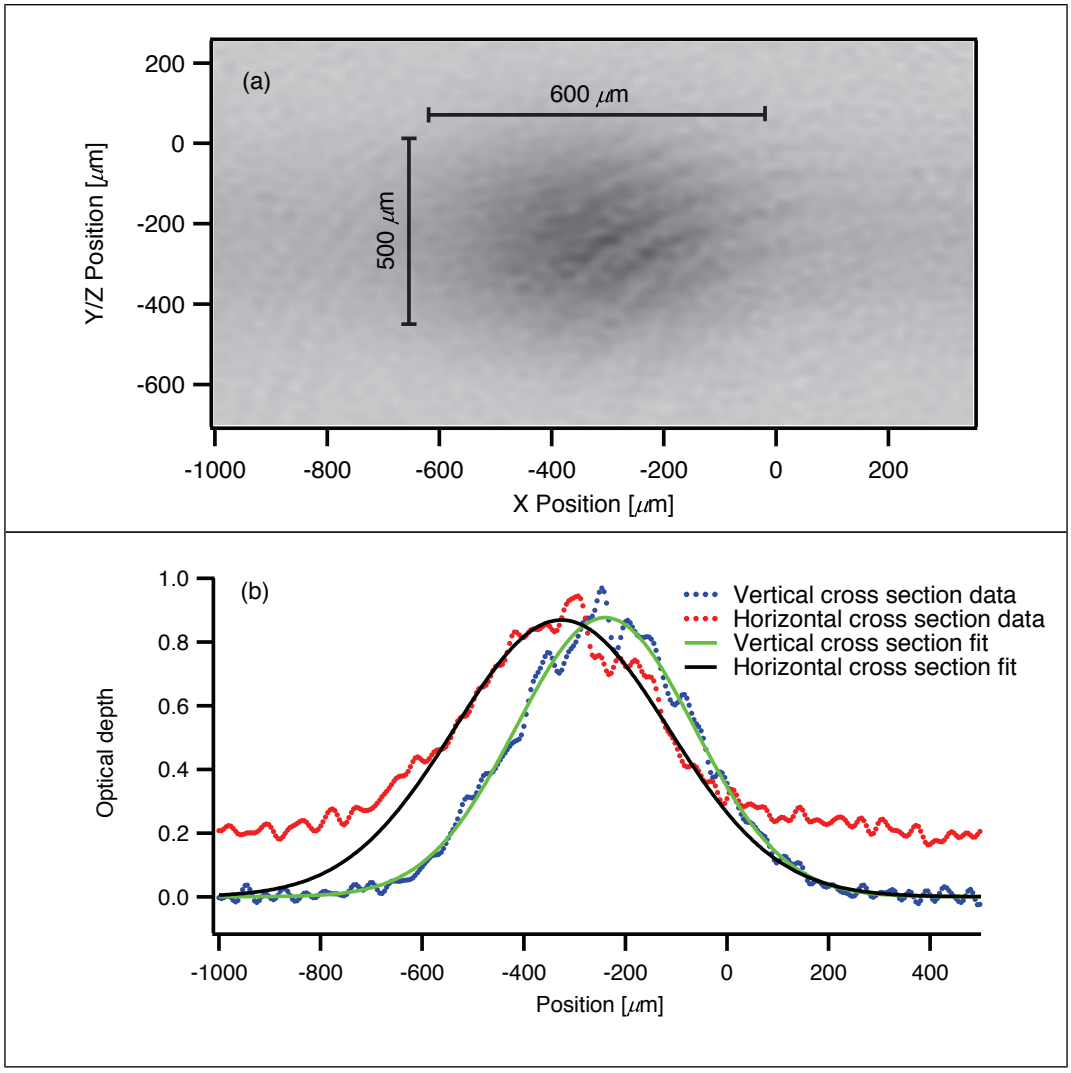

Figure 4.5: (a) CCD image of the fully loaded crossed dipole trap 1 ms after the trapping beams have been turned off, with background subtracted and normalized. (b) Horizontal and vertical cross sections of the center of the dipole trap along with Gaussian fits. The horizontal cross section is actually a double Gaussian, but the fit only shows the Gaussian for the central crossed region with the Gaussian from the single beam extensions removed. This image shows a 290 μK trap.

Trap depth

The dipole beam intensity $I(x, y, z)$ can be expressed as

$$I(x, y, z) = \frac{2P}{\pi w^2(x)} e^{-2\frac{y^2+z^2}{w^2(x)}}, \tag{4.18}$$

where $w(x)$ is the $1/e^2$ beam radius, and P is the power in the beam. An expression for the dipole potential of the ground state of ^{87}Rb in the presence a linearly polarized beam with a large detuning relative to the excited state hyperfine splitting is given as [53]

$$\begin{aligned} U(x, y, z) = \frac{3\pi c^2}{2} \Bigg[& \frac{1}{3\omega_{01}^3} \left(\frac{\Gamma_1}{\omega_{01} - \omega_{laser}} + \frac{\Gamma_1}{\omega_{01} + \omega_{laser}} \right) \\ & + \frac{2}{3\omega_{02}^3} \left(\frac{\Gamma_2}{\omega_{02} - \omega_{laser}} + \frac{\Gamma_2}{\omega_{02} + \omega_{laser}} \right) \Bigg] I(x, y, z). \end{aligned} \tag{4.19}$$

Here, Γ_1 is the natural line width of the D_1 line, ω_{01} is the frequency of the D_1 line, Γ_2 is the natural line width of the D_2 line, ω_{02} is the frequency of the D_2 line, and ω_{laser} is the frequency of the trapping laser beam. The trap depth, U_0, for a single beam dipole trap is simply $U(0, 0, 0)$.

Trap frequencies

An accurate calculation of the trap temperature requires accurate knowledge of the trap frequencies. In a single beam dipole trap where the thermal energy $k_B T$ of the atoms is much smaller than the potential depth U_0, the optical potential seen by the atoms can be well approximated by the harmonic oscillator potential

$$U(x, y, z) \simeq -U_0 \left(1 - \left(\frac{x}{x_R} \right)^2 - \frac{2(y^2 + z^2)}{w_0^2} \right), \tag{4.20}$$

where w_0 is the beam waist, and

$$x_R = \frac{\pi w_0^2}{\lambda} \tag{4.21}$$

is the Rayleigh length. Comparing equations (4.4) and (4.20), it can be seen that the trapping frequencies for a single beam dipole trap are given by

$$\omega_x = \sqrt{\frac{2U_0}{mx_R^2}} \tag{4.22a}$$

$$\omega_{y,z} = \sqrt{\frac{4U_0}{mw_0^2}}. \tag{4.22b}$$

For beam 1, $P = 22$ W and $w_0 = 38\,\mu$m, which gives a trap depth of $U_0/k_B = 1.5$ mK, a trap frequency of $\omega_x = 2\pi \times 19.8$ Hz in the axial direction, and a trap frequency of $\omega_{y,z} = 2\pi \times 3.1$ kHz in the radial direction. For beam 2, $P = 15$ W and $w_0 = 38\,\mu$m, which gives a trap depth of $U_0/k_B = 1.0$ mK, a trap frequency of $\omega_x = 2\pi \times 16.4$ Hz in the axial direction, and a trap frequency of $\omega_{y,z} = 2\pi \times 2.600$ kHz in the radial direction.

For a crossed beam dipole trap with beams of equal waists propagating along the xz-plane and intersecting at their focal points, the optical potential seen by the atoms in the crossed region of the dipole trapping beams looks like

$$U(x,y,z) \simeq -U_0' \left(1 - 2\frac{x^2}{w_x^2} - 2\frac{y^2}{w_0^2} - 2\frac{z^2}{w_z^2}\right), \tag{4.23}$$

where U_0' is the sum of the trap depths of both single beam potentials, $w_x^2 = \left(\sin^2\phi/w_0^2 + \cos^2\phi/2x_R^2\right)^{-1}$, $w_z^2 = \left(\cos^2\phi/w_0^2 + \sin^2\phi/2x_R^2\right)^{-1}$, and $\pm\phi$ is the angle between the beams and the x-axis. Comparing equations (4.4) and (4.23), it can be seen that the trapping frequencies for a crossed beam dipole trap are given by

$$\omega_x = \sqrt{\frac{4U_0'}{mw_x^2}} \tag{4.24a}$$

$$\omega_y = \sqrt{\frac{4U_0'}{mw_0^2}} \tag{4.24b}$$

$$\omega_z = \sqrt{\frac{4U_0'}{mw_z^2}} \tag{4.24c}$$

For a crossed beam configuration with 22 W in beam 1 and 15 W in beam 2, $w_0 = 38\,\mu$m for both beams, and $\phi = \pm 16°$, $U_0'/k_B = 2.5$ mK, $\omega_x = 2\pi \times 1.126$ kHz, $\omega_y = 2\pi \times 4.084$ kHz, and $\omega_z = 2\pi \times 3.926$ kHz. Here we have not implicitly assumed rotational symmetry of the dipole trap, but it can be seen that the trap frequencies in the y and z directions differ by only 4%, making rotational symmetry of the crossed dipole trap a reasonable assumption for the discussions so far.

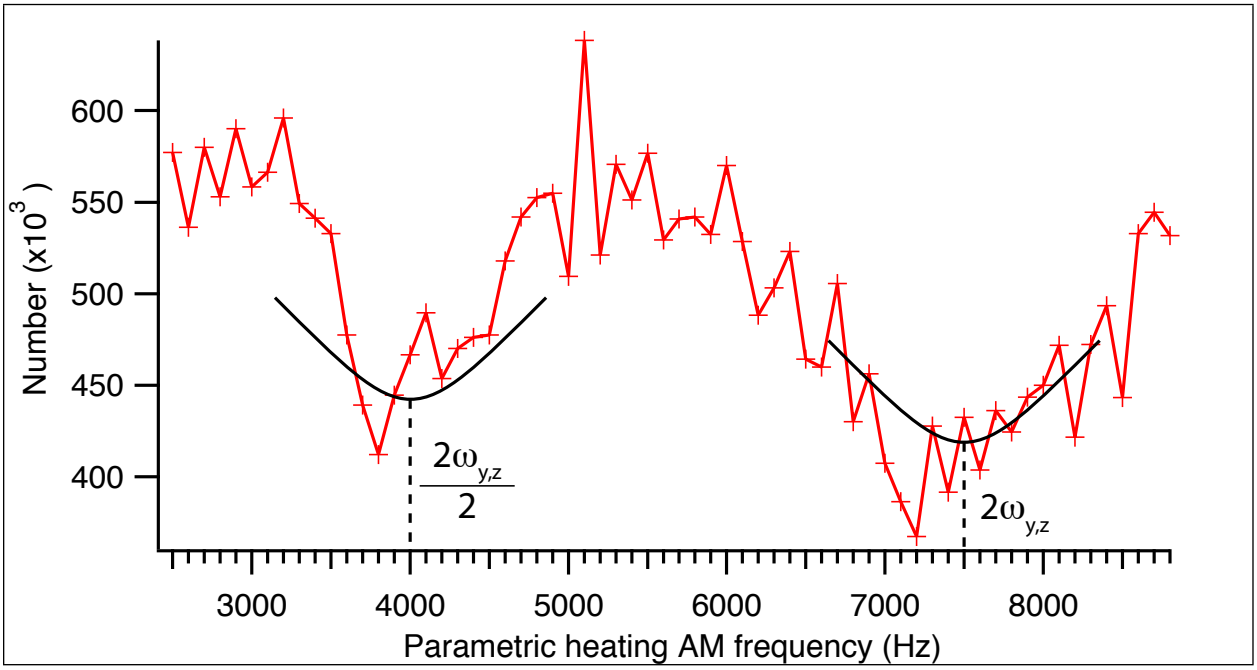

Figure 4.6: Parametric heating experiment showing the crossed trap frequencies in the tightly confined direction. A trap frequency is observed at $2\omega_{y,z} = 7.5$ kHz, as well as a subharmonic at $2\omega_{y,z}/2 = 4.0$ kHz, in reasonable agreement with the calculations based on the beam waists.

Alternatively, the trap frequencies can be measured experimentally using parametric heating [69, 70, 71, 72]. Trap heating can be induced by modulating the trapping laser light intensities at a frequency ω_{mod} such that

$$\omega_{mod} = \frac{2\omega_i}{n},\ n = 1, 2, 3 \ldots . \tag{4.25}$$

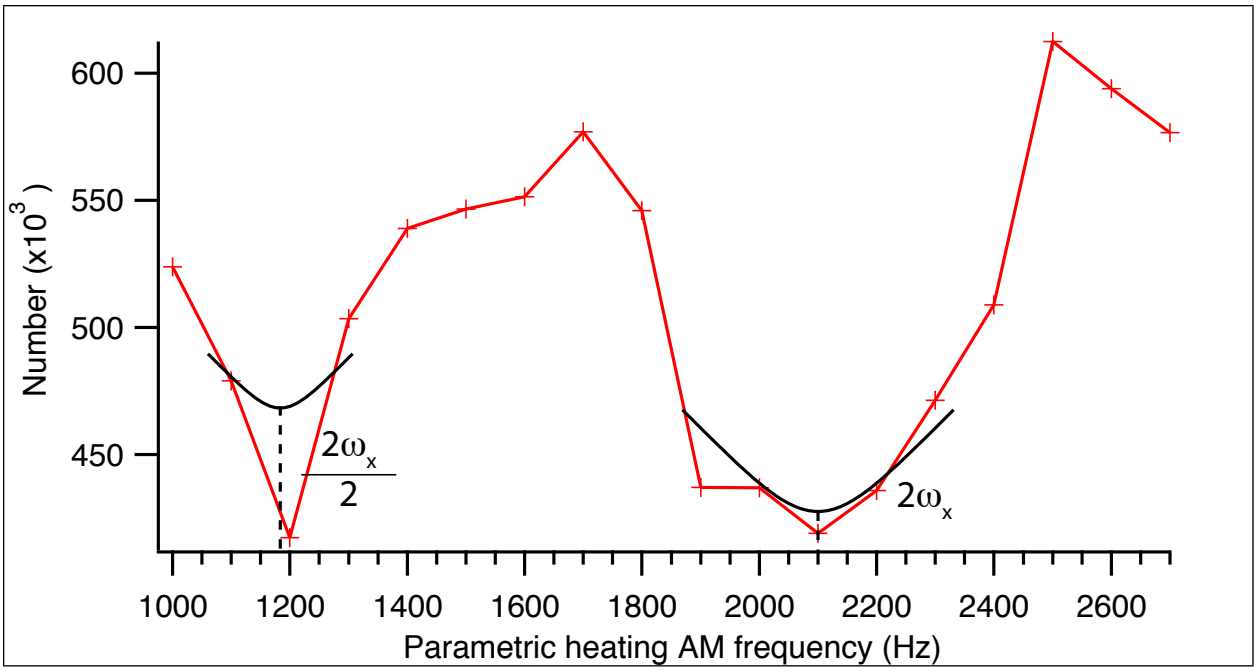

Figure 4.7: Parametric heating experiment showing the crossed trap frequencies in the loosely confined direction. A trap frequency is observed at $2\omega_x = 2.1$ kHz, as well as a subharmonic at $2\omega_x/2 = 1.18$ kHz, in reasonable agreement with the calculations based on the beam waists.

In other words, by scanning the amplitude modulation frequency of the dipole trapping beams, it is possible to observe a decrease in the number of trapped atoms when the modulation frequency is equal to twice the trap frequency (and its subharmonics). Experiments performed in this way, shown in Figure 4.6 and Figure 4.7, show trap frequencies at $2\omega_x = 2\pi \times 2.1$ kHz and $2\omega_{y,z} = 2\pi \times 7.5$ kHz for the crossed beam trap, in good agreement with the trap frequencies calculated from the beam waists. Because ω_y is very close to ω_z, it was not possible to resolve the two individual frequencies. The same experiment was performed on the two single beams. For beam 1, seen in Figure 4.8, a trap frequency $2\omega_{y,z} = 2\pi \times 4.9$ kHz was observed. This does not quite agree with the trap frequency calculated from the beam waist, but the subharmonic observed at $2\omega_{y,z}/2 = 2\pi \times 3.0$ kHz is in better agreement. For beam 2, seen in Figure 4.9, a trap frequency $2\omega_{y,z} = 2\pi \times 4.9$ kHz was observed. Again, this does not quite match the trap frequency calculated from the

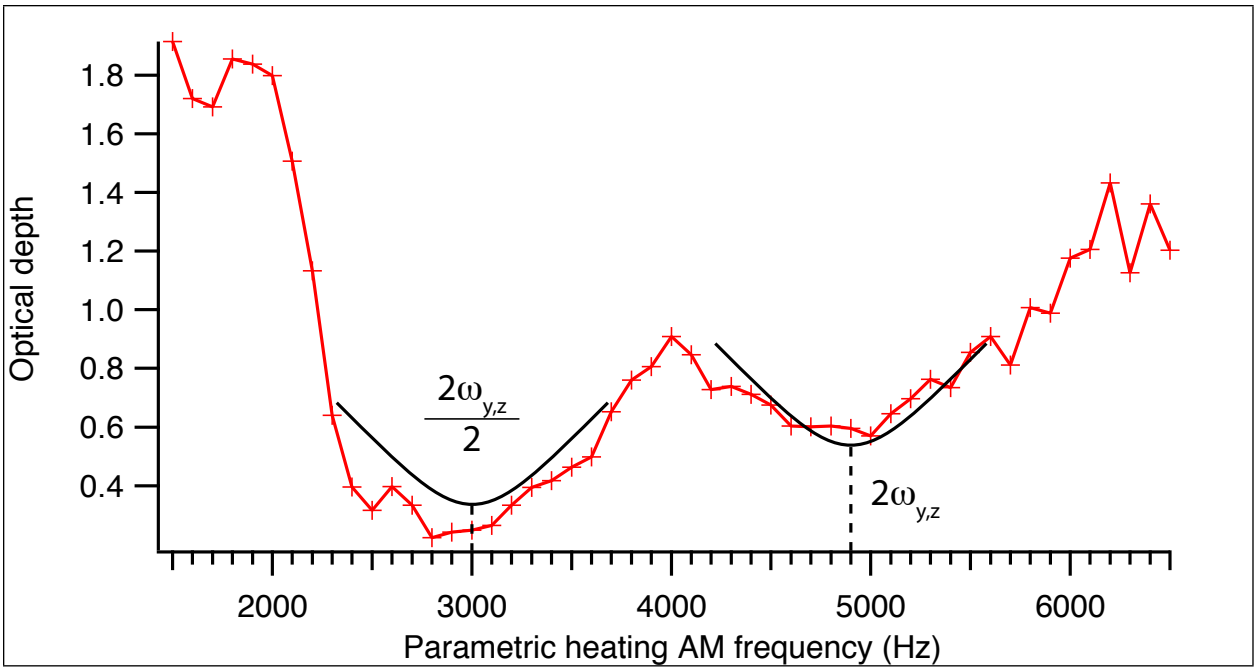

Figure 4.8: Parametric heating experiment showing the single beam trap frequencies for beam 1. A trap frequency is observed at $2\omega_x = 4.9$ kHz, as well as a subharmonic at $2\omega_x/2 = 3.0$ kHz.

beam waist, but the subharmonic observed at $2\omega_{y,z}/2 = 2\pi \times 2.6$ kHz is in better agreement. The low trap frequency ω_x was not observed in the case of the single beams because the time required to modulate the beam intensity enough times to observe a depletion in the trap number was too long relative to the trap lifetime when these measurements were made, at which time the vacuum system was failing, leading to higher background pressures in the UHV chamber.

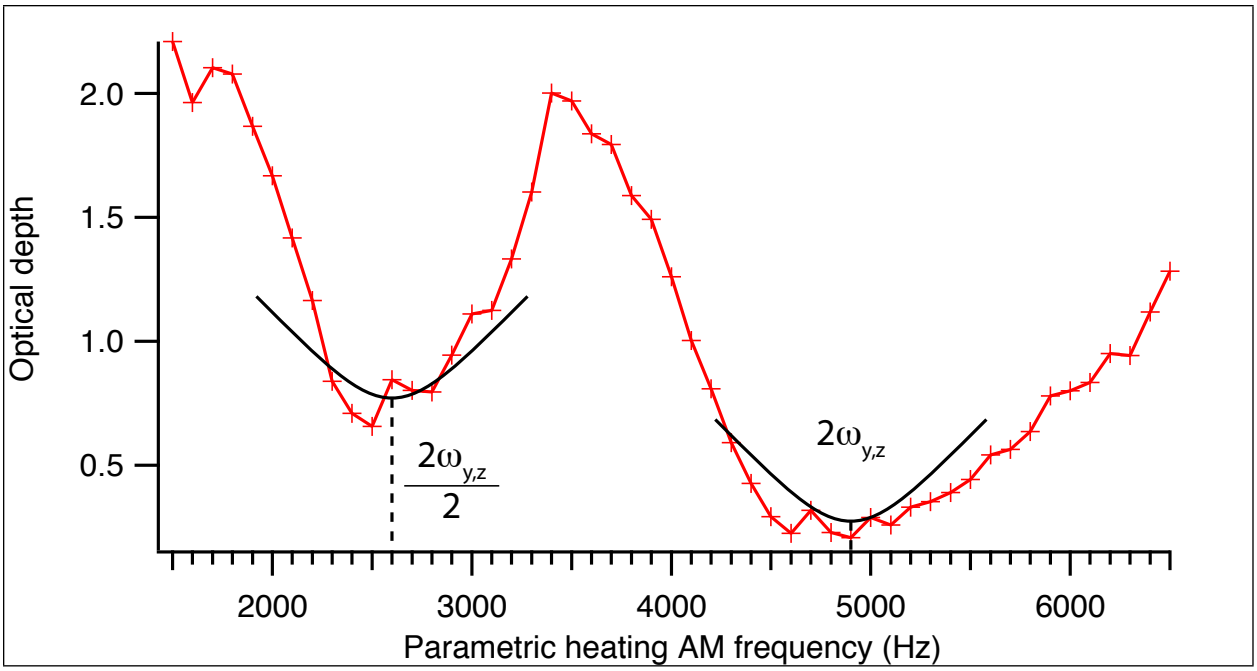

Figure 4.9: Parametric heating experiment showing the single beam trap frequencies for beam 2. A trap frequency is observed at $2\omega_x = 4.9$ kHz, as well as a subharmonic at $2\omega_x/2 = 3.0$ kHz.

Trap lifetime

In a single beam dipole trap with atoms prepared in the lower hyperfine ground state, the primary loss mechanism arises from single particle loss due to collisions with the background gas. This leads to a purely exponential decay in the number of trapped atoms. As long as the trap is prepared in the lower hyperfine ground state, hyperfine-changing collisions between two trapped ground state atoms, which would lead to a rapid non-exponential decay, are negligible [73, 74, 53]. The situation is more complex in the crossed beam dipole trap, however. Due to the significantly higher density in the crossed beam trap, inelastic three-body collisions can play a major role in the decreased trap lifetime. In a three-body collision process, two atoms bond together to form a dimer, and the released energy is absorbed by the third atom. As a result, all three atoms are ejected from the trap as an additional source

of trap loss. This can be observed as a significantly reduced trap lifetime in the crossed beam trap.

For both the single beam and crossed beam dipole traps, the traps are fully loaded as normal, and the atoms are held in the traps for varying amounts of time before being released and imaged. Number measurements are made for each value of the holding time and plotted to get a measurement of the trap lifetime. The single beam traps are observed to have lifetimes in excess of 30 s with very little trap loss, a sign of very low background pressures in the system ($8.8 \times 10^{-}11$ torr based on lifetime measurements of the UHV MOT). The crossed beam trap lifetime, shown in Figure 4.10 for a full power trap, shows a significant loss of atoms at early times due to three-body collisional losses. The lifetime was fit with a double exponential, which shows a short initial lifetime $\tau_{short} = 0.22$ s and a longer lifetime $\tau_{long} = 1.5$ s as the temperature and density decrease due to free evaporation. As one might expect, loading into a crossed dipole trap with lower beam powers increases the trap lifetime due to a reduced maximum density in the trap, as seen in Figure 4.11 and Figure 4.12.

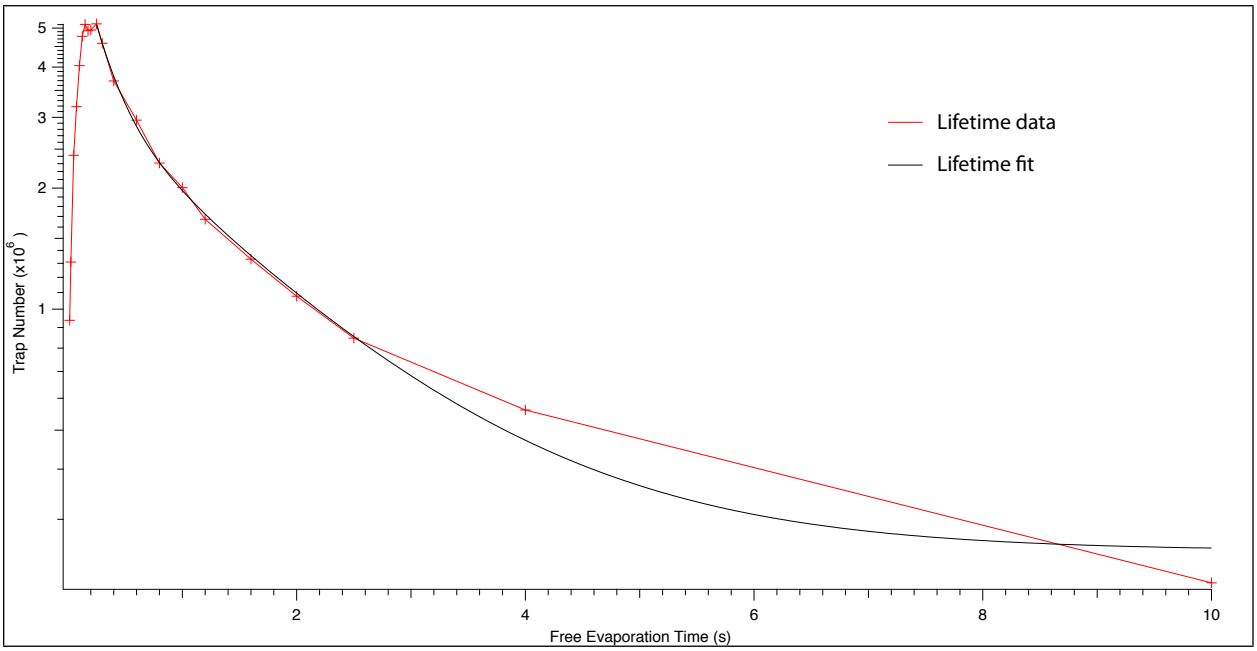

Figure 4.10: A plot showing the lifetime of a full power crossed dipole trap. The red curve is the number of atoms in the trap measured after holding the atoms in the trap for different amounts of time. The black curve is a double exponential fit with $\tau_{short} = 0.22$ s and $\tau_{long} = 1.5$ s. This demonstrates a significantly shortened lifetime compared to the single beam traps due to three-body collisional losses.

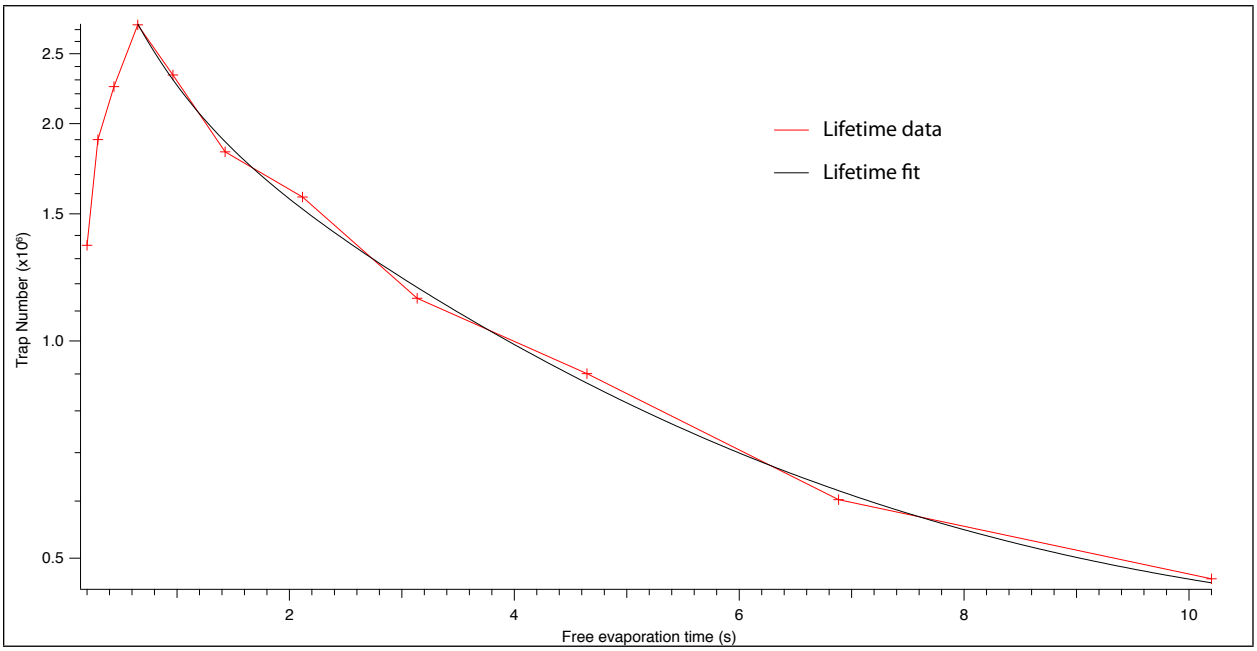

Figure 4.11: A plot showing the lifetime of a crossed dipole trap with 18/25 the full power in the beams. The red curve is the number of atoms in the trap measured after holding the atoms in the trap for different amounts of time. The black curve is a double exponential fit with $\tau_{short} = 0.48$ s and $\tau_{long} = 3.1$ s. The lifetime is slightly longer than a full power crossed dipole trap due to reduced densities.

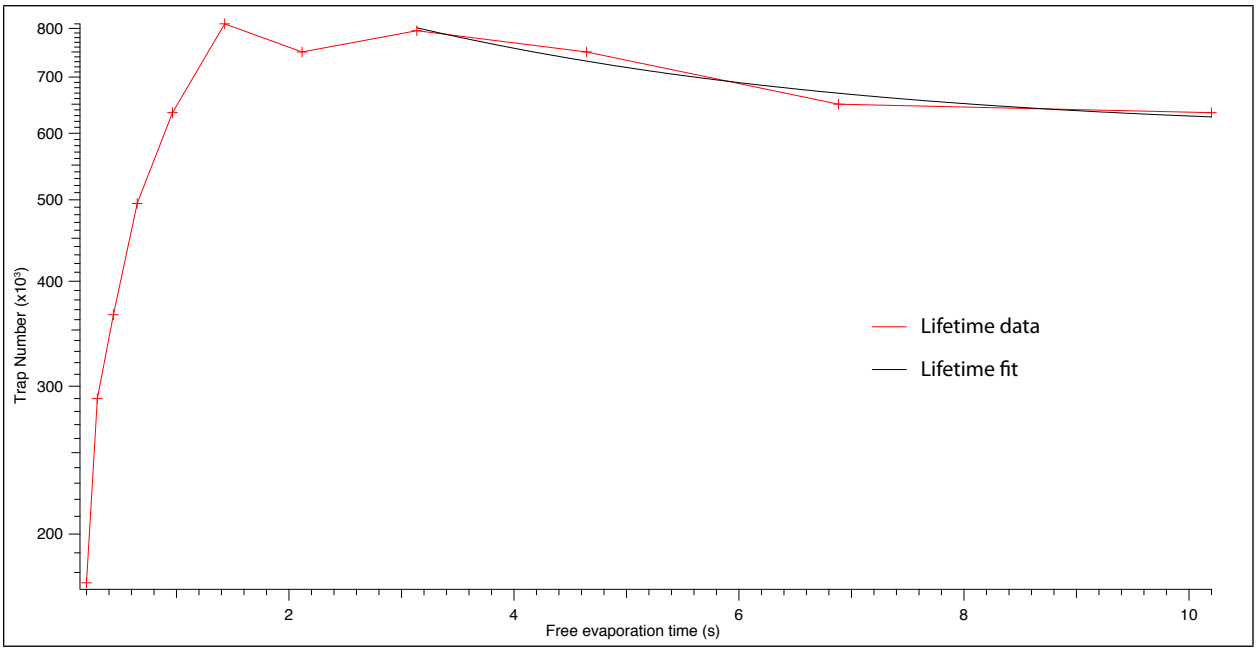

Figure 4.12: A plot showing the lifetime of a crossed dipole trap with 12/25 the full power in the beams. The red curve is the number of atoms in the trap measured after holding the atoms in the trap for different amounts of time. The black curve is now a single exponential fit with $\tau = 5.7$ s. The lifetime is even longer with significantly reduced three-body collisional losses, making a single exponential fit sufficient to measure the lifetime.

4.5 Evaporation and the "wings" problem

Plan for evaporation

While inelastic scattering processes such as two-body collisions between atoms in different internal states or three-body collisions are undesired because they lead to trap heating and loss, elastic scattering processes that only result in a redistribution of kinetic energy are essential during evaporative cooling. Once the fully loaded crossed beam trap reaches densities below the three-body collision threshold, elastic collisions inside the trap become dominant which allow the atoms in the trap to thermalize. The elastic collision process produces a few atoms with energies greater than $k_B T$ on the high-velocity tail of the Boltzmann distribution that escape the trap. The remaining atoms in the trap rethermalize to a lower average temperature. If the trap properties are held constant – so called "free evaporation" – this eventually leads to a steady-state temperature in the trap when the process stagnates. This typically occurs at a maximum trap depth-to-temperature ratio $\eta = U_0/k_B T \approx 14$. If evaporation is to proceed further, the trap depth must be actively reduced – so called "forced evaporation".

The main quantity of interest during evaporation to BEC is the phase space density, ρ, where

$$\rho = n\lambda_{dB}^3, \tag{4.26}$$

n is the density, and λ_{dB} is the thermal deBroglie wavelength

$$\lambda_{dB} = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}. \tag{4.27}$$

The transition to BEC occurs when $\rho \gtrsim 2.6$. Starting with the initial conditions in our trap, free evaporation alone is insufficient to reach such high phase space densities, so a forced evaporation technique must be applied. The typical forced evaporation scheme involves a constant η ramp of the trap depth by decreasing the trapping beam intensities. A major

benefit of a crossed dipole trap versus a single beam trap when using evaporative cooling comes from the significantly higher trap frequencies associated with the tighter confinement of the crossed beam trap. Evaporation in a crossed beam optical dipole traps proceeds quickly at first when large trapping frequencies provide high scattering rates, but as the trapping power decreases, evaporation stalls because the trapping frequencies, and therefore the scattering rates, become too small for the atoms to rethermalize quickly. With our initial trapping conditions of 5.6×10^6 atoms, an average temperature of $290\,\mu$K, and a trap depth of $U_0'/k_B = 2.5$ mK, simulations performed based on a model developed in [75] estimate that evaporation to the BEC phase transition should be possible in about 2 s. Figures 4.13, 4.14, and 4.15 show that at $\rho = 2.6$, we should have a condensate with $\approx 3.5 \times 10^5$ atoms at a temperature of 120 nK.

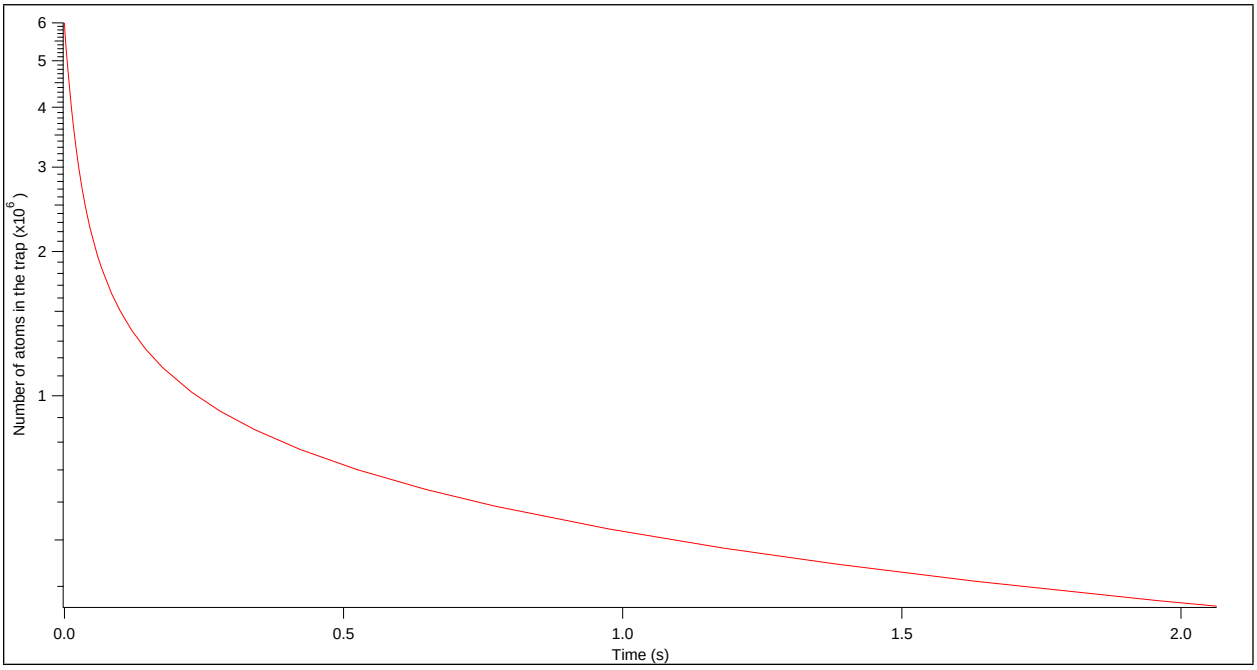

Figure 4.13: This is a plot showing the number of atoms in the crossed beam dipole trap during a simulated evaporation ramp that reaches the BEC phase transition in 2 s. This shows 3.5×10^5 atoms in the BEC.

The choice of η for the evaporation ramp must be a compromise. If η is too large, the rethermalization time for each step of the ramp becomes too long relative to the background gas limited trap lifetime, especially at low trap depths where the trap frequency becomes

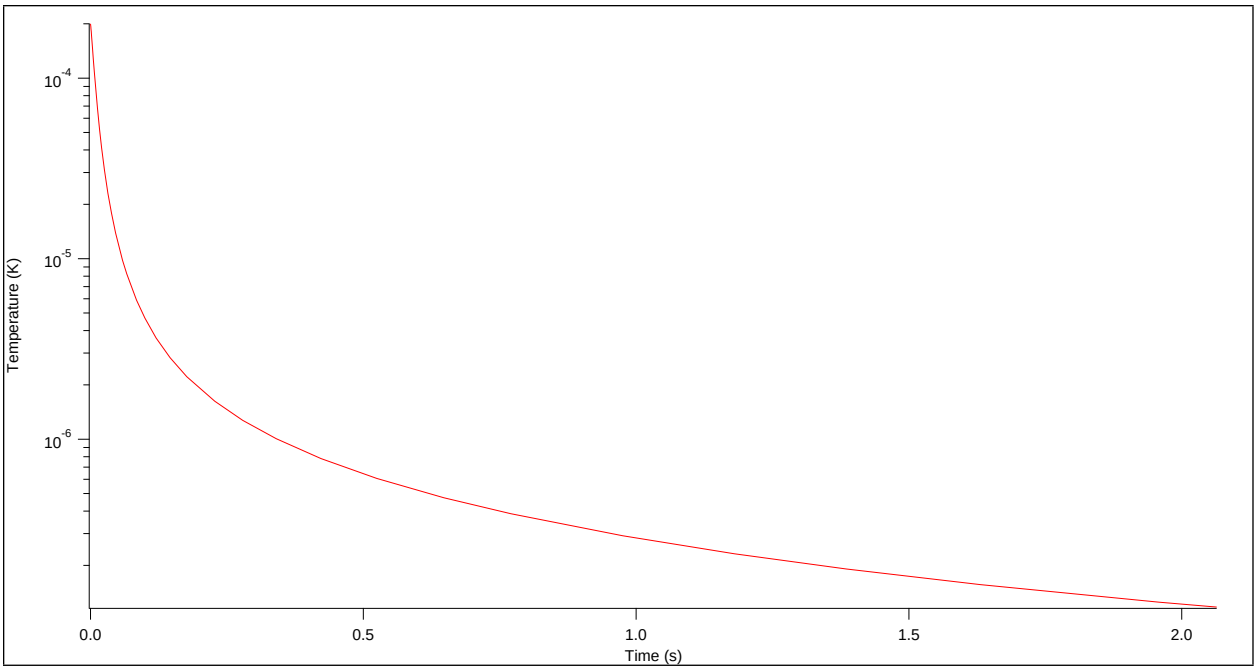

Figure 4.14: This is a plot showing the temperature in the crossed beam dipole trap during a simulated evaporation ramp that reaches the BEC phase transition in 2 s. This shows a final temperature of 120 nK.

prohibitively small. If η is too small, however, a large fraction of the atoms will be ejected from the trap at each step because they have not had a chance to rethermalize to a lower average temperature. Dumping too many atoms out of the trap at once in such a way will leave too few atoms in the trap to observe a BEC.

The original choice for this experiment was to employ a moderate $\eta = 8$ ramp. A reasonable approximation to this smooth ramp is to suddenly drop the trap depth to $\eta = 7$ and then wait at that trap depth until the trap rethermalizes to $\eta = 8$ and repeat the process. At these values, the trap depth will only drop by 12.5% at each step. As the trap depth gets lower, the rethermalization time to reach $\eta = 8$ after each drop will get longer, approximating a smooth curve with constant η. At low trap depths, when U_0 approaches the magnitude of the gravitational potential felt by the atoms, a compensating field is turned on in the electromagnetic coils to keep one third of the atoms from falling out of the trap.

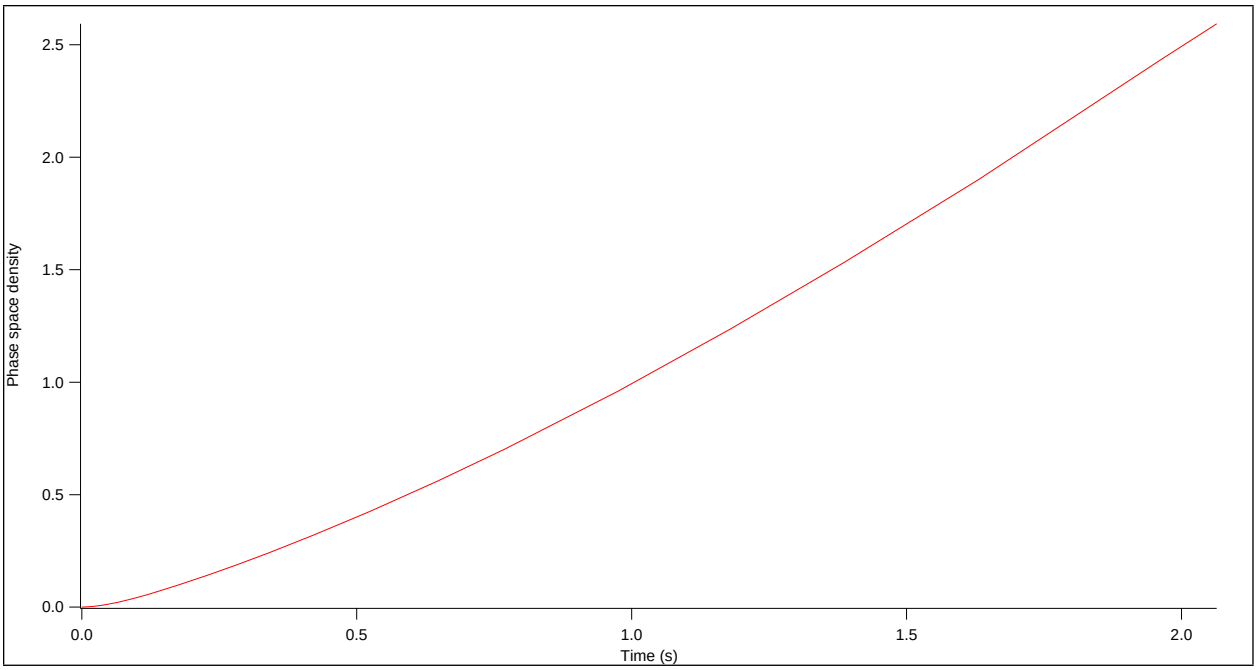

Figure 4.15: This is a plot showing the phase space density in the crossed beam dipole trap during a simulated evaporation ramp that reaches the BEC phase transition in 2 s. This shows that the phase space density does reach the expected value of $\rho = 2.6$ where the phase transition occurs.

Without a levitating field, gravity can tilt the dipole trap potential until it no longer forms a bound potential, causing all the atoms to leave the trap. By introducing a levitating field, the $F = 1$, $m_f = \pm 1$ substates will be Zeeman shifted in opposite directions. Therefore, depending on the direction of the levitating field, one of these substates will still be confined in a bound potential, while the other substate and the $F = 1$, $m_f = 0$ substate will fall out of the trap.

If forced evaporation proceeds as expected, a continuous increase in the phase space density should be observed. Contrary to expectations, however, no increase in phase space density was ever observed for ramps ranging from $\eta = 6$ to $\eta = 9$. The best results were observed for a $\eta = 10$ ramp, which can be seen in Figures 4.16, 4.17, and 4.18. In this case, there was an observed increase in phase space density, but only by a factor of 5 in 1.5 s. The final phase space density data point is still six orders of magnitude below the required

phase space density for a BEC, and if it continues to increase at the same rate, it would take $\approx 3 \times 10^5$ seconds to reach BEC.

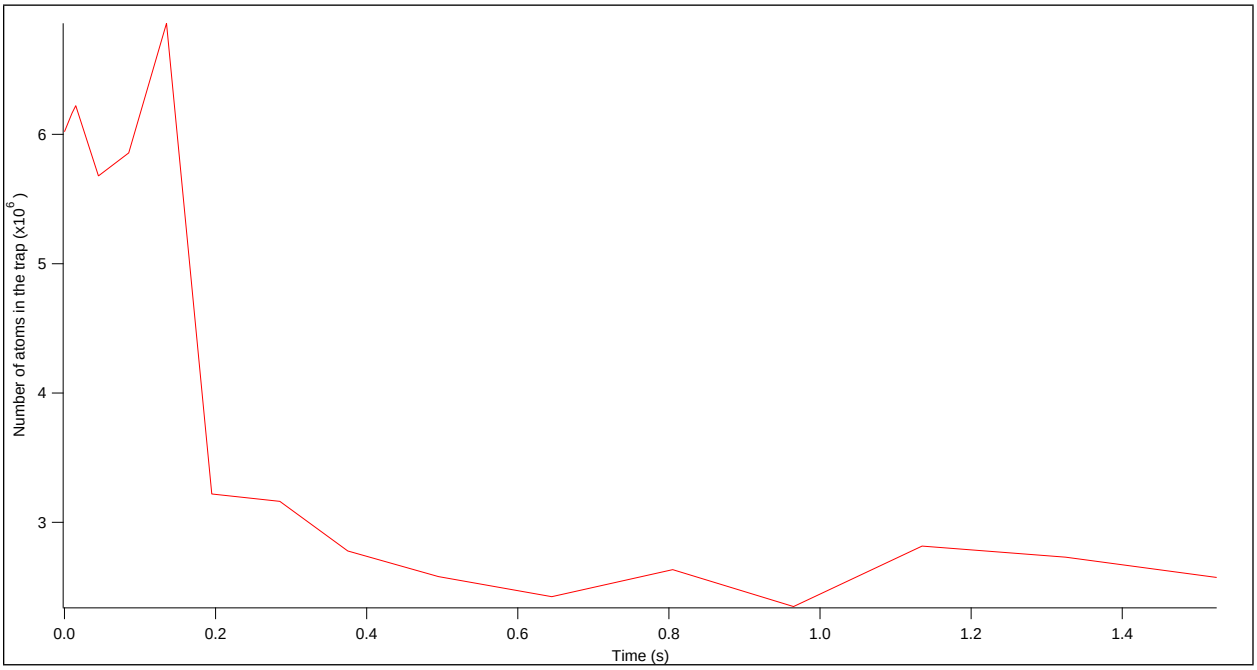

Figure 4.16: This is a plot showing the number of atoms in the crossed beam dipole trap during the experimental evaporation ramp that yields the best results.

It can be seen that none of the experimental values followed the expected trends of the simulation. Forced agreement with the simulation was attempted by manipulating the values of published constant coefficients or by altering the measured values of temperature, number, and phase space density. However, in order to get good agreement for the behavior of one of these values, such as temperature, it came at the expense of further disagreement with the trends in number and phase space density. No manipulation of the model or data resulted in good agreement for all the experimental parameters, and the amount of manipulation of the data required to get good agreement for a single parameter seemed unrealistic. Ultimately it was concluded the model developed in [75] is incomplete and not an appropriate model for this particular set of experimental conditions.

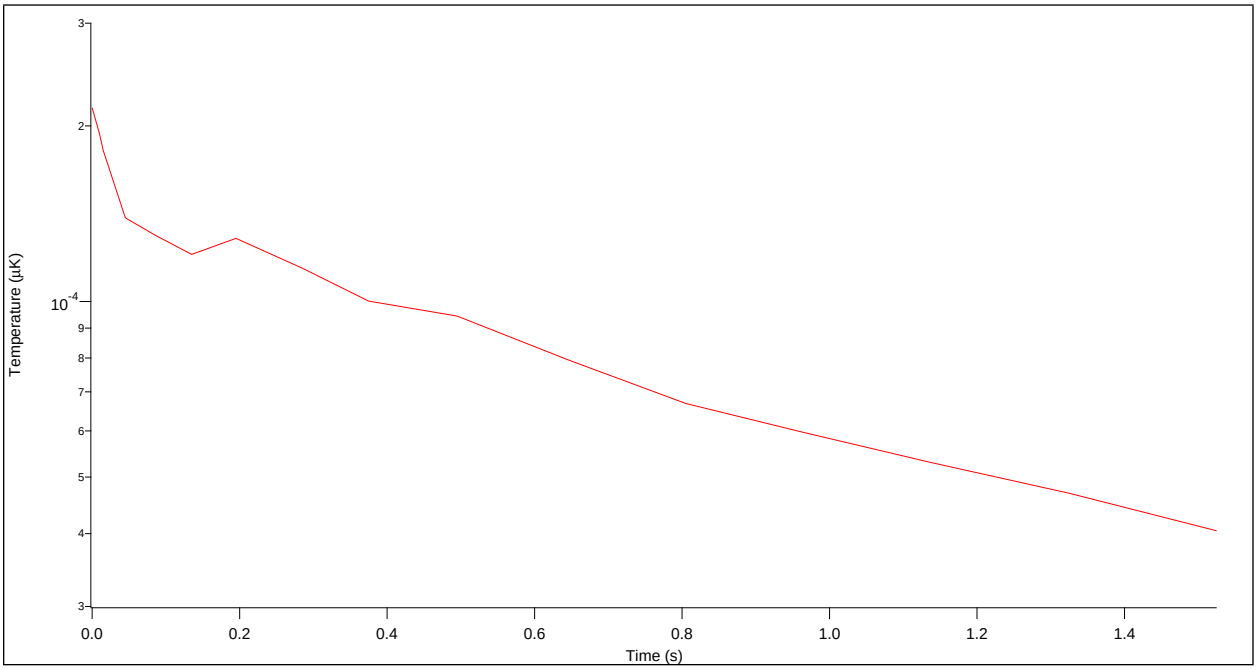

Figure 4.17: This is a plot showing the temperature in the crossed beam dipole trap during the experimental evaporation ramp that yields the best results. Notice that the temperature does not drop nearly as quickly as the simulation in Figure 4.14.

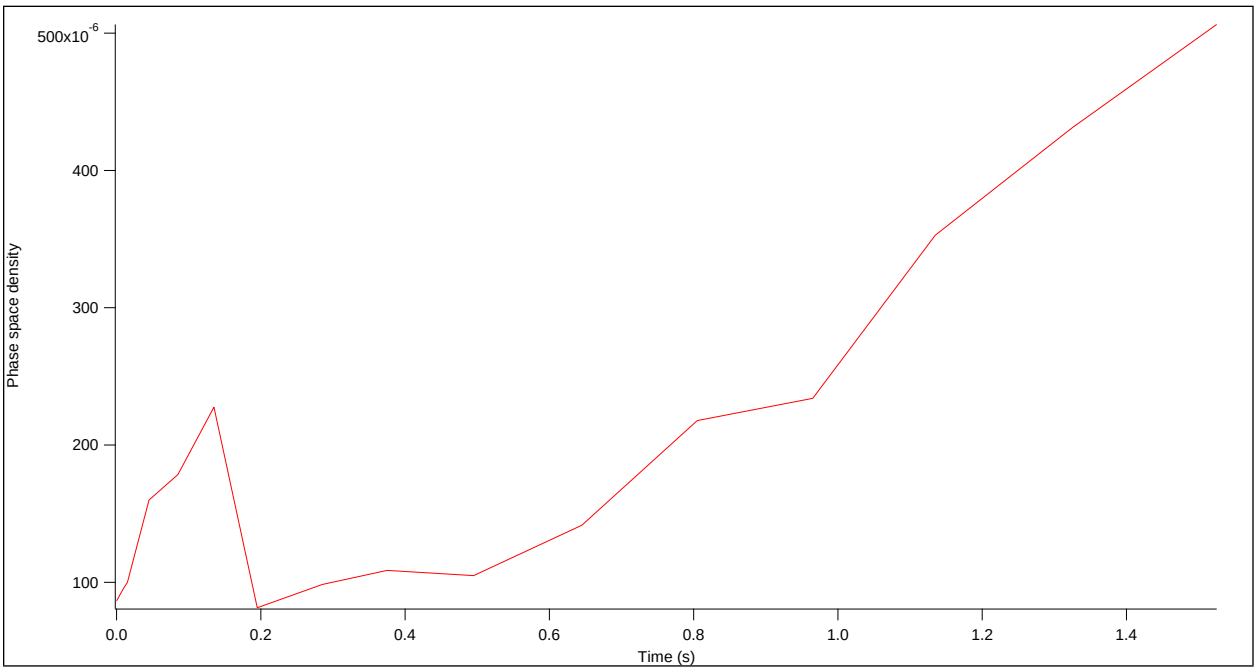

Figure 4.18: This is a plot showing the phase space density in the crossed beam dipole trap during the experimental evaporation ramp that yields the best results. The final data point is still six orders of magnitude below the value required to reach BEC.

A new model for evaporation

Upon further examination of loading into the crossed dipole trap, an observed increase in the number of atoms in the crossed portion of the trap was observed in the first 300 ms after turning off all UHV MOT potentials. The original assumption was that the maximum number of atoms in the crossed dipole trap occurred immediately after turning off the UHV MOT potentials. During the first 300 ms, a rapid drop in temperature was also observed. At 300 ms, when the number of trapped atoms reached a maximum, a slight increase in the temperature occurred before the number of atoms started to decrease, and the temperature fell back down to near its steady state value. This behavior coincides with the onset of three-body collisions and heating of the trap when the trap density becomes large. Also, during evaporation attempts, a change in the ratio of the number of atoms in the crossed portion of the trap, N_{cross}, to the number of atoms in the single beam "wing" portions of the trap, N_{wings}, was observed immediately after a step down to a lower η, as seen in Figure 4.19.

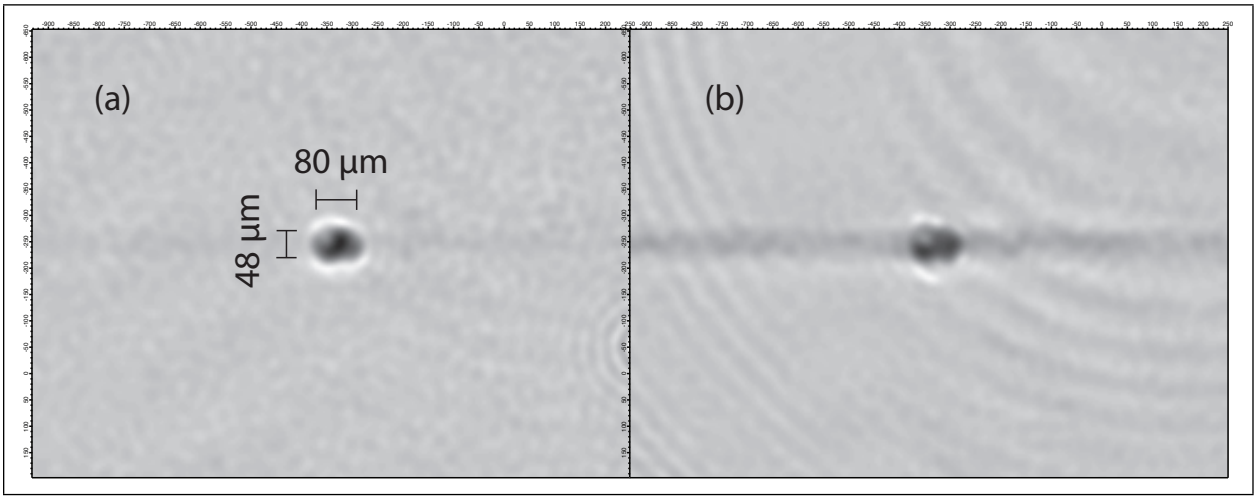

Figure 4.19: (a) A picture of the dipole trap just before the onset of forced evaporation shows a large ratio of the number of atoms in the crossed region to the number of atoms in the wings. (b) A picture of the dipole trap immediately following the first step from $\eta = 8$ to $\eta = 7$ in an attempt at forced evaporation shows a smaller ratio of the number of atoms in the crossed region to the number of atoms in the wings.

The migration of atoms from the wings into and out of the crossed portion of the dipole trap has been seen in previous experiments, but the phenomenon has largely been ignored as inconsequential [37, 76]. Our apparatus is unique, however, in that we use high power 1064 nm light for the dipole trapping beams and have large $w_0 = 38\,\mu$m beam waists to maximize loading from the UHV MOT. This leads to beams with a long Rayleigh length ($x_R = 4.3$ mm) and puts us in a regime where the interaction with the wings actually dominates the evaporation process. [77] is the only direct study of the interaction of the wings with the crossed trap performed to date that we know of in the literature. In our lab, our initial conditions leave the phase transition to BEC out of reach. In order to fully understand the trapping behavior related to our specific experimental setup and what changes are necessary to reach BEC, a new simulation was carried out to fully account for the presence of wings during the evaporation process.

The simulation shows that the wings play a detrimental role during the evaporation process. For two beams crossing on a horizontal plane, atoms in the crossed trap with energies $E < U_0'/2$ remain in the crossed trap. Atoms with energies $U_0'/2 < E < U_0'$ may leave the crossed trap, but remain in the wings. Finally, atoms with energies $E > U_0'$ escape the trap into free space. The trap frequencies of the crossed trap are high due to tight confinement, but the trap frequencies in the wings are much lower due to the highly elongated axial length of single beam potentials. Because of this difference in trap frequencies, the initial rate that atoms flow from the cross to the wings is high relative to the rate atoms flow back from the wings to the cross. However, the filling rate of the wings is also dependent on the population ratio of the wings to the cross. During the initial stages of forced evaporation, the high filling rate of the wings leads to a growing wing population which may eventually become larger than the population in the cross. At some point, the ratio N_{wings}/N_{cross} reaches a steady state value. According to theory, the average energy of atoms recaptured by the cross from the wings is higher than the average energy of atoms flowing from the cross into the wings, which leads to an undesirable net flow of energy into the cross. A simulation was performed,

taking the wings into account with the new model, for our current crossed dipole trapping beam configuration, seen in Figures 4.20, 4.21, and 4.22. It qualitatively agrees with the experimental data better than the simulation performed using the model in [75], although it still does not quite explain the observations made in the lab. Quantitatively there is still disagreement between the simulation with the new model and the experiment. The exact cause of this is unknown, but the simulation is highly dependent on precise knowledge of the operating parameters in the experiment.

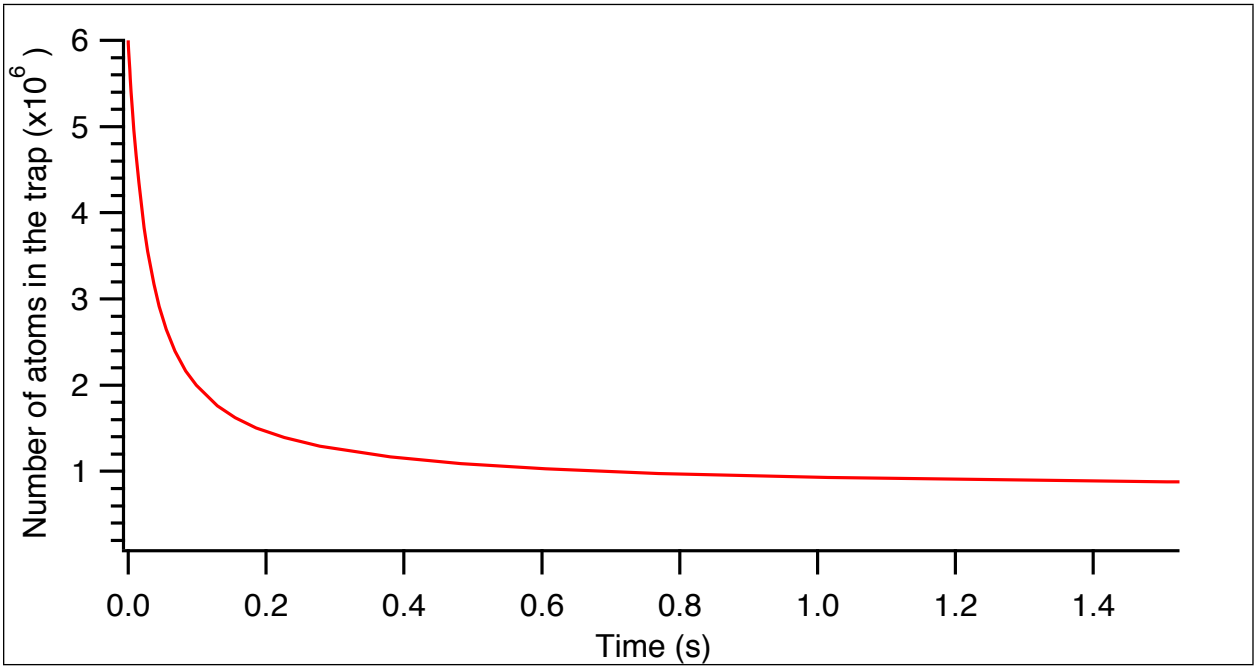

Figure 4.20: This is a plot showing the number of atoms in an evaporation simulation of the crossed beam dipole trap using the same experimental parameters as in Figure 4.16.

In order to solve the problem the wings present to evaporation, a tilted potential arrangement has been suggested [78]. By superimposing either a magnetic gradient or a gravitational gradient to the optical potential, escape channels may be created by which the atoms may flow from the cross into free space instead of the wings. This can be achieved by using either magnetic coils in the current configuration of horizontally crossed beams or by tilting the beams into a vertically crossed configuration. These changes will result in one downsloping lip in the potential and one upsloping lip in the potential along the wings directions, as seen

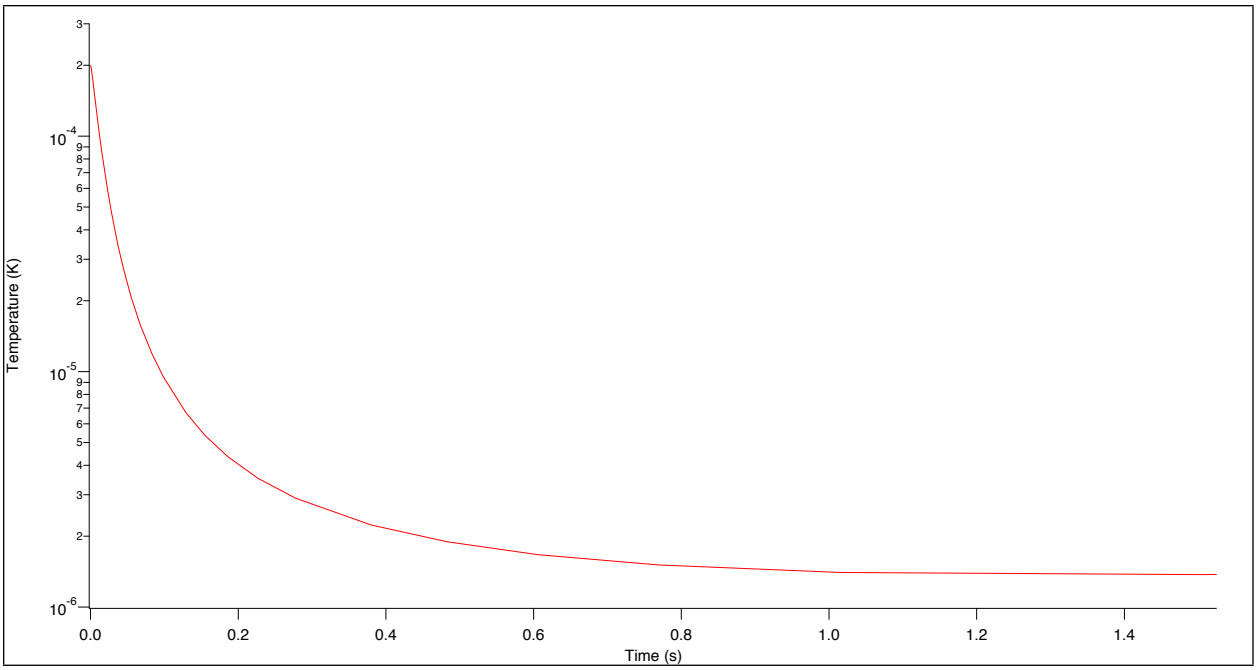

Figure 4.21: This is a plot showing the temperature in an evaporation simulation of the crossed beam dipole trap using the same experimental parameters as in Figure 4.17.

in Figure 4.23. Atoms that would fill the wings in the current experimental configuration would instead be reflected off the upsloping lip and escape to free space via the downsloping lip, the desired behavior during the evaporation process. Whereas the corrected simulation shows that reaching the BEC transition with the currently untilted potential setup might be possible only at unreasonably long evaporation times, using a tilted potential arrangement should produce a BEC with $\approx 3.5 \times 10^5$ atoms in $\approx$ 2s.

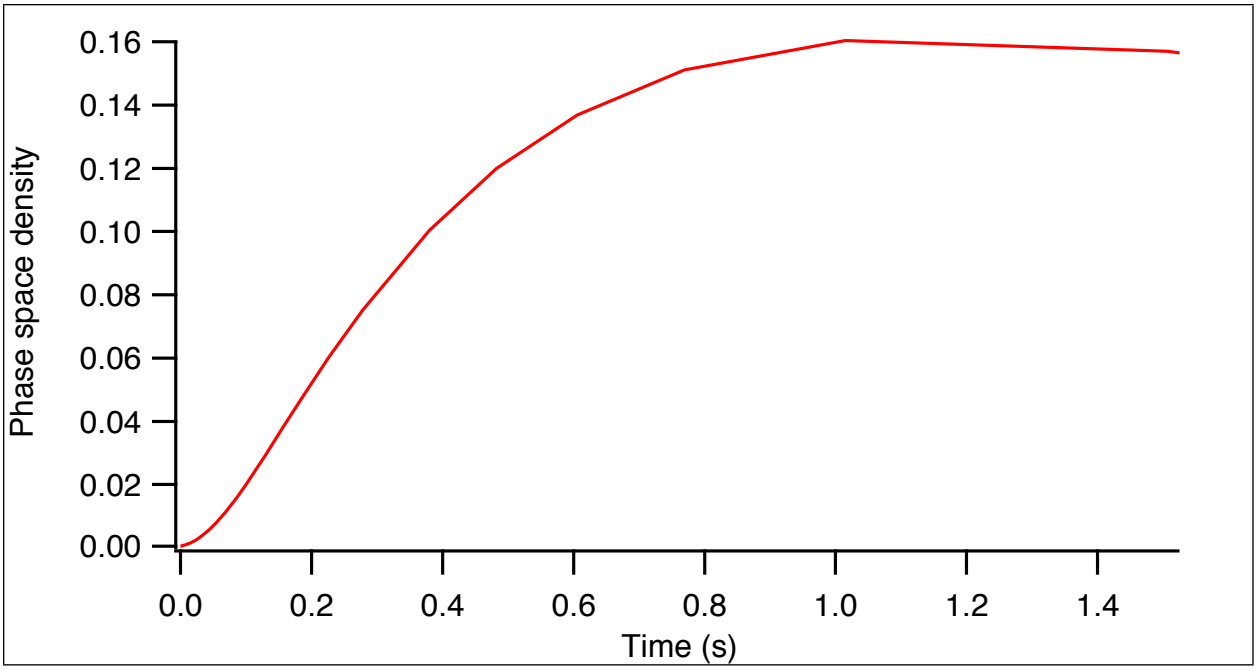

Figure 4.22: This is a plot showing the phase space density in an evaporation simulation of the crossed beam dipole trap using the same experimental parameters as in Figure 4.18.

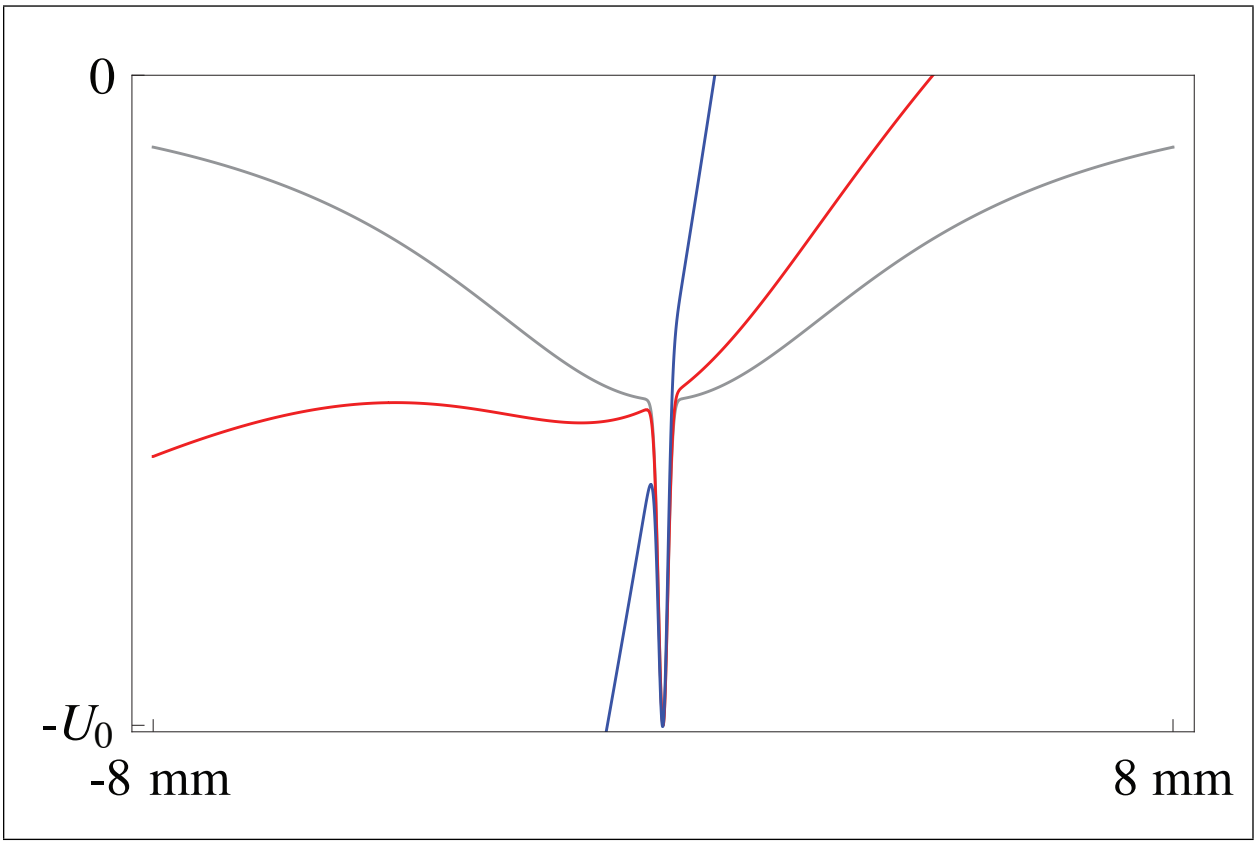

Figure 4.23: This is what a tilted potential would look like for a crossed beam dipole trap. The grey curve is an untilted potential (current arrangement). The red curve is a critically tilted potential where the wing on the left side is at the same height as the lip. The blue curve is a strongly tilted trap. For both the critically tilted and strongly tilted potentials, the left side provides an escape channel for evaporating atoms where they will not remain in the wings.

5 Magnetic Shield

5.1 Design and simulation

A passive magnetic shield is absolutely necessary to reduce low frequency background magnetic fields around the atoms to levels where the linear Zeeman energy is negligible compared to all spin interaction energies. The Earth's background magnetic field has a field strength of ~ 1 G. The regime the experiment needs to be in is in the range of ~ 1 μG or below. Therefore, the magnetic shield must be able to attenuate the background field by at least six orders of magnitude. In our apparatus we accomplish this with a multi-layer enclosed shield using an alloy with a very high magnetic permeability called mu-metal. Due to its high permeability, mu-metal provides a low reluctance path for magnetic flux. Rather than blocking the magnetic fields, mu-metal works by rerouting the magnetic field lines around the shielded area.

In order to reach such high shielding factors, special care must be taken in the design and fabrication of a mu-metal shield. Many factors affect the performance of the shield, including the shape of the shield, number of layers, the spacing between layers, end cap geometry, and holes in the shield. Unfortunately, mu-metal is a very expensive material, so compromises in performance must sometimes be made at the cost of the amount of mu-metal used. Many opinions and rules of thumb exist for the design of effective magnetic shields, but since the intended use is so varied between different experimental setups, each design must be customized for the specified requirements. This is best accomplished by simulating the magnetic shield before fabrication.

The simulation of magnetic shield performance was performed using the free version of the ANSYS Maxwell software. This software can simulate a material with the permeability

characteristics of mu-metal with varying geometries in the presence of a background magnetic field. Longitudinal shielding is always poorer than axial shielding in cylindrical shields with flat end caps with penetrations, and so is of greatest concern when optimizing the design [79, 80, 81]. Unfortunately, the free version of the software only allows simulations of cylindrically symmetric 3D shields. Our shield has four sets of 3 holes each for optical access along the radius of each layer of the shield, meaning it is not quite cylindrically symmetric. In addition to a 1.5 in diameter central hole along the longitudinal axis of the shield on each end cap which can be simulated, there are also four 1 in diameter off-axis holes and two 1.5 in $\times$ 1.5 in square holes on each end cap that cannot be simulated accurately using this software. To account for this lack of an exact simulation for our shield design, the central hole on each end cap was enlarged in the simulation to approximate the existence of the off axis holes that cannot be simulated. Fields that enter the shield through a hole are strongest along the hole axis, and, as a rule, fall off at a distance two hole diameters inside the shield, justifying this approximation [80]. Additionally, tubulated extensions may be added to each hole penetration to mitigate this factor. A picture of the shield detailing the actual shield geometry can be seen in Figure 5.1.

The size of the outer layer of the shield is constrained by space in the experimental setup to be 28 in long with a 26 in diameter. Other variables that can be simulated for effectiveness include the number of layers, the thickness of each layer, the spacing between each layer, and the length of the tubulated extensions for each hole. Shields with four, five, and six layers were all examined. In all geometry arrangements, four layers appeared to provide insufficient shielding for the experimental requirements. Five layers demonstrated a factor $\approx$ 10 times improvement, and six layers, as should be expected, performed even better, providing an additional factor of $\approx$ 10 improvement. However, the cost of material for adding a sixth layer became prohibitive, and five layers performed sufficiently for the experiment according to the simulation. Shield layer thicknesses of 0.025 in and 0.050 in were also tested in different arrangements. Layers with smaller thicknesses saturate more

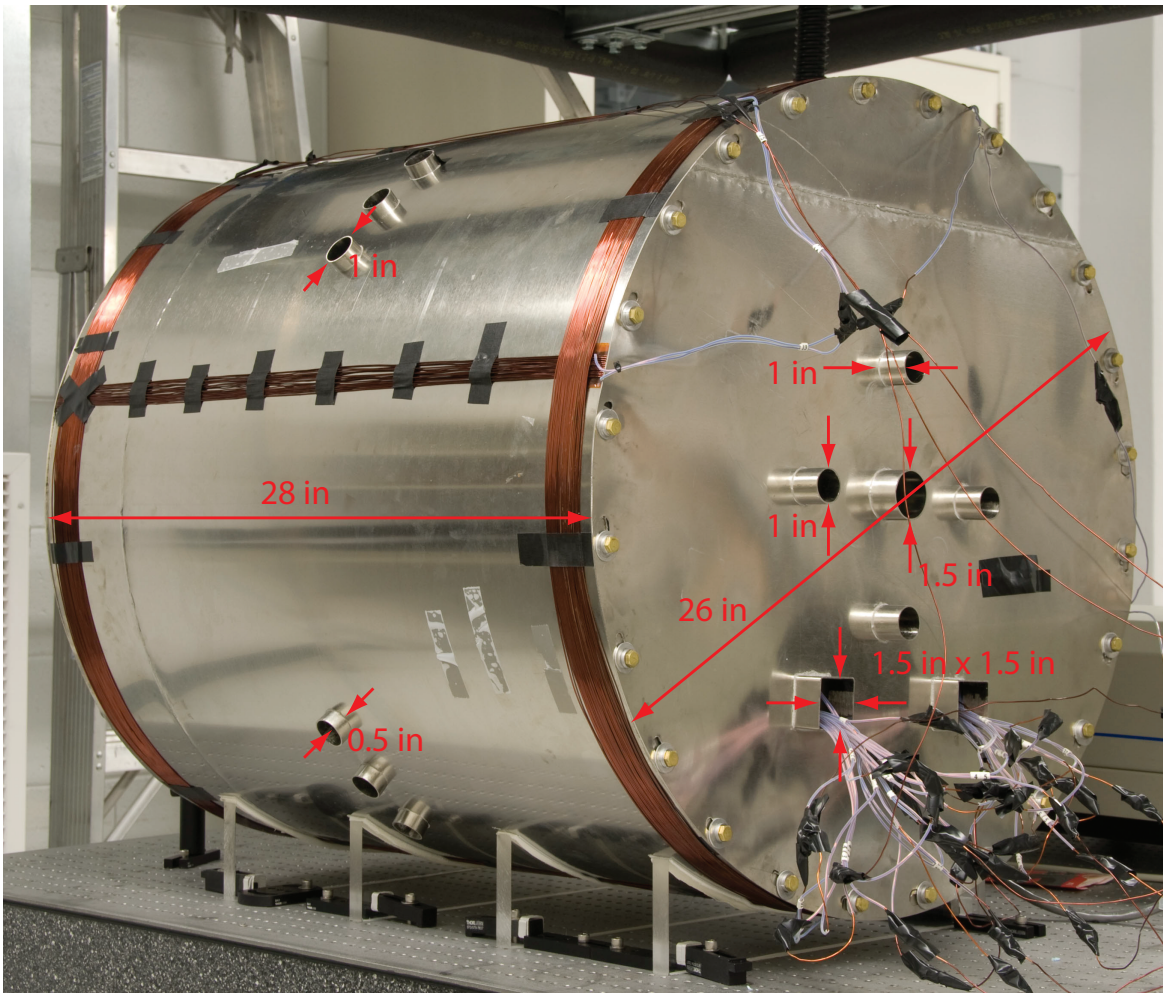

Figure 5.1: This is a picture of the actual mu-metal shield in the lab with dimensions labeled for the outermost layer.

easily than thicker layers and provide poorer shielding performance, but they also require substantially less material and are cheaper. In addition, the outer layer is exposed to the highest flux density, so it was chosen to be thicker. As a compromise, the outer shield layer was chosen to be 0.050 in thick, while the four inner layers are each 0.025 in thick. Contrary to the results published in [81], where they used fewer layers with larger gaps, for our apparatus a geometric series increase in the radial spacing between each layer performed only marginally better than an equal spacing between each layer. Since the average size of

each layer is smaller with an equal spacing between layers than a geometric spacing, the cost savings indicate it is beneficial to use five layers with a 1 in radial gap between each layer. The results are similar in the longitudinal direction, except larger gaps between layers lead to better longitudinal shielding. There is a 1.5 in gap in the longitudinal spacing between each layer. The tubulated extensions at each hole are critical for the longitudinal shielding performance. Without the extensions, the magnetic field lines are not effectively absorbed by the mu-metal in the longitudinal direction at the holes. These field lines can penetrate directly to the center of the shield and prove detrimental to the experiment. As a rule of thumb, the length of the extensions should be about the same size as the diameter of the hole for maximum performance. Again, the outer shield is the most important, so our final design has 1 in extensions on all end cap holes, and the inner shields have 0.75 in extensions on all holes. Figure 5.2 shows an example simulation of an approximation to our actual shield design.

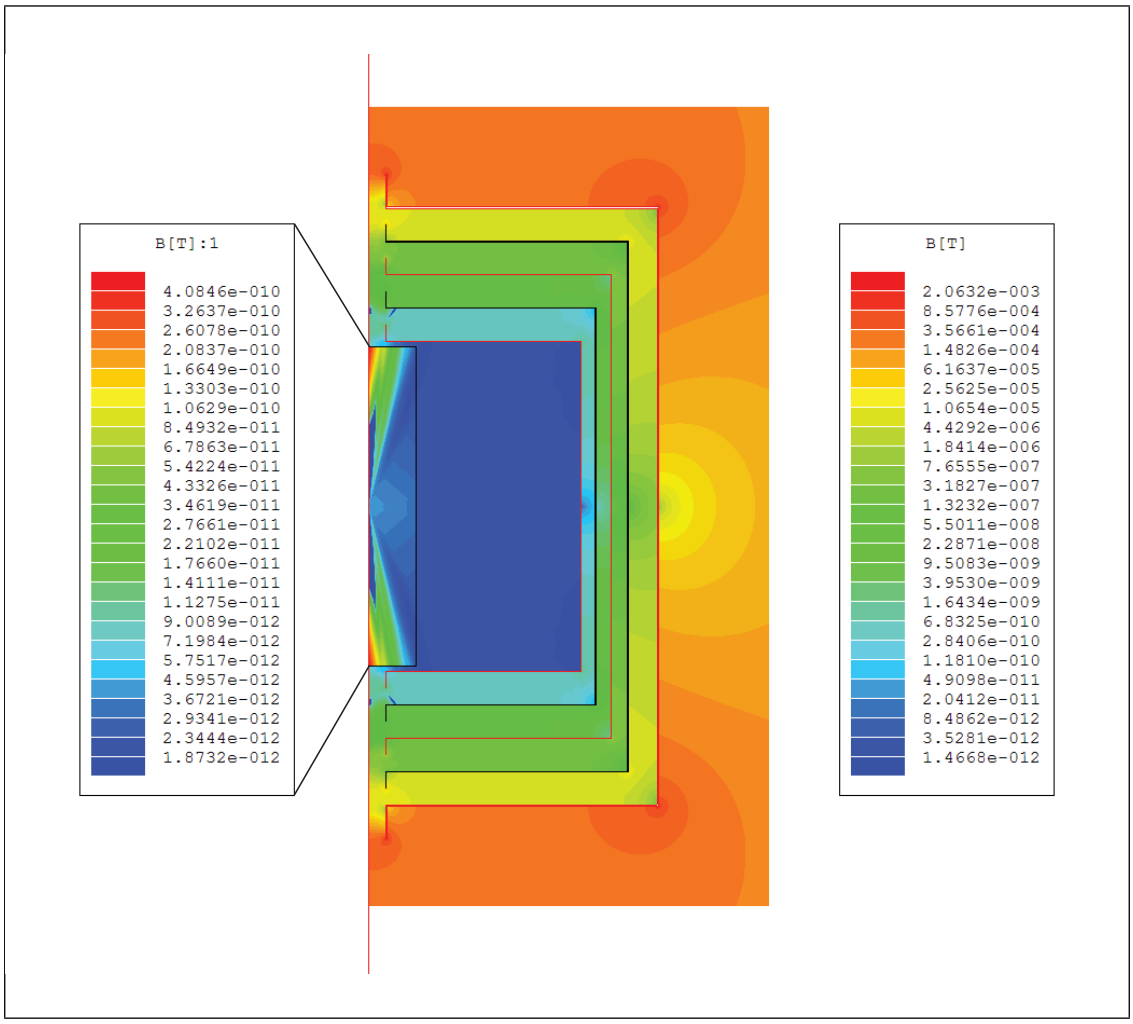

Figure 5.2: This plot shows a simulation of an approximation to our shield design in a 1 G background magnetic field. The color scale on the right shows the field strength for all regions of the plot except the very center of the magnetic shield. The color scale on the left shows the field strength at the center of the shield in the region of the atoms. Both scales are in units of Tesla.

5.2 Construction and initial measurements

Construction

As was reported earlier, the outer layer of the shield is 28 in long with a diameter of 26 in. The shielded volume is defined by the dimensions of the inner layer of the shield, which is 16 in long with a diameter of 18 in. The end caps are flanged and bolted to the shield as opposed to end caps that slide over the cylinder. This is commonly accepted as the best method to attach flat end caps to a cylindrical shield due to imprecise manufacturing standards that can lead to poor mating otherwise. The outer layer of the shield is 0.050 in thick, while each of the four inner layers is 0.025 in thick. The radial spacing between each layer is 1 in equal spacing, and the axial spacing between each layer is 1.5 in equal spacing. The tubulated extensions for each hole on the outer layer end caps is 1 in, regardless of the hole diameter, for ease of construction. Each inner layer uses 0.75 in tubulated extensions for each hole, regardless of the hole diameter. The holes on the radial sides of the shield each have 0.5 in tubulated extensions. The long glass UHV cell of the vacuum chamber slides through the 1.5 in diameter hole at the center of one end cap, and the science chamber rests at the center of the shielded volume. The 1.5 in diameter hole at the center of the opposite end cap is used for optical access of the imaging lens relay system used for CCD images of the dipole trap and BEC. The 1.5 in $\times$ 1.5 in square holes on both end caps are used for an aluminum rail system that supports all the UHV MOT, dipole trap, and optical lattice optics and magnetic coils housed inside the shield. Four of the holes on the radial sides of the shield are used for a rod support system that stabilizes the optical breadboards inside the shield to a rigid cage system external to the shield. The remaining holes in the shield are for optical access for such things as the dipole trapping beams, optical lattice beams, UHV MOT fiber optic cables, and other auxiliary beams that might be required in the future.

Shielding factors

To verify the shielding performance of the actual magnetic shield, large electromagnetic coils were built along the longitudinal axis of the empty shield in a Helmholtz configuration to simulate a uniform magnetic field in the region of the atoms. A fluxgate magnetometer was placed at the center of the shield, and the end caps were sealed. Measurements of the magnetic field strength at the center of the enclosed shield were taken in the presence of both large and small external magnetic fields. For each successive measurement, the polarity of the field was reversed in order to remove any effects of long term linear drift in the magnetometer. The results of these measurements can be see in Figure 5.3 and are summarized in Table 5.1. A diagram of the alignment of the shield with respect to the coils can be seen in Figure 5.4. As is to be expected, the longitudinal shielding factor was worse than the radial shielding factor for both large and small external fields. In the large external field, the longitudinal shielding factor is saturated, providing an upper limit of our longitudinal shielding ability. In the low external field regime, the longitudinal shielding factor is significantly reduced, but the attenuated field strength of $0.24\,\mu$G meets the requirements of the experiment. In both the large and small external field regimes, the measured attenuated radial field is below the noise floor of the magnetometer. It can therefore be concluded that 2 mG DC fluctuations can be sufficiently attenuated to experimental requirements by the shield.

	$\Delta B_{outside}$	ΔB_{inside}	Shielding factor
x	600 mG	7.83 μG	76000
x	2 mG	0.24 μG	8300
z	600 mG	0.17 μG	3.50×10^6
z	2 mG	0.15 μG	12900

Table 5.1: Measured magnetic field strength at the center of the magnetic shield in the presence of an external field. The x direction is the longitudinal axis, and the z direction is the radial axis.

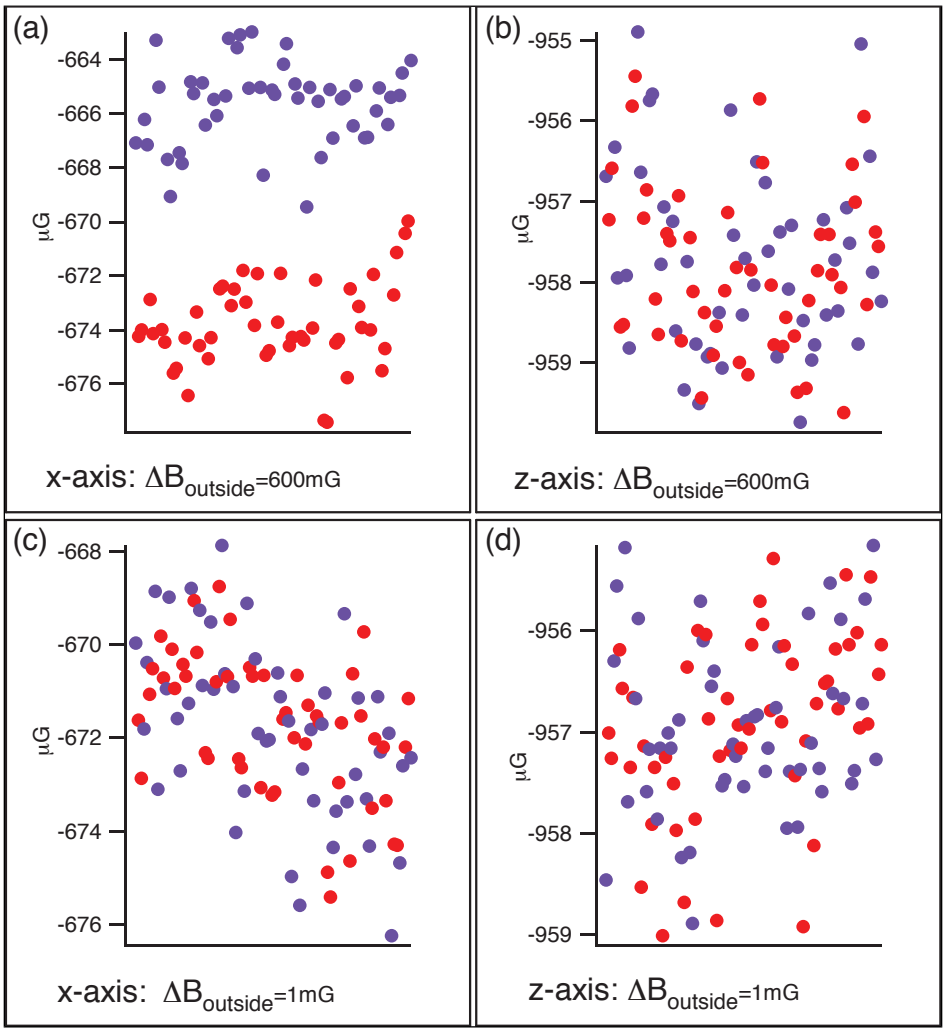

Figure 5.3: Magnetic field measurements at the center of the magnetic shield. The polarity of the external field was reversed after each successive measurement to remove linear drift. The purple points are all taken in external fields of positive polarity, and the red points are all taken in external fields of negative polarity. Plot (a) was measured along the longitudinal axis in a strong external field, (b) is along the radial axis in a strong external field, (c) is along the longitudinal axis in a weak external field, and (d) is along the radial axis in a weak external field. The results of these measurements are summarized in Table 5.1.

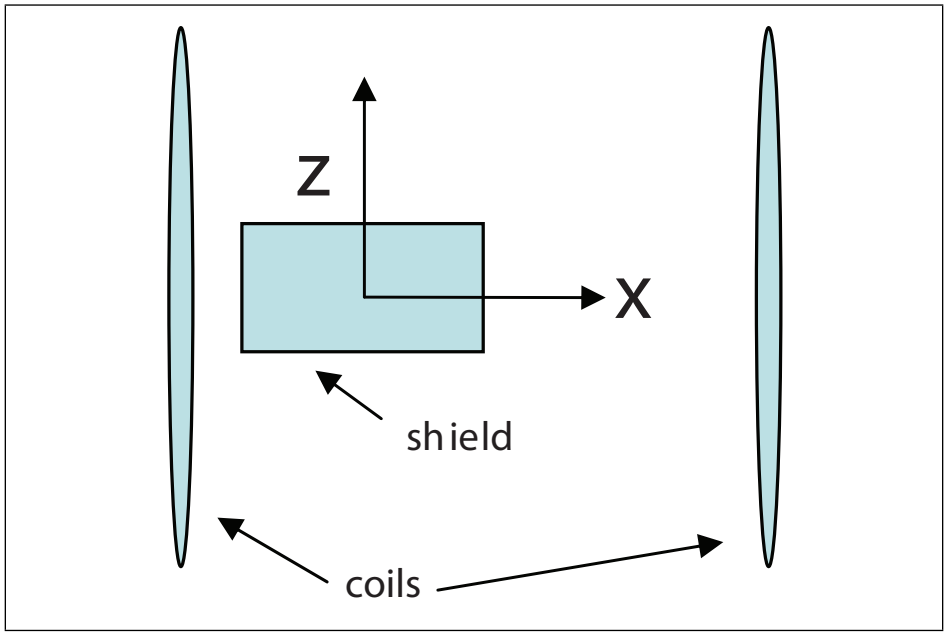

Figure 5.4: This diagram shows the alignment of the magnetic shield with respect to a pair of Helmholtz coils to measure the longitudinal and radial shielding factors.

Noise measurements

The low frequency noise performance of the magnetic shield was also measured. The source of the low frequency noise was taken as the ambient noise in the laboratory, which exhibited fluctuations up to 100 mG. The magnetometer was placed in the center of the enclosed shield, and a set of measurements with integration times $\tau = 0.1, 1$, and 10s was taken using a 7.5 digit digital voltmeter. The linear drift was removed, and the RMS of the residuals was calculated to determine what field noise the atoms would be exposed to. The data collected this way can be seen in Figure 5.5. Figure 5.6 shows an Allan deviation plot comparing the measured data to the published magnetometer rms noise floor of 0.3×10^{-6} G/$\sqrt{\text{Hz}}$. In ultralow magnetic field experiments we do not expect to look at any physics on time scales shorter than 1 s, in which case the shot to shot fluctuations in the background field would lead to uncertanties of $1.6\,\mu$G. These measurements show the magnetic shield is capable of producing a quiet magnetic environment in bandwidths relevant to the experiment.

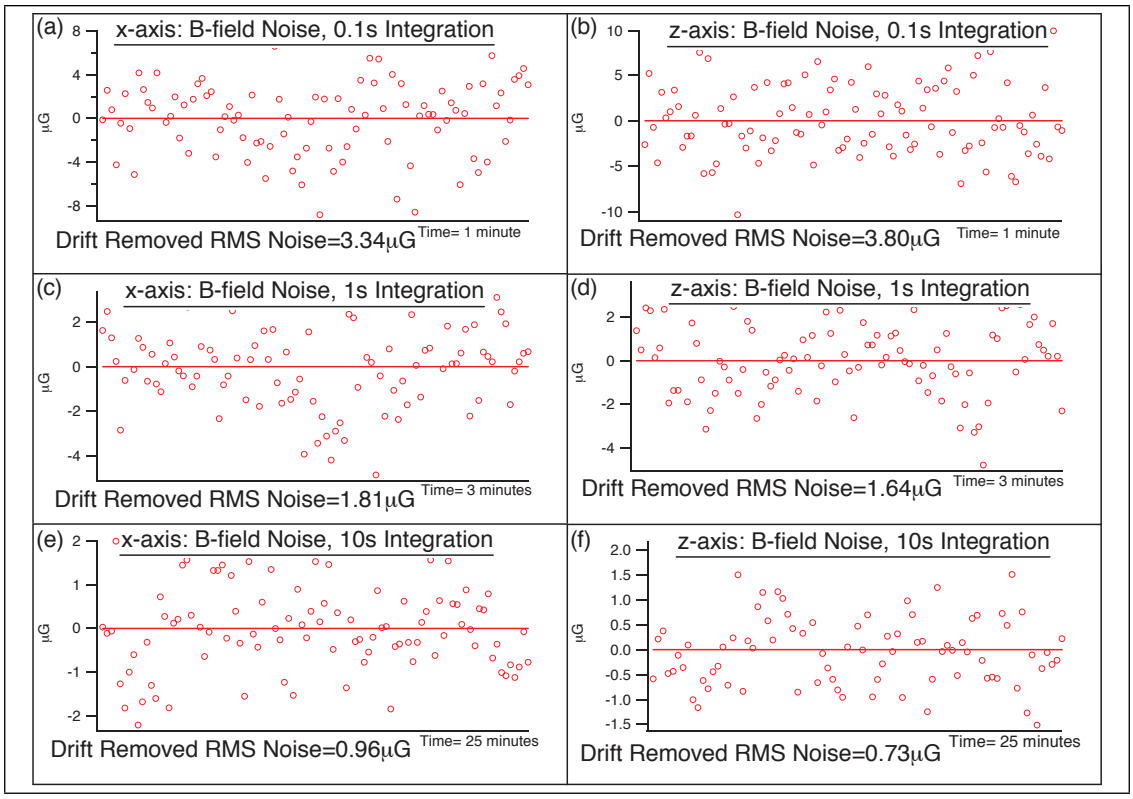

Figure 5.5: Magnetic shield noise measurements. The noise source was taken to be the low frequency ambient noise of the lab. Plots (a) and (b) show drift removed RMS noise measurements for an integration time $\tau = 0.1$ s taken over a 1 minute experiment window. Plots (c) and (d) show noise measurements for an integration time $\tau = 1$ s taken over a 3 minute experiment window. Plots (e) and (f) show noise measurements for an integration time $\tau = 10$ s taken over a 25 minute experiment window. The results of these measurements are summarized in Figure 5.6.

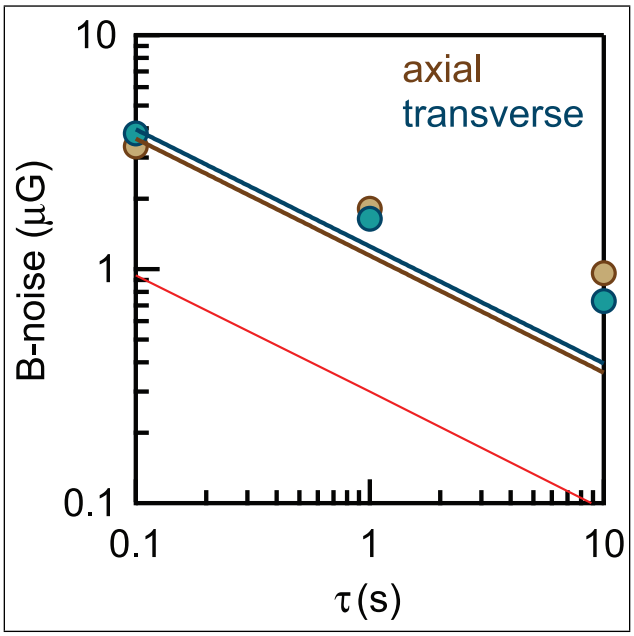

Figure 5.6: This is an Allan deviation plot summarizing the results of Figure 5.5. The points are the Allan deviation of low frequency ambient noise measurements for integration times $\tau = 0.11, 1,$ and 10 s. The red curve is the published noise floor of the magnetometer. The blue and brown curves show what one would expect if the $\tau = 0.1$ data points along each direction behaved as white noise and were integrated down as $1/\sqrt{\tau}$. Because the data points lie above these lines, it indicates the observed noise is slightly pink.

Shield degaussing

Because mu metal is a ferromagnetic material, it can suffer from a remanent magnetization after being removed from an applied field. This remanent magnetization induces a field of its own which can be seen inside the shielded volume. Remanent fields therefore pose a serious problem for experiments working in the low magnetic field regime. By properly degaussing the shield, remanent magnetization should be minimized to reduce the negative impact of any residual fields inside the shield. To degauss the shield, two solenoidal coils and four toroidal coils were wrapped around each layer. The number loops in each coil was increased to deliver a saturating field (a value supplied by the shield manufacturer) at 10 A of current. An exponentially decaying 60 Hz sine wave current was applied to the coils for 30 s. The phase of the sine wave was adjusted to be at a node when the current was abruptly turned off. The radial remanent field after this process was 0.64 mG with a standard deviation of 0.968 μ G. The axial remanent field, however, was 1.83 mG with a standard deviation of 0.42 mG. This lack of reproducibility can pose serious problems in the experiment. A possible explanation for this large variance in the axial remanent field could be that the large inductance of the shield acted as a filter in the presence of the degaussing windings. As a result, the actual flux density driven in the mu metal shield did not follow the driving voltage.

The presence of MOT coils inside the shield can generate large magnetic fields in the vicinity of the inner layer of the shield. If these fields require the shield to be degaussed between successive experimental cycles, the unknown resultant remanent field could be a problem. A slower degaussing process might improve the situation, but that would force the repetition rate of the experiment to be impractically low. Fortunately, after an initial degaussing of the shield, it was discovered that the background remanent field inside the shield reached a constant value after ≈ 10 experimental cycles with no degaussing in between cycles. This constant remanent field can simply be canceled by small shim coils centered around the experiment chamber inside the shield. This critical result simplifies experiment

protocols significantly, removing the need to degauss the shield between each cycle of the experiment.

5.3 Low-noise field canceling coils

A passive magnetic shield only effectively attenuates background magnetic fields in a bandwidth range of 0.1 Hz to 10 Hz. Any field fluctuations occurring at frequencies faster than 10 Hz must be canceled using an active field canceling source. In principle, this could be achieved by any simple shim coil system with a feedback controlled servo loop. In practice, however, the noise written to the coil current by the servo loop is nontrivial when the coils are designed to cancel small magnetic fields with a precision of ~ 1 μG or better. To that end, a low noise current source is required that can meet these requirements. We propose to build a circuit similar to the one described in [82] to perform this task, although it has not yet been built or tested for our experiment.

5.4 Practical considerations for stand alone equipment inside the shield

The UHV MOT, dipole trap, optical lattice, and imaging optics must all reside inside the magnetic shield. Shielding the experiment chamber from external background magnetic fields would do little good, however, if the stand alone equipment inside the shield generated magnetic fields of its own. Therefore, all materials inside the shield must be non-ferrous. Every piece of equipment inside the shield is designed to be made of aluminum, brass, or glass. The magnetic shield slides over two non magnetic aluminum rails. An aluminum breadboard rests on top of the rails inside the shield to hold all the optics necessary for the experiment. Basic Thorlabs post holders and bases are already non magnetic and are used inside the shield. Optical posts and mounting screws are usually stainless steel due to their rigidity and lack of mechanical drift. Stainless steel, however, is magnetic and cannot be used inside the shield. The best alternative is aluminum posts and screws. Aluminum is much softer than stainless steel, so perturbations to the apparatus can easily move components fastened

with aluminum. While this instability is undesirable, perturbations should be minimal since the shield will surround the entire apparatus and prevent unintentional bumps. Typical mirror mounts are non magnetic, but most mounts contain stainless steel springs and/or stainless steel screws that cannot be replaced with non magnetic substitutions. Instead, aluminum flexure mirror mounts from Siskiyou provide an attractive solution. The low profile monolithic mirror mounts are solid aluminum, and the manufacturer supplied stainless steal set and adjustment screws can be easily swapped with standard brass screws. The fiber optic tips are nickel plated brass, and the nickel plating does have a small magnetization. However, the fiber tips are believed to be located far enough away from the experiment chamber that their magnetic fields will fall to sufficient levels in the region of the atoms.

Some experimental components contain magnetic materials that cannot be removed or sufficiently isolated inside the shield. These components must be placed outside the shield with optical access to the shielded region. For example, the CCD camera used for imaging the dipole trap and BEC contains magnetic materials without which it cannot function. A 1:1 imaging lens relay system, shown in Figure 5.7, made of non magnetic materials inside the shield brings the image of the atoms out of the magnetic shield through the center hole on one of the end caps to the CCD camera. The dipole trapping beams are brought into the shield from the exterior because of both the large number of magnetic components they use for proper beam shaping and alignment and the large volume of space these components require.

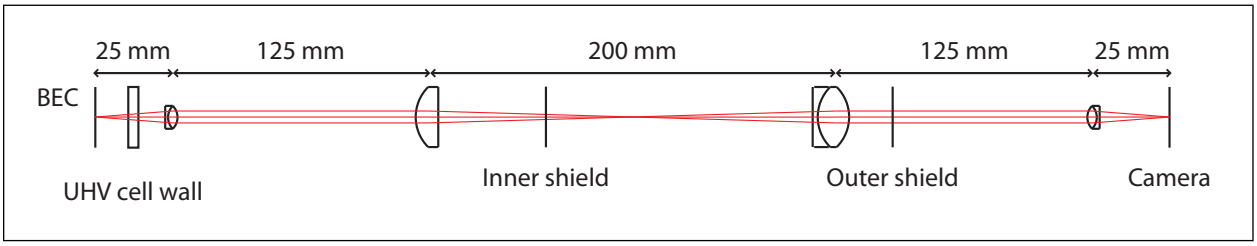

Figure 5.7: This is a diagram of the 1:1 imaging lens relay system that images the dipole trap/BEC inside the magnetic shield onto a CCD camera outside the shield

6 Rabi Oscillations

6.1 Rabi frequency and light shifts

The Hamiltonian for an atom in a radiation field can be written in the form $\mathcal{H}(t) = \mathcal{H}_0 + \mathcal{H}'(t)$, where $\mathcal{H}_0$ is the field-free, time independent atomic Hamiltonian and $\mathcal{H}'(t)$ describes the atomic interaction with a classical electromagnetic field. The Hamiltonian $\mathcal{H}_0$ has the commonly understood eigenvalues $E_n = \hbar\omega_n$. For the case of a two level atom driven by a narrow band laser excitation with ground state $|g\rangle$ and excited state $|e\rangle$, the interaction term of the Hamiltonian can be expressed as the interaction term for a dipole transition in the form

$$\mathcal{H}'(t) = -e\vec{E}(\vec{r}, t) \cdot \vec{r}, \tag{6.1}$$

where e is the charge of the electron and $\vec{E}(\vec{r}, t) = E_0\hat{\varepsilon}\cos(kz - \omega_l t)$ is the electric field operator for a plane wave traveling in the positive z direction. Here, E_0 is the amplitude of the light field, $\hat{\varepsilon}$ is the unit polarization vector, and ω_l is the laser frequency driving the transition. In this case, the probability of finding the atom in the excited state is [39]

$$P_e(t) = |c_e(t)|^2 = \frac{\Omega^2}{\Omega'^2}\sin^2\frac{\Omega' t}{2}. \tag{6.2}$$

Here, $\Omega = (-eE_0/\hbar)\langle e|r|g\rangle$ is the Rabi frequency, which describes the oscillation frequency of the atomic transition in a light field on resonance with the transition. The generalized Rabi frequency, $\Omega' = \sqrt{\Omega^2 + \delta^2}$, describes the oscillation frequency of the atomic transition in a light field at a detuning $\delta = \omega_l - \omega_0$ away from resonance.

In the presence of an electromagnetic field, the eigenenergies E_n of the field-free atomic Hamiltonian are no longer the eigenenergies for the full Hamiltonian. The effect of far-

detuned laser light on the atomic levels leads to an ac-Stark shift, or a light shift. Taking into account all electronically excited states $|e_j\rangle$, the *differential* light shift of an electronic ground state $|g_i\rangle$ ($5^2S_{1/2}\,F=1$ or $5^2S_{1/2}\,F=2$) is

$$\Delta E_i = \frac{3\pi c^2}{2} I \times \sum_j \frac{c_{ij}^2 \Gamma_{ij}}{\omega_{ij}^3(\omega_{laser} - \omega_{ij})}, \tag{6.3}$$

which describes the change in energy of the ground states relative to their field-free energy values E_n [53]. The relevant transition strengths c_{ij}^2 are 2/3 for the D_2 line and 1/3 for the D_1 line. Alternatively, the differential light shift of the ground states can be expressed in terms of a Rabi frequency, which in the limit of $\Omega \ll |\delta|$ is

$$\Delta E_g = \frac{\hbar\Omega^2}{4\delta} \tag{6.4a}$$

$$\Delta E_e = -\frac{\hbar\Omega^2}{4\delta}, \tag{6.4b}$$

where δ is the detuning of the light field from the frequency difference between the hyperfine ground states [39].

6.2 Rabi frequency measurements

We will use the first order magnetic field sensitive transitions from the $F = 1 \rightarrow F = 2$ ground state hyperfine manifolds for high precision magnetic field measurements in our apparatus. By measuring the Rabi frequency in the presence of a small bias field tuned to the transition frequency, the Zeeman shifts can be determined at different field strengths. Since the fluxgate magnetometer cannot see the field at the exact location of the atoms in the glass cell, measuring the Zeeman shifted $5^2S_{1/2}$ states will provide a measurement of the magnetic field experienced by the atoms in the trap. The magnetic sublevels of the $F = 1$ and $F = 2$ states exhibit -0.70 MHz/G and $+0.70$ MHz/G Zeeman shifts, respectively. For $1\,\mu$G fields, this corresponds to Zeeman shifts of ± 0.70 Hz.

As a proof of principle experiment to show that we can drive Rabi oscillations on the magnetic dipole transition, a microwave frequency generator was connected to a microwave waveguide and pointed at the crossed dipole trap. The sine wave with a frequency close to the 6.8 GHz transition was amplified and supplied to the waveguide for varying pulse lengths. At the appropriate frequency, this generated oscillations between the populations of the $F = 1$ and $F = 2$ states inside the dipole trap. The atoms in the dipole trap are initially prepared in the lower $F = 1$ ground state by turning the repumper beams off 1 ms before the MOT beams are turned off during the loading of the dipole trap. Typically, this means an absorption image with a probe beam tuned to the $F = 2 \rightarrow F' = 3$ transition would require the presence of repumper light to see any atoms. If Rabi oscillations are present, however, an absorption image with no repumper light will show atoms that have been excited to the $F = 2$ state by the microwave field.

Figure 6.1 shows a Rabi oscillations experiment for several microwave detunings at constant microwave power. Because this data was taken in the absence of the magnetic shield, all magnetic sublevels with $m_F \neq 0$ are Zeeman shifted far away. Out of convenience, we choose to observe only the transitions that occur between the first order magnetic field insensitive $m_F = 0$ sublevels of the two hyperfine states. As no state preparation was performed on the atoms in the dipole trap, the atoms should initially be spread evenly across the $m_F = 0, \pm 1$ sublevels of the lower hyperfine ground state. Therefore, the peak number of atoms observed in the experiment corresponds to about 1/3 the total number of atoms in the trap. According to Equation (6.2), the resonant microwave frequency will have the lowest observed Rabi frequency as well as the largest oscillation amplitude for a given microwave power. The experiment shows the resonant microwave frequency to be 6.8346890 GHz. Due to the presence of the dipole trapping beams during the Rabi oscillations experiment, the resonance frequency of the microwave transition is shifted by the optical potential. Therefore, the resonant microwave frequency of the magnetic dipole transition will differ slightly from the field free value of 6.834682611 GHz. The expected differential light shifts of the

two $5S_{1/2}$ ground states can be found using Equation (4.19) and taking

$$\Delta E_{F=1} + \Delta E_{F=2} = U_{0,F=2} - U_{0,F=1}. \tag{6.5}$$

For dipole trapping beams with characteristics given in Chapter 4, $U_{0,F=1} = 2\pi\hbar \cdot 51.765319$ MHz and $U_{0,F=2} = 2\pi\hbar \cdot 51.771249$ MHz, giving an expected differential light shift of $\Delta E_{F=1} + \Delta E_{F=2} = 2\pi\hbar \cdot 5.930$ kHz. This is in good agreement with the measured resonant microwave frequency of 6.8346890 GHz, which is 6.389 kHz away from the expected field-free resonance frequency of 6.834682611 GHz.

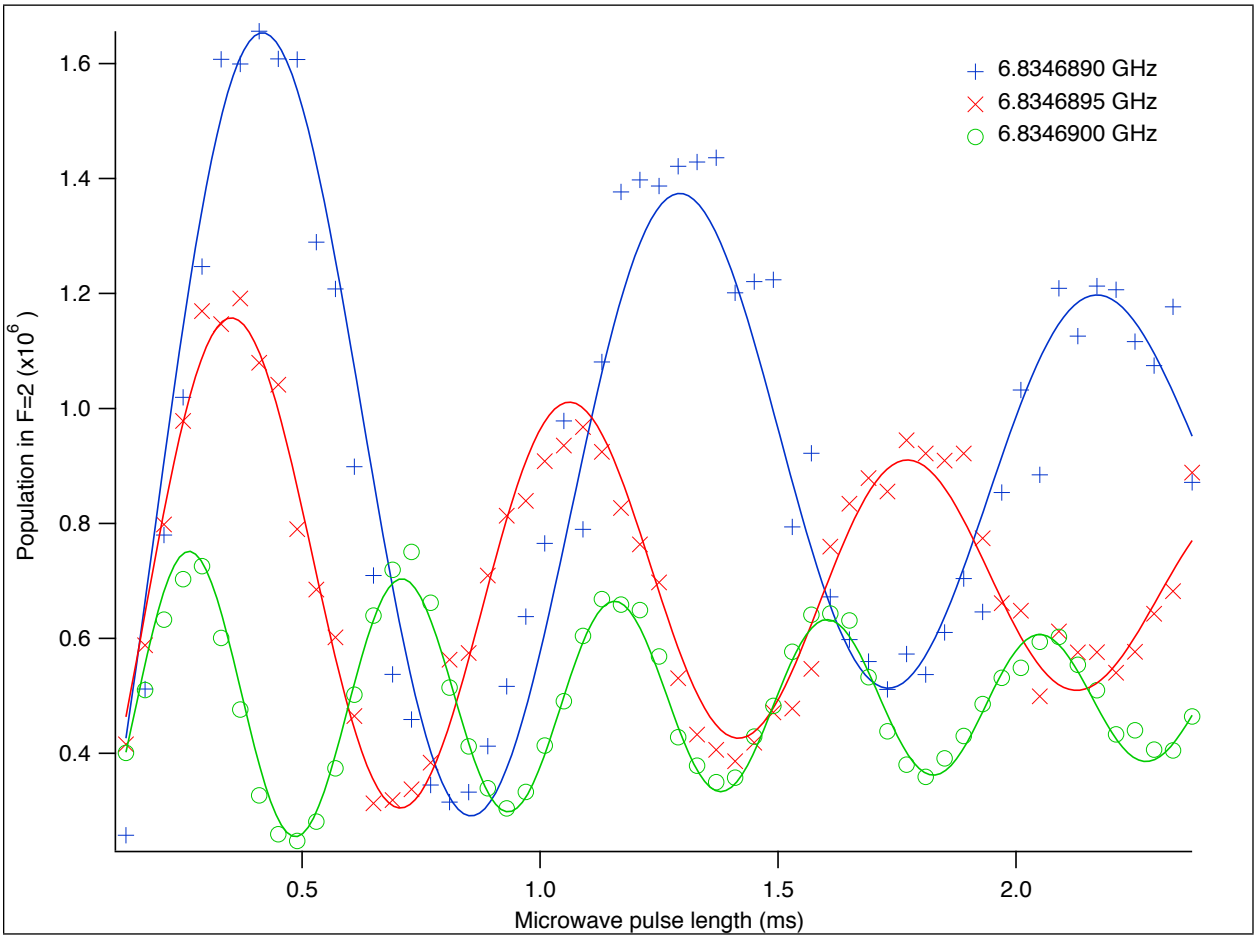

Figure 6.1: Rabi oscillations on the $5S_{1/2}\,F = 1 \rightarrow F = 2$ magnetic dipole transition. The blue trace shows Rabi oscillations at the resonant microwave frequency of 6.8346890 GHz. The red trace shows Rabi oscillations detuned +0.5 kHz from resonance. The green trace shows Rabi oscillations detuned +1.0 kHz from resonance. The data was fit using the function $f(t) = Ae^{-t/\tau}\sin(\Omega' t + \phi) + B$. At resonance (blue), a maximum of 1.6×10^6 atoms are observed in $F = 2$ with a Rabi frequency $\Omega = \Omega' = 2\pi \cdot 1.14$ kHz. The damping time constant for each fit is $\tau = 2$ ms.

Figure 6.1 clearly shows a damped behavior of the Rabi oscillations. Damped oscillations are a result of dephasing and decoherence and are typically characterized by two time constants, τ_1 and τ_2, with a total dephasing rate

$$\frac{1}{\tau} = \frac{1}{\tau_1} + \frac{1}{\tau_2}. \tag{6.6}$$

For our Rabi oscillations measurements, we observe a total dephasing rate of $\tau = 2$ ms. τ_1 is known as longitudinal relaxation and is caused by population decay of atoms in the upper hyperfine ground state decaying to the lower hyperfine ground state. This population decay destroys the coherence of the electronic wave function, which interrupts the Rabi flopping process. The $F = 2$ hyperfine state is metastable, and therefore the decay of the Rabi oscillations we see is not due to τ_1 relaxation.

The dipole trapping beams are detuned from the D_1 and D_2 transitions on the order of -100 THz. The photon scattering rate is [53]

$$\begin{aligned}\Gamma_{sc} = \frac{\pi c^2}{2\hbar}\Bigg[&\frac{1}{\omega_{01}^3}\left(\frac{\Gamma_1^2}{(\omega_{01}-\omega_{laser})^2} + \frac{\Gamma_1^2}{(\omega_{01}+\omega_{laser})^2}\right) \\ &+ \frac{2}{\omega_{02}^3}\left(\frac{\Gamma_2^2}{(\omega_{02}-\omega_{laser})^2} + \frac{\Gamma_2^2}{(\omega_{02}+\omega_{laser})^2}\right)\Bigg] I(x,y,z).\end{aligned} \tag{6.7}$$

For our experiment, $\Gamma_{sc} = 2$ Hz, so dephasing due to spontaneous scattering of trap photons is considered negligible in the damping process. The crossed dipole trap lifetime $\tau_{CDT} = 1.5$ s is three orders of magnitude larger than the damping time constant, so it only contributes to the damping process at the 0.2% level.

τ_2 is known as transverse relaxation, caused by dephasing processes. Dephasing can occur due to inhomogeneities in the trapping environment. In the presence of trapping laser beams with Gaussian intensity profiles, atoms in a dipole trap experience a continuous distribution of light shifts. This leads to Rabi oscillations of different frequencies, weighted in amplitude with the density distribution of the dipole trap, that interfere [83, 84, 85]. As atoms undergo collisions with other trapped atoms, therefore, their wave functions lose their

phase relationship. Figure 6.2 shows the results of a Mathematica simulation of Rabi oscillations of 1000 randomly sampled atoms in an inhomogeneous light field for the conditions of our dipole trap at the light shifted resonance frequency. The dephasing time constant in the simulation, $\tau = 1.9$ ms, is in very good agreement with the results of the experiment, which had $\tau = 2$ ms. Therefore, we can say that the primary cause of dephasing is due to inhomogeneities in the trapping environment. This dephasing may be reversed, however, and coherence reestablished, using a spin-echo technique described in [86].

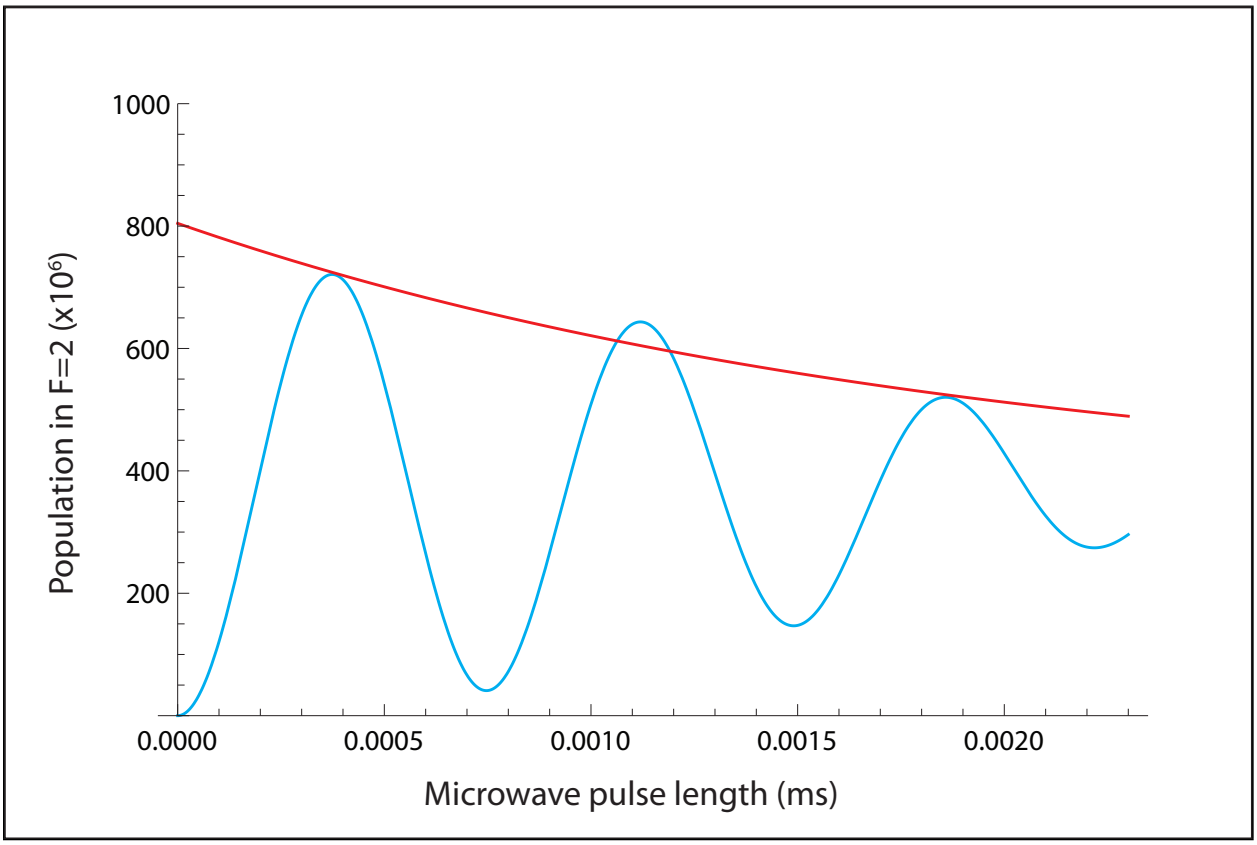

Figure 6.2: This is a plot showing the results of a Mathematica simulation in which Rabi oscillations were calculated for 1000 atoms in an inhomogeneous light field matching the conditions in our crossed dipole trap. The blue curve shows the Rabi oscillations for the light shifted resonance frequency $2\pi \cdot 6.8346890$ GHz, including damping due to dephasing. The red curve is an exponential decay with time constant $\tau = 1.9$ ms, which matches the decay of the simulation.

7 Conclusion

In an effort to probe new regimes of quantum magnetism in spinor BECs, we have devised a method to quickly and efficiently load an optical dipole trap in an ultralow magnetic field environment. The dipole trap forms in the central region of a five layer mu-metal magnetic shield which has been shown to reduce background magnetic fields to the 1 μG level. The dipole trap is unique in that we can dynamically change the crossing angle of the dipole trapping beams to induce changing trap anisotropies. This will permit future searches for the formation of spontaneous spin textures in the ultralow field regime, which have been predicted but never observed. Due to the nature of our dipole trap design, atoms trapped in the non-intersecting regions of the dipole trapping beams heat the atoms trapped in the intersecting region and prevent evaporation to a BEC. We propose, however, that this problem can be circumvented by tilting the dipole trapping beams into a vertical orientation to allow gravity to tilt the trapping potential and provide an escape channel for trapped atoms. We have developed a new technique to digitally count an optical beat note frequency between two laser sources to make an agile, high performance electronic frequency offset lock. In the future, the techniques and schemes we have established should provide a solid framework for anyone developing quantum simulation experiments using spinor BECs in ultralow magnetic fields.

Bibliography

[1] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen(De), and U. Sen. Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond. *Advances in Physics*, 56(2):243–379, 2007.

[2] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, 80:885–964, Jul 2008.

[3] M. Lewenstein, A. Sanpera, and V. Ahufinger. *Ultracold Atoms in Optical Lattices: Simulating quantum many-body systems*. Oxford University Press, 2012.

[4] T. Ho. Spinor bose condensates in optical traps. *Phys. Rev. Lett.*, 81:742–745, Jul 1998.

[5] C. K. Law, H. Pu, and N. P. Bigelow. Quantum spins mixing in spinor bose-einstein condensates. *Phys. Rev. Lett.*, 81:5257–5261, Dec 1998.

[6] H. Pu, C. K. Law, S. Raghavan, J. H. Eberly, and N. P. Bigelow. Spin-mixing dynamics of a spinor bose-einstein condensate. *Phys. Rev. A*, 60:1463–1470, Aug 1999.

[7] E. Demler and F. Zhou. Spinor bosonic atoms in optical lattices: Symmetry breaking and fractionalization. *Phys. Rev. Lett.*, 88:163001, Apr 2002.

[8] K. Gross, C. P. Search, H. Pu, W. Zhang, and P. Meystre. Magnetism in a lattice of spinor bose-einstein condensates. *Phys. Rev. A*, 66:033603, Sep 2002.

[9] A. Imambekov, M. Lukin, and E. Demler. Spin-exchange interactions of spin-one bosons in optical lattices: Singlet, nematic, and dimerized phases. *Phys. Rev. A*, 68:063602, Dec 2003.

[10] K. V. Krutitsky and R. Graham. Spin-1 bosons with coupled ground states in optical lattices. *Phys. Rev. A*, 70:063610, Dec 2004.

[11] Z. Xie, W. Zhang, S. T. Chui, and W. M. Liu. Magnetic solitons of spinor bose-einstein condensates in an optical lattice. *Phys. Rev. A*, 69:053609, May 2004.

[12] S. Tsuchiya, S. Kurihara, and T. Kimura. Superfluid˘mott insulator transition of spin-1 bosons in an optical lattice. *Phys. Rev. A*, 70:043628, Oct 2004.

[13] A. Imambekov, M. Lukin, and E. Demler. Magnetization plateaus for spin-one bosons in optical lattices: Stern-gerlach experiments with strongly correlated atoms. *Phys. Rev. Lett.*, 93:120405, Sep 2004.

[14] J. J. García-Ripoll, M. A. Martin-Delgado, and J. I. Cirac. Implementation of spin hamiltonians in optical lattices. *Phys. Rev. Lett.*, 93:250405, Dec 2004.

[15] S. Jin, J. Hou, B. Xie, L. Tian, and M. Ge. Superfluid˘mott-insulator transition of spin-2 cold bosons in an optical lattice in a magnetic field. *Phys. Rev. A*, 70:023605, Aug 2004.

[16] K. V. Krutitsky, M. Timmer, and R. Graham. First- and second-order superfluid˘mott-insulator phase transitions of spin-1 bosons with coupled ground states in optical lattices. *Phys. Rev. A*, 71:033623, Mar 2005.

[17] M. Rizzi, D. Rossini, G. De Chiara, S. Montangero, and R. Fazio. Phase diagram of spin-1 bosons on one-dimensional lattices. *Phys. Rev. Lett.*, 95:240404, Dec 2005.

[18] D. Jaksch and P. Zoller. The cold atom hubbard toolbox. *Annals of Physics*, 315(1):52–79, 2005.

[19] M. Chang, Q. Qin, W. Zhang, L. You, and M. S. Chapman. Coherent spinor dynamics in a spin-1 bose condensate. *Nature Physics*, 1(2):111–116, 2005.

[20] F. Zhou and G. W. Semenoff. Quantum insulating states of $f = 2$ cold atoms in optical lattices. *Phys. Rev. Lett.*, 97:180411, Nov 2006.

[21] A. Widera, F. Gerbier, S. Fölling, T. Gericke, O. Mandel, and I. Bloch. Precision measurement of spin-dependent interaction strengths for spin-1 and spin-2 87rb atoms. *New Journal of Physics*, 8(8):152, 2006.

[22] Y. Kawaguchi, H. Saito, and M. Ueda. Spontaneous circulation in ground-state spinor dipolar bose-einstein condensates. *Phys. Rev. Lett.*, 97:130404, Sep 2006.

[23] S. Yi and H. Pu. Spontaneous spin textures in dipolar spinor condensates. *Phys. Rev. Lett.*, 97:020401, Jul 2006.

[24] K. Eckert, L. Zawitkowski, M. J. Leskinen, A. Sanpera, and M. Lewenstein. Ultracold atomic bose and fermi spinor gases in optical lattices. *New Journal of Physics*, 9(5):133, 2007.

[25] R. V. Pai, K. Sheshadri, and R. Pandit. Phases and transitions in the spin-1 bose-hubbard model: Systematics of a mean-field theory. *Phys. Rev. B*, 77:014503, Jan 2008.

[26] J. Simon, W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner. Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472(7343):307–312, 2011.

[27] M. Ueda. Bose gases with nonzero spin. *Annual Review of Condensed Matter Physics*, 3(1):263–283, 2012.

[28] J. Struck, C. Ölschläger, R. Le Targat, P. Soltan-Panahi, A. Eckardt, M. Lewenstein, P. Windpassinger, and K. Sengstock. Quantum simulation of frustrated classical magnetism in triangular optical lattices. *Science*, 333(6045):996–999, 2011.

[29] P. Soltan-Panahi, J. Struck, P. Hauke, A. Bick, W. Plenkers, G. Meineke, C. Becker, P. Windpassinger, M. Lewenstein, and K. Sengstock. Multi-component quantum gases in spin-dependent hexagonal lattices. *Nature Physics*, 7(5):434–440, 2011.

[30] G. Jo, J. Guzman, C. K. Thomas, P. Hosur, A. Vishwanath, and D. M. Stamper-Kurn. Ultracold atoms in a tunable optical kagome lattice. *Phys. Rev. Lett.*, 108:045305, Jan 2012.

[31] D. Clément, A. F. Varòn, J. A. Retter, L. Sanchez-Palencia, A. Aspect, and P. Bouyer. Experimental study of the transport of coherent interacting matter-waves in a 1d random potential induced by laser speckle. *New Journal of Physics*, 8(8):165, 2006.

[32] L. Fallani, C. Fort, and M. Inguscio. Bose–einstein condensates in disordered potentials. *Advances In Atomic, Molecular, and Optical Physics*, 56:119–160, 2008.

[33] M. White, M. Pasienski, D. McKay, S. Q. Zhou, D. Ceperley, and B. DeMarco. Strongly interacting bosons in a disordered optical lattice. *Phys. Rev. Lett.*, 102:055301, Feb 2009.

[34] M. Pasienski, D. McKay, M. White, and B. DeMarco. A disordered insulator in an optical lattice. *Nature Physics*, 2010.

[35] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of bose-einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, 1995.

[36] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969–3973, Nov 1995.

[37] M. D. Barrett, J. A. Sauer, and M. S. Chapman. All-optical formation of an atomic bose-einstein condensate. *Phys. Rev. Lett.*, 87:010404, Jun 2001.

[38] William D Phillips and Harold Metcalf. Laser deceleration of an atomic beam. *Physical Review Letters*, 48(9):596, 1982.

[39] H. J. Metcalf and P. Van der Straten. *Laser cooling and trapping*. Springer, 1999.

[40] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Simplified system for creating a bose–einstein condensate. *Journal of low temperature physics*, 132(5-6):309–367, 2003.

[41] D. A. Steck. Rubidium 87 d line data, 2010.

[42] A. M. Steane, M. Chowdhury, and C. J. Foot. Radiation force in the magneto-optical trap. *JOSA B*, 9(12):2142–2158, 1992.

[43] E. Dimova, O. Morizot, G. Stern, C. L. G. Alzar, A. Fioretti, V. Lorent, D. Comparat, H. Perrin, and P. Pillet. Continuous transfer and laser guiding between two cold atom traps. *The European Physical Journal D*, 42(2):299–308, 2007.

[44] S. Bali, K. M. O'Hara, M. E. Gehm, S. R. Granade, and J. E. Thomas. Quantum-diffractive background gas collisions in atom-trap heating and loss. *Physical Review A*, 60(1):R29, 1999.

[45] J. J. Arlt, O. Marago, S. Webster, S. Hopkins, and C. J. Foot. A pyramidal magneto-optical trap as a source of slow atoms. *Optics communications*, 157(1):303–309, 1998.

[46] J. M. Kohel, J. Ramirez-Serrano, R. J. Thompson, L. Maleki, J. L. Bliss, and K. G. Libbrecht. Generation of an intense cold-atom beam from a pyramidal magneto-optical trap: experiment and simulation. *JOSA B*, 20(6):1161–1168, 2003.

[47] K. H. Kim, K. I. Lee, H. R. Noh, W. Jhe, N. Kwon, and M. Ohtsu. Cold atomic beam produced by a conical mirror funnel. *Physical Review A*, 64(1):013402, 2001.

[48] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman. Low-velocity intense source of atoms from a magneto-optical trap. *Physical review letters*, 77(16):3331, 1996.

[49] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven. Two-dimensional magneto-optical trap as a source of slow atoms. *Physical Review A*, 58(5):3891, 1998.

[50] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. B. Ovchinnikov, and T. Pfau. Intense source of cold rb atoms from a pure two-dimensional magneto-optical trap. *Physical Review A*, 66(2):023410, 2002.

[51] R. S. Conroy, Y. Xiao, M. Vengalattore, W. Rooijakkers, and M. Prentiss. Compact, robust source of cold atoms for efficient loading of a magnetic guide. *Optics communications*, 226(1):259–266, 2003.

[52] W. Wohlleben, F. Chevy, K. Madison, and J. Dalibard. An atom faucet. *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics*, 15(2):237–244, 2001.

[53] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. *Advances In Atomic, Molecular, and Optical Physics*, 42:95–170, 2000.

[54] P. Bouyer, T. L. Gustavson, K. G. Haritos, and M. A. Kasevich. Microwave signal generation with optical injection locking. *Optics Letters*, 21(18):1502–1504, Sep 1996.

[55] M. Kasevich and S. Chu. Laser cooling below a photon recoil with 3-level atoms. *Phys. Rev. Lett.*, 69(12):1741–1744, Sep 1992.

[56] E. A. Cummings, M. S. Hicken, and S. D. Bergeson. Demonstration of a 1-w injection-locked continuous-wave titanium : sapphire laser. *Applied Optics*, 41(36):7583–7587, Dec 2002.

[57] V. Boyer, L. J. Lising, S. L. Rolston, and W. D. Phillips. Deeply subrecoil two-dimensional raman cooling. *Phys. Rev. A*, 70(4):1–8, Oct 2004.

[58] T. Stace, A. N. Luiten, and R. P. Kovacich. Laser offset-frequency locking using a frequency-to-voltage converter. *Measurement Science and Technology*, 9(9):1635–1637, 1998.

[59] U. Schunemann, H. Engler, R. Grimm, M. Weidemuller, and M. Zielonkowski. Simple scheme for tunable frequency offset locking of two lasers. *Rev. Sci. Instrum.*, 70(1, Part 1):242–243, Jan 1999.

[60] M. Zhu and J. L. Hall. Stabilization of optical-phase frequency of a laser system - application to a commercial dye-laser with an external stabilizer. *Journal Of The Optical Society Of America B-Optical Physics*, 10(5):802–816, May 1993.

[61] G. Santarelli, A. Clairon, S. N. Lea, and G. M. Tino. Heterodyne optical phase-locking of extended-cavity semiconductor-lasers at 9 ghz. *Optics Communications*, 104(4-6):339–344, Jan 1994.

[62] L. Ricci, M. Weidemuller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. Konig, and T. W. Hansch. A compact grating-stabilized diode-laser system for atomic physics. *Optics Communications*, 117(5-6):541–549, Jun 1995.

[63] H. Muller, S. W. Chiow, Q. Long, and S. Chu. Phase-locked, low-noise, frequency agile titanium: sapphire lasers for simultaneous atom interferometers. *Optics Letters*, 31(2):202–204, Jan 2006.

[64] M. O. Mewes, M. R. Andrews, N. J. van Druten, D. M. Kurn, D. S. Durfee, and W. Ketterle. Bose-einstein condensation in a tightly confining dc magnetic trap. *Phys. Rev. Lett.*, 77(3):416–419, 1996.

[65] M. Drewsen, P. Laurent, A. Nadir, G. Santarelli, A. Clairon, Y. Castin, D. Grison, and C. Salomon. Investigation of sub-doppler cooling effects in a cesium magneto-optical trap. *Applied Physics B*, 59(3):283–298, 1994.

[66] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard. Phase-space density in the magneto-optical trap. *Phys. Rev. A*, 52:1423–1440, Aug 1995.

[67] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard. High densities of cold atoms in a dark spontaneous-force optical trap. *Physical review letters*, 70(15):2253, 1993.

[68] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. Behavior of atoms in a compressed magneto-optical trap. *JOSA B*, 11(8):1332–1335, 1994.

[69] LD Landau and EM Lifshitz. Classical mechanics, 1960.

[70] T. A. Savard, K. M. O'hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Physical Review A*, 56(2):R1095, 1997.

[71] M. E. Gehm, K. M. O'hara, T. A. Savard, and J. E. Thomas. Dynamics of noise-induced heating in atom traps. *Physical Review A*, 58(5):3914, 1998.

[72] S. Friebel, C. D'andrea, J. Walz, M. Weitz, and T. W. Hänsch. Co2-laser optical lattice with cold rubidium atoms. *Physical Review A*, 57:20, 1998.

[73] D. Sesko, T. Walker, C. Monroe, A. Gallagher, and C. Wieman. Collisional losses from a light-force atom trap. *Physical review letters*, 63(9):961–964, 1989.

[74] C. D. Wallace, T. P. Dinneen, K.-Y. N. Tan, T. T. Grove, and P. L. Gould. Isotopic difference in trap loss collisions of laser cooled rubidium atoms. *Phys. Rev. Lett.*, 69:897–900, Aug 1992.

[75] D. Comparat, A. Fioretti, G. Stern, E. Dimova, B. L. Tolra, and P. Pillet. Optimized production of large bose-einstein condensates. *Physical Review A*, 73(4):043410, 2006.

[76] Y. Takasu, K. Maki, K. Komori, T. Takano, K. Honda, M. Kumakura, T. Yabuzaki, and Y. Takahashi. Spin-singlet bose-einstein condensation of two-electron atoms. *Physical review letters*, 91(4):040404, 2003.

[77] K. J. Arnold and M. D. Barrett. All-optical bose–einstein condensation in a 1.06¡ i¿ μ¡/i¿ m dipole trap. *Optics Communications*, 284(13):3288–3291, 2011.

[78] M. Williams and C. Fertig. A multipartite model of evaporative cooling of atomic clouds in optical dipole traps. Unpublished manuscript, 2013.

[79] A. Mager. Magnetic shielding efficiencies of cylindrical shells with axis parallel to the field. *Journal of Applied Physics*, 39(3):1914–1914, 1968.

[80] D. U. Gubser, S. A. Wolf, and J. E. Cox. Shielding of longitudinal magnetic fields with thin, closely spaced, concentric cylinders of high permeability material. *Review of Scientific Instruments*, 50:751, 1979.

[81] E. A. Burt and C. R. Ekstrom. Optimal three-layer cylindrical magnetic shield sets for scientific applications. *Rev. Sci. Instrum.*, 73(7):2699–2704, 2002.

[82] S. Linzen, T. L. Robertson, T. Hime, B. L. T. Plourde, P. A. Reichardt, and J. Clarke. Low-noise computer-controlled current source for quantum coherence experiments. *Review of scientific instruments*, 75(8):2541–2544, 2004.

[83] J. C. Camparo, R. P. Frueholz, and C. H. Volk. Inhomogeneous light shift in alkali-metal atoms. *Physical Review A*, 27(4):1914, 1983.

[84] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Analysis of dephasing mechanisms in a standing-wave dipole trap. *Physical Review A*, 72(2):023406, 2005.

[85] P. J. Windpassinger, D. Oblak, U. B. Hoff, J. Appel, N. Kjærgaard, and E. S. Polzik. Inhomogeneous light shift effects on atomic quantum state evolution in non-destructive measurements. *New Journal of Physics*, 10(5):053032, 2008.

[86] M. F. Andersen, A. Kaplan, and N. Davidson. Echo spectroscopy and quantum stability of trapped atoms. *Phys. Rev. Lett.*, 90(2):023001, 2003.