2011-12-14

Development of a Strontium-87 Ion Interferometer

Christopher Joseph Erickson
Brigham Young University - Provo

Development of a $^{87}Sr^{+}$ Interferometer

Christopher Joseph Erickson

A dissertation submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Dallin S. Durfee, Chair
Scott D. Bergeson
Michael J. Ware
Justin B. Peatross
Jean-Francois S. Van Huele

Department of Physics and Astronomy

Brigham Young University

December 2011

ABSTRACT

Development of a $^{87}Sr^{+}$ Interferometer

Christopher Joseph Erickson
Department of Physics and Astronomy
Doctor of Philosophy

I present the construction of a low-velocity intense source (LVIS) of laser-cooled neutral strontium using permanent ring magnets. The LVIS consists of a magneto-optical trap from which cold strontium is extracted in a well-collimated beam. I also present the development and implementation of a full suite of low-noise, high-bandwidth laser control electronics including a microcontroller unit. This microcontroller remotely controls and monitors the current driver, temperature controller, and PID lock circuit for each diode laser simultaneously. The current driver output is accurate to within 2 μA and repeatable to with a few nA. The noise spectral density of the current driver hits a floor of 10^{-10} A/$\sqrt{\text{Hz}}$ at $\sim$ 50 Hz and has a modulation bandwidth of $\sim$ 50 MHz. The PID lock-circuit includes a scan-balancing option that we have used to scan an AR coated laser diode $\sim$ 30 GHz mode-hop free. I describe the construction of an 80 mW frequency doubled 461 nm laser system using PPKTP for cooling and trapping neutral strontium in the LVIS. The LVIS, the electronics systems, and the 461 nm laser system represent major milestones on the way to producing a matter-wave interferometer using ^{87}Sr ions. The interferometer is based on an optical Raman transition between the hyperfine ground states of $^{87}Sr^{+}$. The ions will be produced by exciting the strontium LVIS beam to an auto-ionizing state in the continuum. In the interferometer two $\pi/2$ pulses of light and one π pulse will be delivered to the ions to split and recombine their wave functions. I present calculations of the predicted sensitivity and a discussion of the possible applications. I present a method for locking a 407.8 nm laser to the 5s $^2S_{1/2}$ to 5p $^2P_{3/2}$ strontium ion transition in a neutral vapor. I present calculations for the necessary vacuum levels for the experiment and describe the preparation and assembly of the vacuum apparatus. The major vacuum system consists of two connected elastomer sealed chambers: one at 10^{-7} Torr and the other at 10^{-10} Torr separated by a region of low conductance. I present a Sr vapor cell constructed from standard CF fittings that allows the strontium to be heated to $\sim 730^{\circ}$ C, which can also be run as a thermal beam. I present a method for protecting the viewports on small-form alkali-earth vapor cells using lead or indium foil during the evaporation of oxide layers. Finally, I report on the current status of the experiment as well as detail future work on the apparatus.

Keywords: strontium, matter-wave, interferometry, LVIS, charged particle, cold atoms, magneto-optical trap, MOT, laser cooling

ACKNOWLEDGMENTS

Many friendships developed out of my graduate experience that sweetened the experience overall and made the long hours bearable in the short term. In particular the fond memories of Dan Christensen, Aaron Bennett, and David McKenna will always be associated with this work. I will forever treasure the friendship with Will Somers that was forged in the first years of graduate school. Overall, the many lunches, Friday morning golf outings, and co-commiseration sessions with Michael Besselman absolutely preserved my sanity throughout this degree.

I am also appreciative of the National Institute of Standards and Technology, the National Science Foundation, BYU Graduate Studies, the BYU College of Physical and Mathematical Sciences, the Science, Mathematics, and Research for Transformation (SMART) Fellowship, and the Center of Excellence (RDL/AFRL) at Kirtland Air Force Base for the financial support that made this work possible.

I would like to acknowledge several of the people who have helped and inspired me throughout my education. Scott Daniels, John Ellsworth, Wes Lifferth, and Nan Ah You have provided great support. Wes's skill and prowess were integral to designing and manufacturing several critical parts of the apparatus. John's advice and knowledge was necessary in solving many of the problems with the vacuum chambers and he was always quick to provide needed parts and service. Scott's foresight and efforts in acquiring electronics equipment for the department has saved the lab untold tens of thousands of dollars. Nan's efficiency and managerial skills have also helped me acquire needed supplies and equipment. The friendship, wit, stories, and example of all four will be forever remembered.

Among the faculty who have left an especially significant impression are Dr. Allred, Dr. Davis, Dr. Hirschmann, Dr. Nielsen, Dr. Bergeson, Dr. Van Huele, Dr. Peatross, Dr. Ware, and Dr. Durfee. My educational experiences have indeed been influenced by many people. All of these men exhibit an interest and excitement in physics that I have found contagious. In addition to their

example, these people have also contributed a great deal of support. David Allred taught me where to find and how to interpret phase diagrams detailing the interactions of multiple elements, and was always willing to help me with those efforts. Scott Bergeson was never too busy to discuss and help solve the problems I encountered and was always willing to loan me equipment and resources. Jean-Francois Van Huele has been constantly attentive to my progress and supportive throughout the degree as have been the other members of my committee whose advise and encouragement proved invaluable.

Above all in the physics department, I owe many thanks to my advisor, Dallin Durfee. With great patience (not an exaggeration), he has helped me to come to a better understanding of the physical principles that govern this world. He stands as a great example of moral resolution, devotion to higher principles, and selfless action. He has become a great friend and a mentor.

I would like to thank my wonderful family. My family and I have received great support from our extended family during this time in graduate school. My in-laws have a great love of knowledge and learning and have been extremely encouraging and helpful. My parents are great examples of hard work who instilled in me a hunger for education, and taught me to ever pursue my dreams. The support from siblings and others of my education is more appreciated than they will ever know.

More important than any of the afore mentioned is my wonderful wife Katie. Without her support, love, patience, and forgiveness I would be an incompetent wreck. Why she has stayed by my side through such a demanding education, is beyond my comprehension. She is loving, kind, charitable, and understanding. She is my hero, and I love her. While the pregnancies of Benjamin, Kallie, and Andrew have created much stress, their love and adoration has brought me great peace and joy. The pregnancy of soon-to-be-born Zachary at the end of this degree has been particularly difficult, but it has taught me much about self control and selflessness. I love my family and am grateful for their love and the lessons they have taught me.

Finally, it is God to whom all the credit for this work is really due. He has sustained me through my whole life and blessed me with strength and faculty enough to achieve all my accomplishments.

All that I do is dedicated to Him and for His purposes. I rely on Him for all my needs, and it is through the Atonement of His Son that He has made possible my highest hopes of returning to His presence with my family. His hand has been present throughout this work and will be present throughout any work I shall do hereafter.

Contents

List of Figures

List of Tables

Foreword

While as a thesis this is mainly a technical work, I cannot help but add what we might call a human interest section right at the beginning. This is somewhat unconventional, but since nearly all of my research experience has been gained in Dallin Durfee's lab I feel it is appropriate to give a short history.

I first joined the lab as an undergraduate double majoring in physics and music. The music was to chase the dream that every teenager with a small amount of musical inclination has of making it in that industry. The physics was because I was still smart enough to pursue a passion that could realistically put bread on the table, so to speak. Shortly before my declaration as a physics major, the BYU Physics and Astronomy Department had begun requiring undergraduates to produce a senior thesis. I am grateful to the department for the wisdom in implementing that requirement. By committing to pursue two of the most credit-hour intensive programs on campus it is unlikely that I would have found the time to involve myself in undergraduate research on my own. As a result, I would never have gained many of the rich experimental physics opportunities I have enjoyed.

At that time the lab was really the combined efforts of Scott Bergeson and Dallin Durfee and the students involved often referred to it as the "Burfeson" lab, working on both professor's projects. Despite my inexperience, I was hired at the standard department wage for 10 hours a week to cut my teeth on a variety of electronics projects. Thanks to this experience I was taught the basics

of soldering and the functions of passive and active circuits long before I encountered them in any lab class. My first substantial project, and the subject of my senior thesis, was to construct a high-temperature calcium vapor cell to perform spectroscopy on the partially forbidden intercombination line [1, 2]. This project taught me that even the most simple-sounding project could be complex and time consuming but also rewarding as the work was honored with a spot in the undergraduate session of the American Physical Society's (APS) Division of Atomic, Molecular, and Optical Physics (DAMOP) annual meeting in 2005. This project also taught me the basics of experimental vacuum techniques and instilled in me a fascination with vacuum systems.

Near the end of this project I was preparing for and thinking about graduate schools and programs. There was great deal of excitement and adventure in the prospect of working in any one of the several AMO graduate programs at other institutions. However, I began to feel strongly about staying at BYU. Despite the risk of acquiring a stigma from completing all of my education at one university and the fact that BYU is primarily an undergraduate institution, after much prayer and meditation is seemed clear to me that "this [BYU] is the place". Yes, I just used a corny LDS reference, but what can I say? I have spent a decade at BYU. So I sent some inquiring, if not misguided, applications to schools located near where my wife's family lived before embarking in the fall of 2005 on a course of study towards Ph.D. candidacy at BYU .

On the tail end of finishing my undergraduate degree I joined Brian Neyenhuis and Greg Doermann in working on a narrow line-width laser for use with the calcium vapor cell I had developed. All of this work was part of the larger project of building a calcium interferometer based on the $4s^2\ {}^1S_0$ to the $4s4p\ {}^3P_1$ intercombination line. In retrospect, much of the work done on this laser system led to the development of critical parts of the ion system described in this thesis.

I took over the project of designing and constructing a laser-diode current driver [3, 4]. In the end we developed a driver that exhibits lower noise and greater stability than any of the commercially available systems and lab built designs that are publicly available. This project was the

beginning of what we envisioned as an array of lab equipment for diode laser control that could be simultaneously monitored and managed from a single micro-processing unit. While I built and characterized the current driver, Marshall Van Zjill deciphered a several hundred page manual in order to design a circuit for and program a USB programmable PIC controller, and Brian Neyenhuis started work on a proportional-integral-differential (PID) lock circuit.

One of the main design points for these circuits was taking advantage of the low noise and high speed of certain integrated circuit chips (ICs) that were only available in surface mount packages. In the beginning, we were forced to hunch over the tiny features of custom printed circuit boards (PCBs) with a magnifying glass, a pair of tweezers, and some very fine low melting-point solder. The work was tedious, time consuming, and frustrating. Debugging these circuits often lead to destruction of the PCB.

I was desperate to find a better way to do this and happened across the website of a 30-something robotic hobbyist who claimed to solder surface mount chips in the toaster oven of his landlady. Encouraged by this, I purchased a toaster oven and, to my delight, found that you could indeed solder an entire PCB with hundreds of components with very little mess by using a water based no-flux solder paste. A description of our method for working with surface mount parts as well as tips and tricks can be found in Appendix B.

While Marshall and I continued development of our laser diode control devices, Brian began developing a specialized detector to pair with the PID circuit. Most laser systems in precision metrology require a feedback system to keep their frequency stable for driving a specific atomic resonance. Commonly these systems consist of modular radio frequency (RF) components that are used to create a feedback or error signal which the PID controller uses to keep the laser diode at a specific frequency. Brian's circuit condensed all of the RF components onto one PCB creating a low-noise and compact homodyne photodiode detector.

After the initial implementation of each of these four circuits I took over their development and

optimization. Additionally, we added a laser-diode temperature controller, which also interfaced with our micro-controller, as well as a piezo driver. Final credit for the homodyne detector goes to Aaron Bennet. James Archibald developed the first version of the piezo driver, I developed the subsequent versions, and Michael Hermansen has been doing the work of characterizing our latest version. This suite of circuits constitute the basis for all of the laser systems critical to the work described in thesis and their development and implementation represent a very significant amount of my time as a graduate student. The latest designs and performance are detailed in Chapter 5.

While I was in my first year as a graduate student, Brian Neyenhuis was preparing to take the Physics GRE. A practice question arose concerning Coulomb's Law which piqued Brian's interest. He found that very little improvement had been made in testing the theory in recent history. He began considering the development of scientific instruments and whether or not matter-wave interferometry might provide a new breakthrough in this area. Together with Dallin, they worked out that a significant improvement could be made by doing the measurement interferometrically with a charged particle [5]. Shortly thereafter, Dan Christensen calculated that the stray and fringing fields expected in a physical apparatus would not be a limiting factor in the experiment [6]. Dan also numerically verified the results of Brian's analytical calculation. On the basis of this work Dallin applied for and won a NIST Precision Measurement Grant to begin the project.

In the meantime, I continued constructing the neutral calcium interferometer. The many challenges and frustrations of that project are detailed in [7]. Dan Christensen worked many hours with me to help solve and commiserate over the problems of that apparatus. I have fond memories of that experiment and the many afternoons working on it with the lights off to make work with the laser easier. Dan would always play Barry White or sports radio from Philadelphia in the background. It seems like every time Dallin entered the room if Barry White was not crooning, Philly sports radio was playing an add for a hot bar girl contest or some other racy sounding event.

We were in a race to finish the calcium interferometer when the underground lab was shut down

for renovations. Most of the project was shelved while we worked on what could fit onto one optics table in a shared room the Chemistry Department lent us. By this time I had received a SMART fellowship. The fellowship had come with the stipulation that I convert my Ph.D. track into a Masters program and use what I had already accomplished to finish within one year. Gratefully, enough work was completed on the calcium interferometer to earn a Masters degree and I was granted a second fellowship to pursue a Ph.D.

Work on the calcium interferometer was often painstaking and tedious. A lot of patience was required to build the kHz linewidth 657 nm laser that the project required. Much of the work assembling the apparatus consisted of working within the confines of the heavy aluminum box that housed the laser or hunched over in the small space between the massive vacuum chamber and its heavy lid aligning optics. Occasionally, moments of excitement occurred like when Dan Christensen sealed a stainless steel Conflat nipple containing hydrochloric acid and a large lump of calcium. Work progressed on the interferometer until the only piece lacking was a frequency doubled laser to serve as a calcium fluorescence probe.

As funding for the calcium experiment dwindled, another grant was awarded for the ion interferometer. Dallin was reluctant to switch my project thanks to a less than desirable memory of his graduate advisor doing the same to him. After undergraduates destroyed three tapered amplifiers for the calcium project and there was no money left to replace them, he gave in and I began working on the strontium ion interferometer. While it had been the desire of my heart for a while to work on the ion interferometer, I have to admit a guilty pleasure in occasionally mentioning, within Dallin's hearing, how he changed my project mid-degree.

I began work on this project very near the beginning of the construction phase. The only existing parts at the time consisted of a custom-built vacuum chamber to house a magneto-optical trap (MOT), and the first basic components for building a frequency-doubled laser system to trap neutral strontium. From there on I designed and constructed most of the apparatus as it exists

today. Notable exceptions include the basis for a laser locking electronic circuit built by Jarom Jackson, a 407.8 nm diode laser built by Dean Anderson, calculation of the MOT and Zeeman slower fields by Kelvin Blaser, a 405.2 nm laser system built by Dallin, and characterization and wavelength selection of stock 405 nm diodes by Mark Cunningham. I also initially designed and purchased parts for a laser system to drive Raman transitions in the interferometer that was given to James Archibald to complete for his Master's Thesis. I would also like to thank Daniel Merrill, Aaron Bennet, Michael Hermansen, Ben Francis, Stuart Harper, and Mary Lyon for their efforts.

This dissertation is focused on the work I have put towards building a strontium ion interferometer. If you had told me a year ago that I would still not have a fully functioning interferometer by the writing of this thesis I would have laughed at you. Instead, there are a couple of specific pieces of lab equipment that are laughing at me. I now have a much more profound understanding of the seemingly outrageous price tags of commercial laser systems. In retrospect, given the internship requirements of my fellowship and the demands of two high-risk pregnancies that have occurred since I began work on the ion interferometer, it is no small miracle what has been accomplished. I am grateful to have been part of this work and for the many lessons and much experience working on this project has afforded me.

I performed all experimental procedures in this thesis unless specified below:

Chapter 3, Section 3.2: The mounts for the ring magnets discussed in this section were designed by Dallin Durfee and machined by Wes Lifferth.

Chapter 3, Section 3.3: The magnetic coils described in this section were designed, constructed, and mounted by Kelvin Blaser.

Chapter 3, Figures 3.4, 3.7, and 3.8: The magnetic fields described in these figures for trapping, slowing, and accelerating, and the corresponding velocity of the Sr LVIS were calculated by Kelvin Blaser and Dallin Durfee.

Chapter 3, Figure 3.6: The interferograms that I took of the elliptical mirrors shown in Figure 3.6 were computationally analyzed by Mary Lyon.

Chapter 4, Section 4.2: The wavelength selection of the diodes in this section was performed by Mark Cunningham.

Chapter 4, Section 4.3: The laser system described in this section was built by Dallin Durfee.

Chapter 4, Section 4.4: The laser system described in this section was built by Dean Anderson.

Chapter 4, Section 4.6.1: The initial circuit used to drive the AOM in the 461 nm laser system lock described in this section was built by Jarom Jackson.

Chapter 5, Figures 5.4 and 5.5: The noise measurement described in part (b) of both figures was performed by Daylin Troxel.

Chapter 5, Section 5.4: The performance of the homodyne detector in this section was characterized by Aaron Bennett.

Chapter 9, Figure 9.6: The window clamp shown in Figure 9.6 was designed by Dallin Durfee.

Chapter 9, Figures 9.10 and 9.11: The HV electrodes and their fields shown in the figures were designed and calculated by Dallin Durfee.

Chapter 1

Introduction

This work sets forth in detail the construction of a slow strontium source and the work done towards using it to realize a strontium ion interferometer. To best appreciate the design and function of the ion source and interferometer, we begin with a review of some basics about matter wave interferometers in general.

Any interferometer is, simply put, a device through which a passing wave is coherently split along two distinct paths and then recombined. Phase differences between the two paths, or arms, of the interferometer are revealed by constructive or destructive interference fringes in the exiting signal. Such devices can measure to within a fraction of the wavelength used any influence that can affect the phase of that wave.

With that in mind, a matter-wave interferometer may best be described by simple analogy with an optical interferometer, such as a Michelson-Morley interferometer [8]. In an optical interferometer mirrors and beam splitters are used to split and recombine light. Likewise, a matter-wave interferometer is a device that spatially splits and recombines the quantum wave function of a particle with mass.

Light interferometers are used in multiple applications including velocimetry, navigation, surface profiling, stress measurements, and even the search for gravity waves [9–14]. Matter-wave

interferometers are often used in fundamental studies and applications including but not limited to inertial force sensing, precision metrology of fundamental constants, and local gravity measurements [15–17].

1.1 Ion Interferometry

1.1.1 Motivation for Building an Ion Interferometer

In certain applications atom interferometers have several advantages over other interferometers. Since atoms at thermal velocities have extremely small wavelengths, atom interferometers have the potential to perform many types of experiments with greater precision than optical interferometers. While continuous-wave coherent photon sources can have wavelengths as low as a couple hundred nano-meters, atomic wavelengths can be on the order of pico-meters for atoms traveling at hundreds of meters per second.

Since atoms have internal structure, inertial mass, and magnetic moments, atom interferometers can perform experiments that cannot be done with optical interferometers whose wave functions are not affected by as many parameters. Examples include measuring time variance in the fine structure constant [18–24], time dilation and Kennedy-Thorndike experiments [25–28], and Lorentz violation studies [29, 30].

In addition to exploiting the natural properties of neutral atoms, atoms can also be ionized to give them an overall charge. This allows for the realization of matter-wave interferometers that are also capable of making precision measurements of electromagnetic fields. An ion interferometer is useful for characterizing materials and performing fundamental tests of electromagnetic theory. While no device like the one presented in this work have ever been built, charged particle interferometers have existed for some time in the form of electron interferometers [31, 32].

Electron interferometers, however, suffer from several limitations as they require physical grat-

ings or bi-prisms and fairly high-energy electrons in order to have good signal-to-noise ratios. Electron interferometers are also based off of thermal sources of electrons, which limits the coherence length of the beam. Only one ion interferometer has been previously realized but it also suffered from the same problems as electron interferometers as it was based on the use of a thermal ion source and electron bi-prisms [33].

This work describes the development of a charged particle interferometer of unprecedented sensitivity. Ultimately, the ion interferometer described here will be used to search for a photon rest mass and violations of Coulomb's Law as proposed in [5,6]. In building this ion interferometer we are opening up a new sub-field in precision metrology for conducting fundamental studies of electro-magnetic theory, precision measurement of free-space electric fields, and characterization of material properties.

1.1.2 A Short History of Matter-wave Interferometry

Atom interferometry has its roots in the doctoral work of Louis de Broglie and the idea that matter not only has a wavelength but can exhibit wavelike behavior [34]. The wavelike behavior of atoms was first observed in the early part of the last century by Esterman and Stern who diffracted a beam of sodium atoms off of a sodium chloride crystal [35]. The first atom interferometer was developed by Norman Ramsey as an atomic clock. Even though the clock's atom beam never split spatially, its wave function was nonetheless split and recombined [36–39].

The first true separated beam atom interferometer was by Carnal and Mlynek, who used narrow slits to split and recombine metastable helium atoms [40]. This was followed almost immediately by Pritchard, who used micro-fabricated diffraction gratings to realize a sodium interferometer [41]. Shortly thereafter, Bordé published a report of an optical Ramsey spectrometer using beams of light to split and recombine a beam of thermal atoms [42].

Since then, matter-wave interferometers of varying design have been developed that have con-

tributed to the field of precision measurement [43]. In addition to the previously mentioned applications these devices have performed a variety of fundamental experiments including but not limited to: measuring the photon recoil of an atom [44], measuring the local acceleration g due to gravity [16, 45], measuring the Earth's gravity gradient [46, 47], making precision rotation measurements [48–51], and measuring the electric polarizability of atoms [52].

1.1.3 Overview of Basic Matter-wave Interferometer Schemes

The principal concerns with any interferometer are sensitivity, accuracy, and signal-to-noise ratio (SNR). Sensitivity is the principal motivation for building a matter-wave interferometer in the first place. The sensitivity of an interferometer depends strongly on the wavelength and the length of time a particle spends in the interferometer. It can also depend on the separation of the arms and the size of the enclosed area depending on the purpose and setup of the interferometer.

This can mean that the sensitivity can often be enhanced by cooling the atom down in order to send it through the interferometer more slowly, even though this increases the wavelength. When measuring fields, a slower atom will spend more time interacting with the field resulting in a greater phase shift in the interferometer. Enlarging the enclosed area of the interferometer also increases sensitivity when measuring fields. The area enclosed can be increased by lengthening the interferometer, which also could increase the time an atom spends interacting with a field. Causing the interferometer's arms to split at a greater angle will also increase the enclosed area resulting in a greater potential difference between the arms.

Different ways of splitting and recombining an atom's wave function have been used including slits [40], diffraction gratings [41], traveling light waves [42], the diffraction of atoms off of standing light waves [53, 54], and Raman transitions [46, 47, 49, 51] (which we use in this work).

Despite the size of atomic wavelengths, matter-waves can be diffracted off of macroscopic slits and physical gratings. In these setups multiple diffraction orders are generated at each stage of the

interferometer of which only a few contribute to the fringes generated at the output, which much be resolved spatially. The contrast of the output fringes in this setup is subject to the mechanical stability of the gratings as well as the ability to spatially detect the fringes.

Improvements can be made to this setup by replacing the physical grating with an optical one. A laser, usually red detuned relative to a resonant transition of the species used, is retro-reflected across the path of an atomic beam to create a standing wave. The atoms will diffract off of this optical grating, which has a period of half the optical wavelength used. The benefit is that the atoms can be kept far from physical surfaces that can generate unwanted fields. The stability is usually no better than physical gratings and multiple diffraction orders are still generated at each grating.

A matter-wave interferometer may also be realized by directly driving an atom between two energy states with a resonant laser as shown in Figure 1.1. An atom first receives a pulse of light tailored to place it in an equal quantum superposition of two energy states. A second pulse drives each half of the superposition into the opposite state and a third pulse recombines the superposition. Depending on the phase difference between the two paths of the interferometer the exiting atoms are measured to be in either one state or the other.

The benefit of this scheme is that nearly all the atoms contribute to the signal and the phase of the interferometer is encoded in the internal state of the atoms. This is much easier to detect with high SNRs than spatially resolving interference fringes. Instead, the interferometer phase can be detected with fluorescence probe. The difficulty lies in choosing an atomic transition with a long enough lifetime so that atoms can traverse the interferometer before there is significant chance of spontaneous emission. This usually requires building an extremely stable laser in order to drive a narrow linewidth transition.

A Raman transition can be used to realize the interferometer. This setup is similar to Figure 1.1 except that each pulse is formed by two counter-propagating laser beams. The atom is excited on a

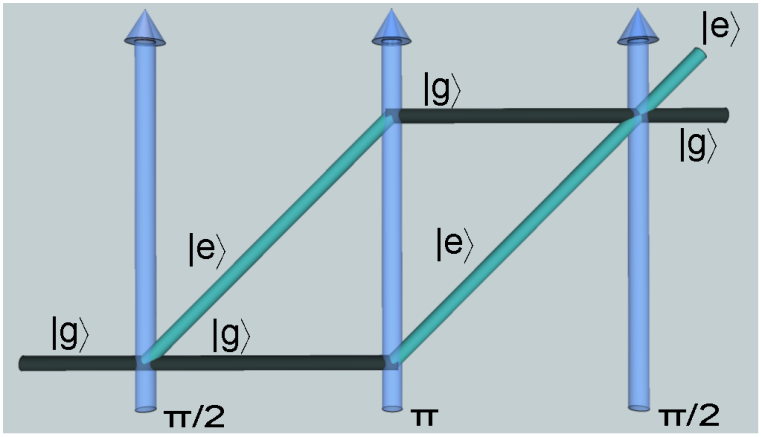

Figure 1.1 Drawing of an optical interferometer scheme for matter-waves. A beam of atoms starts out in the ground state on the left hand side of the figure. As the atoms cross a laser beam they receives a $\pi/2$ pulse of light and are placed in a quantum superposition of ground and excited states. The $\pi/2$ pulse imparts momentum to the excited state part of the superposition, which recoils away from the original beam. Two more beam crossings serve to bring the arms of the interferometer back together (π pulse), and then mix them together (last $\pi/2$ pulse).

transition with an extremely long lifetime by absorption of a photon from one beam and stimulated emission of the photon into the other beam. The two lasers are detuned from an intermediate energy level that usually has a relatively broad linewidth. An additional benefit of this design is that the excited state part of the superposition receives a two photon momentum kick, doubling the separation of the interferometer's arms relative to a single photon transition. Also, much less stringent requirements are placed on the laser system. With a Raman transition the stability of the relative phase and frequency of the lasers are of greatest importance, which is much easier to control than the overall stability and phase.

1.1.4 Description of Ion Interferometer Scheme

In our design, a coherent stream of ions with a very narrow velocity distribution is created by ionizing laser cooled strontium atoms and optically pumping them into the same hyperfine ground state. The ions then cross three counter-propagating sets of continuous-wave (CW) laser beams.

The CW laser beams effectively deliver a pulse of light to the atoms as they cross the beams just as in Figure 1.1. The intensity of the first "pulse" places the atoms in a superposition of the two hyperfine ground states of $^{87}Sr^{+}$. The second pulse then drives each part of the superposition to the opposite state and the third recombines them. With each pass through a laser beam the transverse momentum of the ions is manipulated to create two physically separate paths through the interferometer. The detected state of the atoms as they exit the interferometer depends on the relative phase of the two arms of the interferometer. A more detailed theory of the operation, function, and known physical properties of the interferometer are described in Chapter 8.

The interferometer will be able to perform a variety of fundamental and practical studies. It will be able to characterize the electrical properties of materials such as hydrogen storage materials, optical materials. The interferometer could also be used for measuring time-dependence of patch charge potentials and probing electron collisions in superconducting materials. In the future, the interferometer will be adapted to perform more fundamental studies of the currently accepted electromagnetic theory. Specifically, an ion interferometer could be used to detect deviations from Coulomb's law with a sensitivity two orders of magnitude greater than current experiments [5].

1.2 Cold Ion Source

The two major components required to construct the ion interferometer are an ion source and a laser system that can drive the Raman transition. This dissertation covers in detail the construction of the ion source, the design of the Raman laser system, and the construction and development of all of the other systems and procedures necessary for the realizing the interferometer.

1.2.1 Basic Requirements of the Ion Source

In order for the interferometer to function properly the ions traversing the apparatus must have certain characteristics. As seen in Figure 1.1, CW lasers are used to deliver pulses of light to the atoms by directing the lasers perpendicular to the path of the atoms. Even though the lasers are CW, since the atoms pass quickly into and out of the beam they effectively see the light field as a pulse. In order for this "light pulse" to drive the ions into an equal superposition of ground and excited states the intensity and size of the beam must be carefully controlled to match the time of flight of the ions through the beam.

If the ion beam has a large velocity distribution, only a small percentage of the ions will be able to receive the correct pulse while the bulk of the atoms will be driven into a range of unequal quantum superpositions, destroying the contrast of the interferometer. The phase shift of the ion due to an electric field is also velocity dependent, causing broad velocity distributions to further reduce contrast. In order to achieve the highest contrast, the velocity distribution of the ions must be small enough so that, given the Fourier broadening of the laser, all of the ions effectively receive the same pulse.

In order to provide a mono-energetic beam of ions we start by creating a stream of laser cooled neutral atoms. This neutral atom source is based off of the low-velocity-intense source (LVIS) first developed by Carl Wieman's group [55].

1.2.2 Basic Design of the Ion Source

The LVIS is in essence a specialized magneto-optical-trap (MOT). Three counter-propagating sets of laser beams and a magnetic field serve to trap neutral atoms emitted from an effusive thermal source. In order to extract a beam of atoms from this trap a mirror with a hole in its center is placed in the vacuum to retro-reflect one of the beams. This hole effectively opens a region of the trap which atoms are free to leave through. The benefit of this configuration is that it creates

a highly collimated beam of cold atoms with a narrow longitudinal velocity distribution. Similar configurations using alkali atoms report typical longitudinal velocity distributions of a few meters per second full-width at half-max (FWHM) [55–58].

This neutral atom beam is then converted to an ion beam after it exits the trapping region. The atom beam passes between a pair of electrodes where lasers drive the atoms to an auto-ionizing state in the continuum. The potential across the electrodes can velocity tune the ion beam to any desired energy up to 20 keV.

1.2.3 Why Strontium?

The Atomic Structure of ^{87}Sr and ^{87}Sr^{+}

For many experiments the mass of the ion is not a critical parameter, as shown in [5], allowing us to use the element of our choice. Strontium was chosen for this experiment for several reasons. First, Sr has been extensively cooled and studied [59–61] and is similar in atomic structure to calcium, which we had previously used in the lab. Aside from this, ^{87}Sr is a naturally occurring isotope with ground state hyperfine structure as shown in Figure 1.2. The hyperfine ground states can be driven via a Raman transition detuned from the ion's 407.8 nm 5s $^2S_{1/2}$ to 5p $^2P_{3/2}$ resonant transition.

As seen in Figure 1.2 ^{87}Sr also has an 1 nm wide auto-ionizing transition [62] that can be reached from the 461 nm 5s^2 $^1S_0$ to 5s5p $^1P_1$ cooling transition with a 405 nm photon. This makes ionizing Sr a relatively easy process using the same laser technology already required by for the cooling and Raman transitions.

Only a small fraction of the elements have been laser cooled, and only a few of those have ions with a resonance transition that can be driven with a diode laser. The fact that strontium has been well studied and has these properties is extremely lucky. The possibility of using 405 nm diodes is particularly beneficial since they are widely available and diode lasers are a very stable laser technology that can be worked with easily in the laboratory. Additionally, 405 nm is one of

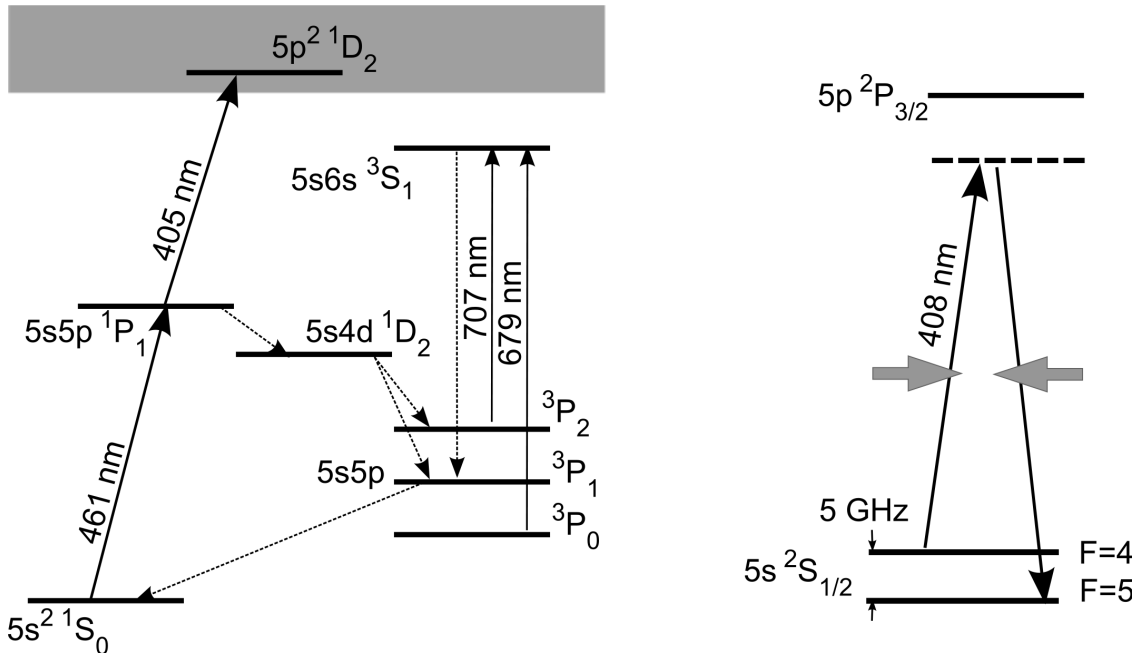

Figure 1.2 The energy level diagram for neutral ^{87}Sr is on the left and the diagram for $^{87}Sr^{+}$ is on the right. The only hyperfine structure shown is that of the ion's ground state.

the shortest available laser diode wavelengths, resulting in nearly the smallest grating period and widest ion beam separation available with this technology.

1.3 Thesis Layout

Chapter 2 reviews the basic theory behind the construction of the ion source, from trapping neutral strontium to creating a coherent, mono-energetic ion beam, while Chapter 3 details the Sr MOT and LVIS constructed. Chapter 4 describes the different laser systems built for the experiment and Chapter 5 describes the electronics and control circuits used in the experiment. Chapter 6 provides an overview of the design and operation of the Sr vapor cell used for frequency locking the laser systems. Chapter 7 details the requirements for and construction of the major vacuum systems required. Chapter 8 presents the theory, layout, and design I completed of the full ion interferometer yet to be realized. Chapter 9 presents the successes of the work to date. The appendices provide

supplemental theory, procedures, and code used to build and design this experiment.

Chapter 2

Theory of Laser Cooling and Trapping

The practice of laser cooling and trapping atoms has been extensively used in many sub-fields of atomic physics. It is a powerful technique that lends itself well to many varied applications including the cold ion source in this experiment. Despite the fact that the technology is well understood, the ion source we have constructed will be the first of its kind. Since laser cooling and trapping of neutral atoms is a critical part of the ion source, I present here a brief overview of the relevant theory.

2.1 Atom-Photon Interactions

Several good resources exist at just about all levels of understanding for gaining a more in depth knowledge of the interactions between light and matter for cooling and trapping atoms [63–66]. This overview is meant to engage those whose discipline lies outside AMO physics as well as to present important parameters specific to our apparatus.

Photons of frequency ν have an associated energy $E = h\nu$ and thereby a characteristic momentum $p = h\nu/c = \hbar k$. When an atom emits or absorbs a photon it does so in accordance with the laws of conservation of energy and momentum. For example, when an atom emits or absorbs

a photon, since the photon carries a momentum of $\hbar k$ the atom's momentum must change by the same amount in the appropriate direction. While this change in momentum for one event is very small, by absorbing (and scattering) many photons within a short time frame an atom can experience significant forces.

The length of time an atom spends in an excited state after absorbing a photon is known as the lifetime of the state. This is the characteristic time it takes to spontaneously emit the absorbed photon and the inverse is known as the linewidth Γ. For the strontium cooling transition $\Gamma/(2\pi) =$ 32 MHz and the lifetime of the state is $\tau = 4.97$ ns.

For spontaneously emitted photons there is no directional preference and they are emitted randomly over the entire solid angle. For an atom in a resonant laser beam the atom experiences an overall net force in the direction of beam propagation by absorbing photons from one direction and emitting them over all directions. The spontaneously emitted photons cause the atom to perform a random walk in the plane orthogonal to the beam axis.

The average force exerted on our strontium atom sitting in a resonant light field is given by

$$< \mathbf{F} >= Q_s \hbar \mathbf{k}, \tag{2.1}$$

where Q_s is the rate of photon scattering by the atom. For a given transition, Q_s is equal to the population in the excited state times Γ. In the rest frame of the atom, the excited state population depends on the laser beam frequency and intensity, and it can be calculated by solving the optical Bloch equations in steady state [66, 67]. The result gives an expression for the scattering rate of

$$Q_s = \frac{\Gamma}{2} \frac{I_L/I_{Sat}}{1 + I_L/I_{Sat} + (2\Delta/\Gamma)^2}, \tag{2.2}$$

where I_L is the intensity of the laser, $I_{Sat}$is the saturation intensity, and Δ is the detuning of the laser frequency $\omega_L = 2\pi\nu_L$ from the resonant frequency ω_0. The quantity I_L/I_{sat} is called the on-resonance saturation parameter S_0.

The saturation parameter can take on different definitions depending on the particular definition of I_{Sat}. Here we use

$$I_{Sat} = \frac{2\pi^2\hbar\Gamma c}{3\lambda_0^3}, \tag{2.3}$$

where λ_0 is the wavelength of the atomic transition [68]. In this formalism, the atom scatters photons at half of the maximum rate when $S_0 = 1$, effectively placing $1/4$ of the atoms in the excited state at any given time. For the strontium cooling transition I_{Sat} is equal to 42.5 mW/cm^2.

The fluorescence scattering rate has a maximum of $\Gamma/2$ for $S_0 \gg 1$ and small detuning. This means that a strontium atom in a powerful on resonance laser beam can experience a maximum acceleration of $\mathbf{a}_{max} = Q_{max}\hbar\mathbf{k}/m_{Sr} \approx 9.9 \times 10^5$ m/s^2 $= 1.01 \times 10^5$ g.

2.2 Doppler Cooling

The atom's velocity relative to the propagation axis of the laser can shift the light frequency seen by the atom due to the Doppler effect. This leads to a velocity dependent radiation force on an atom moving in a light field. If $v \ll c$ then the frequency shift due to the Doppler effect is approximately $\mathbf{k}\cdot\mathbf{v}$ leading to

$$< \mathbf{F_{Dop}} > = \hbar\mathbf{k}\frac{\Gamma}{2}\frac{S_0}{1+S_0+(2(\Delta-\mathbf{k}\cdot\mathbf{v})/\Gamma)^2}. \tag{2.4}$$

As Eq. 2.4 indicates, the detuning of a laser can be used to exert this force on atoms of a particular velocity relative to the laser beam.

The velocity range of atoms likely to be strongly affected by this force is small due to its Lorentzian lineshape. The range depends on the linewidth of the atomic transition driven and is given by the FWHM of Eq. 2.4, $|\Delta - \mathbf{k}\cdot\mathbf{v}| \leq \Gamma$. Therefore, atoms at the right velocity to shift the laser frequency into resonance are quickly pushed out of resonance to a different velocity class. For the strontium cooling transition, atoms within a velocity range of only $v_{cool} \approx 14.8$ m/s will significantly feel this force.

2.2.1 Optical Molasses

The velocity dependent force of Eq. 2.4 indicates a laser beam can provide an opposing force to an atom's motion in one direction. Let us consider the case of two counter-propagating, red-detuned laser beams. In this one dimensional case, a strontium atom moving along the light field will see the frequency of one laser beam Doppler shifted closer to resonance while the other beam will be further red-shifted. In the limit of low intensity ($S_0 \ll 1$) the resulting force can be written as

$$\begin{aligned} < \mathbf{F}_v > &= < \mathbf{F}_{Dop}(\mathbf{k}) > - < \mathbf{F}_{Dop}(-\mathbf{k}) > \\ &= \frac{\hbar \mathbf{k} \Gamma S_0}{2} \left[\frac{1}{1 + \left(\frac{2(\Delta - \mathbf{k}\cdot\mathbf{v})}{\Gamma} \right)^2} - \frac{1}{1 + \left(\frac{2(\Delta + \mathbf{k}\cdot\mathbf{v})}{\Gamma} \right)^2} \right]. \end{aligned} \tag{2.5}$$

In the limit of small $\mathbf{k} \cdot \mathbf{v}$ this force can be approximated as

$$< \mathbf{F}_v > \approx \frac{8\hbar k^2 S_0 \Delta}{\Gamma(1 + (2\Delta/\Gamma)^2)^2} \mathbf{v}. \tag{2.6}$$

Figure 2.1 shows the velocity dependence of $< \mathbf{F}_v >$ from Eq. 2.5 in one dimension for strontium. The optimal detuning for the laser beam so as to maximize the damping force occurs when $\Delta = \Gamma/2$ [66, 69]. With the correct configuration and detuning, three pairs of counter-propagating laser beams can significantly damp the atom's motion in three dimensions: a setup known as optical molasses.

2.2.2 Cooling Limits

Equation 2.6 allows us to rewrite the Doppler force as a damping force $< F_v > = -\alpha v$ for small velocities. The average rate at which the atom loses kinetic energy is then $\dot{E}_{cool} = < F_v > v = -\alpha v^2$. An atom's kinetic energy would damp to zero if not for the discrete nature of the force. The atom is cooled by a series of small momentum transfers as the atom absorbs and emits photons. This means that ultimately the atom's momentum fluctuates about an average value with the fluctuations tending to heat the atom.

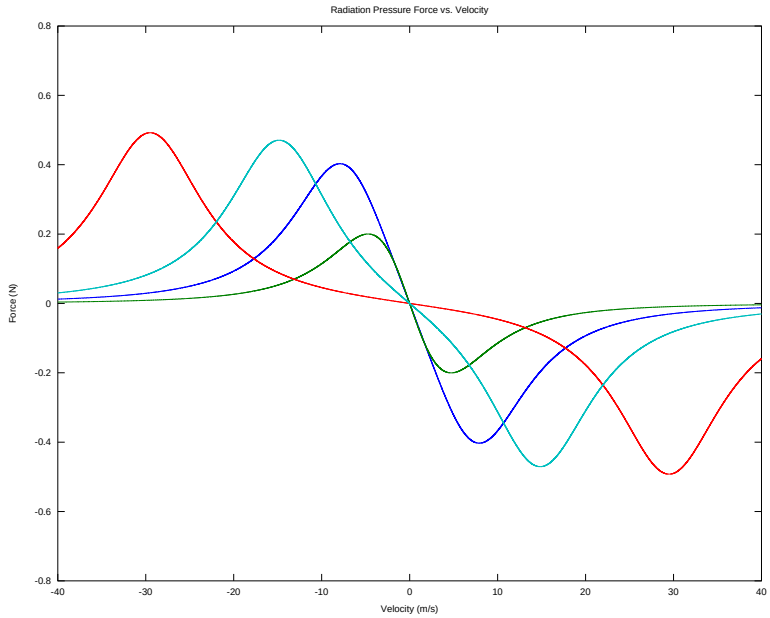

Figure 2.1 The radiation pressure force for different values of Δ.The red line is $\Delta = 2\Gamma$, the light blue line is $\Delta = \Gamma$, the dark blue line is $\Delta = \Gamma/2$, and the green line is $\Delta = \Gamma/6$.

This heating can be thought of by considering that as the atom is cooled both beams become equally detuned. Therefore, the probability of the atom absorbing a photon from either of the beams is equal. In the limit of low intensity each absorption and emission event is uncorrelated with any others resulting in a random walk.

As the atom performs this random walk the mean square of the momentum increases with time. From basic statistical theory, after N momentum steps of size $\hbar k$ the mean square momentum is $< p^2 >= N\hbar^2k^2$. The mean square change in momentum per unit time is then given by the rate of absorbing and emitting photons Q_S:

$$\frac{d}{dt} < p^2 >= \hbar^2k^2\frac{\Gamma}{2}\frac{S_0}{1+S_0+(\frac{2\Delta}{\Gamma})^2}. \tag{2.7}$$

The average rate of heating is related to mean square of the change in momentum by $\dot{E}_{heat} = \frac{d}{dt} < p^2 > /2M$.

In steady state the average rates of cooling $\dot{E}_{cool}$ and heating $\dot{E}_{heat}$ can be equated,

$$\frac{\frac{d}{dt} < p^2 >}{2M} = \alpha < v^2 > . \tag{2.8}$$

A characteristic kinetic temperature for the atoms in the molasses is then given by the mean square velocity at which the heating rate equals the cooling rate. In our one dimensional analysis there is only one degree of freedom so that

$$\frac{M < v^2 >}{2} = \frac{k_B T}{2}, \tag{2.9}$$

where k_B is Boltzman's constant and the temperature T is in Kelvin. Again taking $S_0 \ll 1$, we can write the temperature for the molasses as

$$T = \frac{\hbar\Gamma}{4k_B}\left(\frac{\Gamma}{2\Delta} + \frac{2\Delta}{\Gamma}\right). \tag{2.10}$$

It can be shown that in a symmetrical three dimensional case the temperature is the same as Eq. 2.10 [70]. This equation yields a minimum when $\Delta = \Gamma/2$ of

$$T_D = \frac{\hbar\Gamma}{2k_B}. \tag{2.11}$$

The Doppler limit is typically several hundred μK. For strontium $T_D = 122\ \mu$K with an associated velocity of $v_D = 10.7$ cm/s.

2.3 Trapping

The random nature of scattering photons in steady state ensures that eventually the atom will wander out of the optical molasses. In order to trap the atoms the energy level structure of the atoms is manipulated by the Zeeman effect to provide a position dependent force in addition to the velocity dependent force.

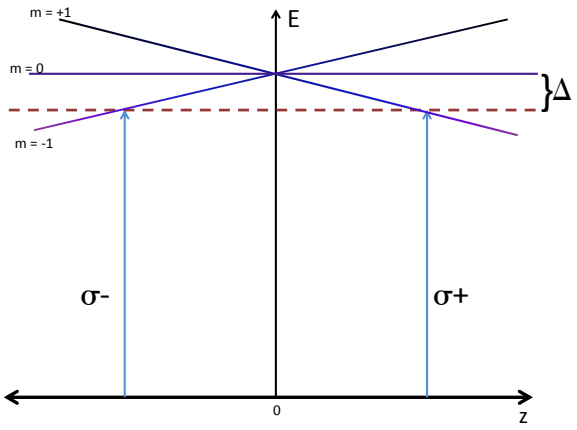

Figure 2.2 Simple diagram of the Zeeman shift that creates the trapping force in a MOT in one dimension. The magnetic field is zero at $z = 0$. The increasing field to either side shifts opposite magnetic sublevels into resonance with the laser beam of appropriate polarization creating a position dependent force.

In the one dimensional case, a magnetic field is generated along the propagation direction of two counter-propagating cooling beams, which are given circular polarizations of opposite helicity. The polarizations then act on the $J = 0$ to $J = 1$ cooling transition in the presence of a magnetic field that splits the magnetic sublevels of the transition as seen in Figure 2.2. The magnetic field is designed to have a zero at the center of the trap and increases in strength moving outward. An atom displaced from the origin finds that one of its magnetic sublevels is Zeeman shifted closer to resonance with the beam that would push it back to the center, while the sublevel that can be driven by the other beam is further shifted out of resonance.

The energy shift experienced by the magnetic sublevels is related to the strength of the magnetic field by

$$\Delta E_{shift} = m_j g_e B \mu_B \tag{2.12}$$

where μ_B is the Bohr magneton, g_e is the Landé factor, and m is the shifted sublevel. The associated frequency shift is $\Delta E_{shift}/\hbar$. This setup is known as a magneto-optical trap (MOT), and the total

force can be written

$$< \mathbf{F}_{MOT} >= \frac{\Gamma\hbar\mathbf{k}}{2}\left[\frac{S_0}{1+S_0+(\frac{2(\Delta-\mathbf{k}\cdot\mathbf{v}+\mu_\mathbf{B}\mathbf{B}/\hbar)}{\Gamma})^2}-\frac{S_0}{1+S_0+(\frac{2(\Delta+\mathbf{k}\cdot\mathbf{v}-\mu_\mathbf{B}\mathbf{B}/\hbar)}{\Gamma})^2}\right]. \tag{2.13}$$

2.4 Zeeman Slowing

A MOT is typically loaded from a thermal source of atoms with a wide velocity distribution. As can be seen from Figure 2.1 atoms that are moving very fast feel little to no effect from laser beams that are detuned by $\Gamma/2$. Unfortunately, most of the atoms in a thermal atom beam are fast enough that they pass right through the MOT setup just described. As such, the number of atoms loaded into the MOT can be vastly increased by using similar radiation pressure forces to slow a thermal atom beam as it approaches the trapping region.

As discussed in section 2.2, a red-detuned laser beam counter-propagating with the atoms can quickly push a small range of atoms to a slower velocity class and out of resonance. In order to continuously cool atoms, a magnetic field along the axis of propagation is tailored such that the Zeeman shift equals the changing Doppler shift as the atoms slows.

The necessary magnetic field gradient can be related to the change in velocity experienced by the atom. An atom will come into resonance with the laser beam when the Zeeman shift equals the difference between the detuning and the Doppler shift:

$$\frac{g_e B\mu_B}{\hbar} = \mathbf{k}\cdot\mathbf{v}_{atom} - \Delta.$$

If the atoms are defined to be traveling in the $\hat{z}$ direction, then $\mathbf{k}$ is in the $-\hat{z}$ direction so that

$$\mathbf{k}\cdot\mathbf{v}_{atom} = -\frac{2\pi}{\lambda}v_{atom}$$

$$\Rightarrow v_{atom} = -\frac{\lambda}{2\pi}\left(\Delta+\frac{g_e B\mu_B}{\hbar}\right). \tag{2.14}$$

In our slower the magnetic field starts near zero at the atom source and increases to a field maximum B_{max}. In order for the atoms to have a velocity near zero at the end of the slower, the detuning is set to be

$$\Delta = -\frac{g_e B_{max} \mu_B}{\hbar}.$$

In this configuration, the maximum velocity of atoms that can be slowed is

$$v_{max} = -\frac{\lambda}{2\pi}\Delta.$$

The acceleration experienced by the atom $a_{atom} = \frac{dv}{dt} = \frac{dv}{dz}\frac{dz}{dt} = v_{atom}\frac{dv_{atom}}{dz}$ can be calculated in terms of the magnetic field gradient.

$$\frac{dv_{atom}}{dz} = -\frac{\lambda}{2\pi}\frac{g_e \mu_B}{\hbar}\frac{dB}{dz} \tag{2.15}$$

$$\begin{aligned} a_{atom} = v_{atom}\frac{dv_{atom}}{dz} &= \left(\frac{\lambda}{2\pi}\right)^2 \frac{g_e \mu_B}{\hbar}\frac{dB}{dz}\left(\frac{g_e B_{max}\mu_B}{\hbar} + \Delta\right) \\ &= \left(\frac{\lambda}{2\pi}\right)^2 \frac{g_e \mu_B}{\hbar}\frac{dB}{dz}\left(\frac{g_e B_{max}\mu_B}{\hbar} - \frac{v_{max}2\pi}{\lambda}\right) \end{aligned} \tag{2.16}$$

The field must be engineered so that the magnitude of the acceleration given above does not exceed the maximum acceleration that the light force can produce given in Eq. 2.4. Specifically, $a_{atom} \leq a_{max}$ where $a_{max} = Q_{max}\hbar k/m_{Sr}$. This places an upper limit on the magnetic field gradient that can usefully be generated.

Chapter 3

^{87}Sr MOT and LVIS

3.1 Overview

Neutral strontium based MOTs have already been created in several different groups [71–75]. The design and function of our apparatus differs significantly from these previous experiments as the purpose is to produce a beam of laser cooled strontium rather than study the trap itself. This chapter explains the technique and parameters for creating of a low-velocity intense source (LVIS) of strontium from a MOT [55].

A cutaway drawing of the LVIS operating in its vacuum chamber is shown in Figure 3.1. All of the MOT beams and a slowing beam are derived from the output of the 461 nm laser system described in section 4.1. The laser is frequency locked to the correct detuning for cooling ^{87}Sr as described in section 4.6, and the slowing light is further detuned by an acousto-optic modulator (AOM).

A functional drawing of the optics setup for the MOT/LVIS is shown in Figure 3.2 while a picture of the actual setup is shown in Figure 3.3. To create the trap and slower the laser is divided into four equal beams with the use of $\lambda/2$-waveplates and polarizing beam cubes. Three of the

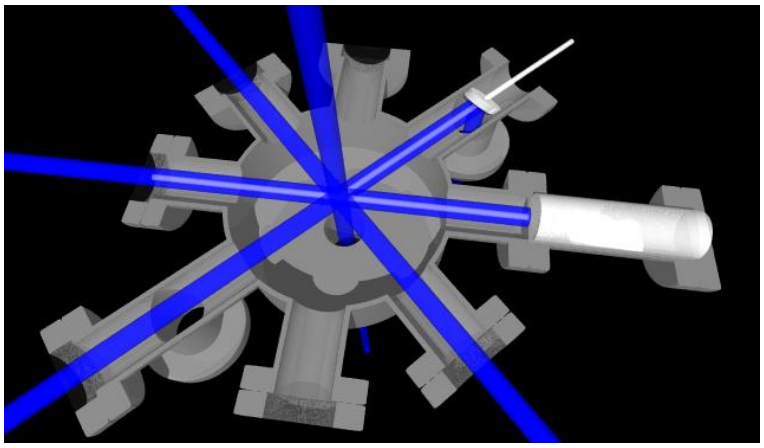

Figure 3.1 Illustration of a functioning LVIS with the vacuum chamber cut away. The only optic shown is the elliptical mirror that sits inside the vacuum chamber.

beams are sent through the chamber along orthogonal directions: two in the horizontal plane of the setup and one in the vertical. These beams are collimated to $\sim$ .4 cm beam waist before passing into the vacuum chamber through a $\lambda/4$-waveplate. Two of the beams exit on the opposite side of the chamber and through a second $\lambda/4$-waveplate. They are then retro-reflected to create circularly polarized counter-propagating beams whose electric fields oscillate in the opposite direction.

The third beam reflects off of an elliptical mirror inside the vacuum before exiting and retro-reflecting through a $\lambda/4$-waveplate. The mirror has a hole drilled through its center, which casts a shadow in the center of the retro-reflected beam. For atoms within this shadow the forces are no longer balanced and atoms are pushed out of the trap and through the mirror by the incident laser to create the LVIS.

The slower beam is red-shifted by a 400 MHz AOM, expanded, and then focused through a heated viewport onto the thermal Sr source in the MOT chamber. The viewport is heated to prevent atoms from condensing on the glass and blocking the slowing beam.

The magnetic field for the MOT is provided by two fixed N42 neodymium ring magnets. The magnets are currently mounted flush to the top and bottom surfaces of the chamber with their "north" poles facing each other to create the desired field. The magnetic field for the slowing beam

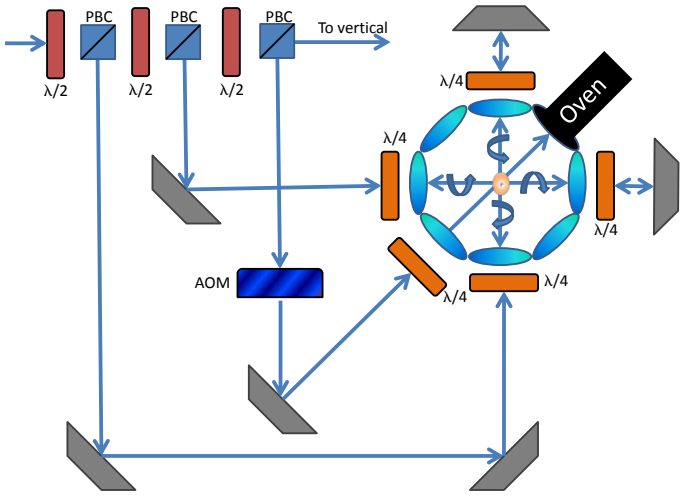

Figure 3.2 Functional drawing for the optics setup of a MOT. The beams for the horizontal axis are created by using $\lambda/4$ waveplates and retro-reflection to create the trap in the plane of the drawing. The beams and optics for the vertical axis are not shown but are the same as for the horizontal. A slower beam is created by further red-shifting a beam from the trapping light with an AOM and aligning it with the thermal source. Note that this drawing does not include any magnetic fields or the mirror with a hole in it to create an LVIS.

is created by two coils of wire that encircle the thermal Sr source. These coils together with the ring magnets create the necessary magnetic field gradient for slowing the atoms as they approach the trap.

3.2 Generating the Magnetic Fields

The strontium MOTs described in [71–75] generate magnetic field gradients of $\sim 40-70$ G/cm using anti-Helmholtz coils. While this is a typical setup, anti-Helmholtz coils require a current source as well as thermal management and for strontium the coils can be bulky. The benefit is that the magnetic field can be manipulated with a knob and even turned off all together. Since we are not performing studies requiring constant manipulation of the magnet field, we use a simpler design based off of permanent neodymium ring magnets. The magnets are placed opposing each

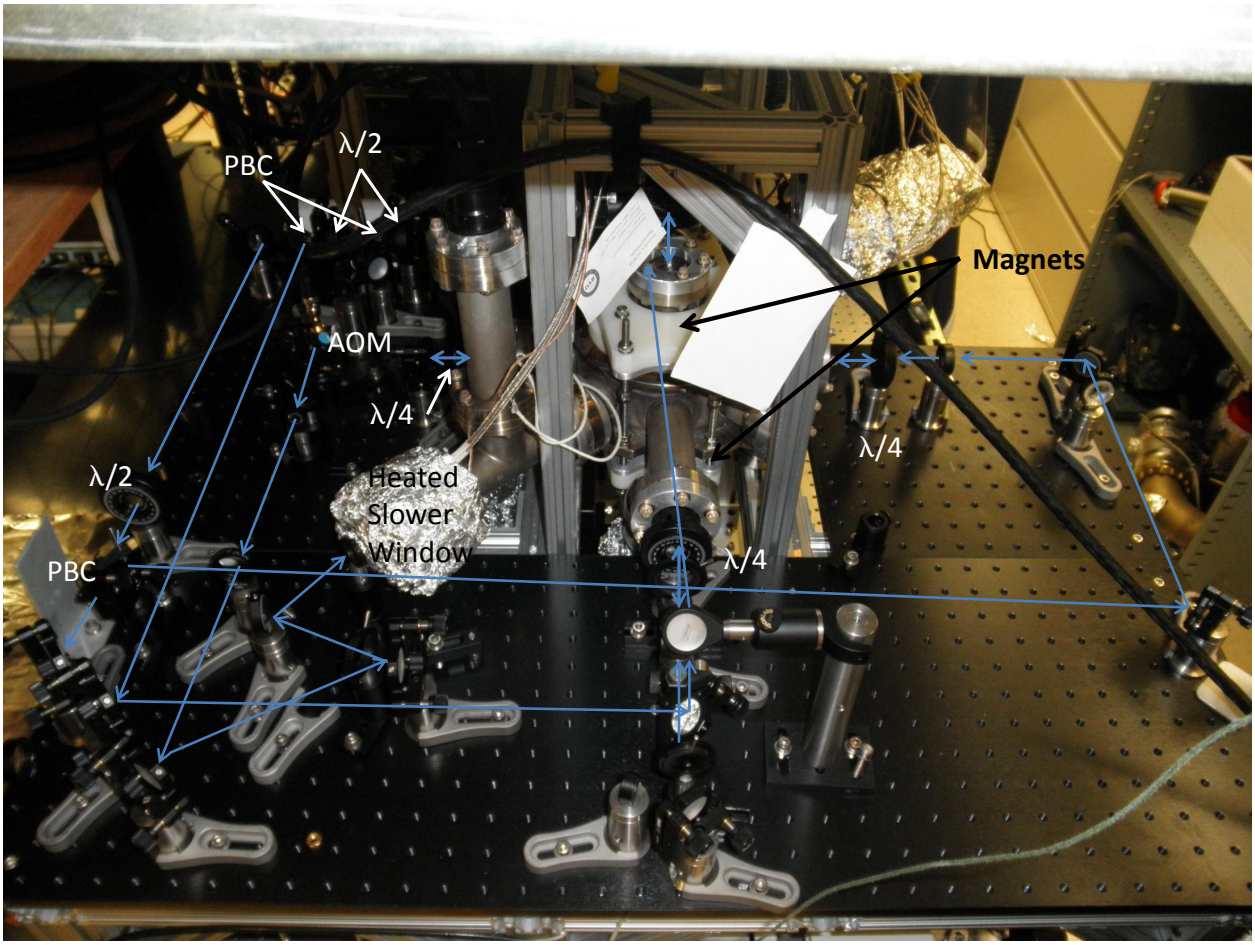

Figure 3.3 Picture of the actual optics setup for the MOT/LVIS with overlaid beams. Two optical breadboards have been mounted vertically for the optics in that dimension and are not seen.

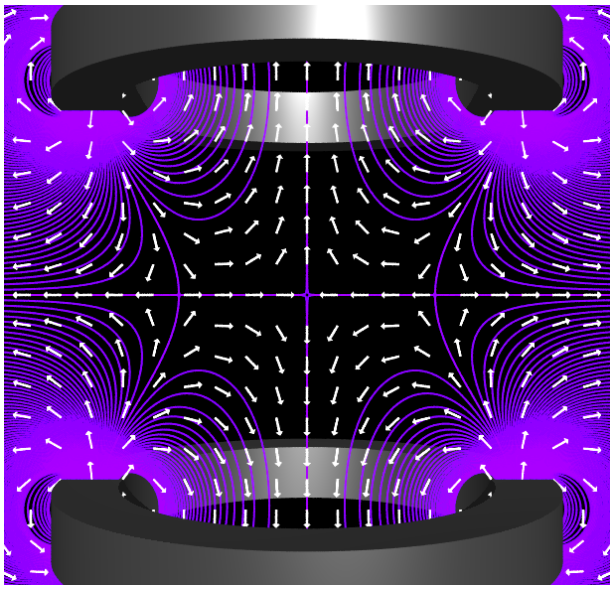

Figure 3.4 Magnetic field lines for the MOT. Note that at a distance from the center near the inner radius of the magnets the field changes sign. With current carrying coils in an anti-Helmholtz configuration the field dies off without changing sign.

other to create the appropriate field zero and gradient.

The ring magnets can be thought of simply as two current density loops. Assuming that the magnetization of the entire ring is uniform, then the current density $J_M = \nabla \times \mathbf{B} = 0$ everywhere except on the surface of the magnets. Additionally, tiny loops of current across the top and bottom surfaces of the magnet will cancel out except at the edges of the surface. Only at the inner and outer diameters of the ring does the J_M have a non-zero value and the magnets can be thought of as two concentric ribbons of current with opposite sign. The magnetic field can then be calculated via the Biot-Savart Law:

$$\mathbf{B} = \iint \frac{\mu_0}{4\pi} \frac{(\mathbf{J}dV) \times \mathbf{r}}{r^3}. \tag{3.1}$$

Figure 3.4 shows the calculated magnetic field generated by our permanent magnets in this setup.

The downside of the ring magnets is that to adjust the field gradient the magnets separation must be changed while keeping each of them equidistant from the desired position of the trap center. We calculated that the separation between the two magnets should be around 8 cm to create

a magnetic field gradient of ~ 50 G/m. The magnets have an outer diameter of 4 in., an inner diameter of 3 in., and are 0.5 in. thick.

3.3 Zeeman Slower

As seen by the field lines in Figure 3.4 and the one dimensional plot of the magnetic field in Figure 3.7, the ring magnets do not create the exact same magnetic field as anti-Helmholtz coils. Instead of the magnetic field simply fading away far from the trapping region, the ring magnets produce another zero point crossing and maximum before dying off. The slowing laser beam is detuned to use this second field gradient for Zeeman slowing the atoms.

As is, the gradient of the field is too steep for the maximum deceleration of the atoms as discussed in 2.4. In order to create a usable field gradient a small set of coils has been placed concentric to the axis of the thermal atom beam. The coils are made from 10 gauge magnet wire and consist of one 50 turn coil 12.5 cm from the trap center and another 25 turn coil 17.2 cm from the center. Each coil has a mean radius of 3.8 cm. The coils are designed to produce the correct field gradient with a current of 9 A.

Unfortunately, over time our slower contributes less and less to the number of atoms captured by the MOT. The slowing window is a standard glass viewport, which the strontium chemisorbs into at higher temperatures (see Appendix D). A sapphire viewport has been purchased to replace it the next time the MOT chamber is vented and serviced.

The field used for slowing has a maximum of ~ 200 G. This corresponds to slowing atoms with up to a maximum velocity of ~ 300 m/s down to the capture range for the MOT. When the slowing light is working optimally as many as 46.89% of the atoms can be slowed to the capture velocity. If the slowing beam is completely blocked, only 7.38×10^{-4}% of the atoms can be slowed to the capture velocity.

Thanks to the robust and forgiving nature of a MOT setup, the slower is actually not an essential part of the trapping. Even the first MOTs reported capture of atoms from the background gasses and residual vapors in vacuum [76, 77]. As it turns out, our trapping scheme will capture enough atoms from the tail end of the velocity distribution given a high enough density of atoms in the thermal beam.

3.4 Low Velocity Intense Source (LVIS)

Once a MOT has been established it is deceptively simple in theory to extract an extremely well collimated atomic beam. The LVIS technology was first pioneered by Lu et. al. at JILA [55].

The atomic beam is created by effectively punching a hole in the trap. This is done by drilling a hole in the center of a mirror used to retro-reflect one of the trapping beams. The shadow cast by the hole provides a region where atoms see an unbalanced force down the axis of the reflected beam. Atoms with an appreciable transverse velocity will re-enter the light field at the edge of the shadow and can be recycled before passing through the mirror, creating a highly collimated beam of atoms.

3.4.1 Elliptical Mirror

The mirror with the hole must be mounted inside the vacuum chamber so that the beam of atoms remains accessible and useful across the length of the chamber. Retro-reflecting directly off of this mirror is not ideal as it would require a waveplate with a hole in the vacuum as well. Careful setting and alignment of both would be required before evacuating the chamber and there would be no way to adjust them afterward.

In our setup a fixed elliptical mirror with a hole drilled at 45º to its plane is mounted within the vacuum chamber for the MOT as seen in Figure 3.1. The diagram of the mount used for mirror

is shown in Figure 7.8 of Chapter 7. The mirror directs the beam out a port on the chamber to be retro-reflected and polarization adjusted outside the vacuum.

Our mirror is designed to have the same circular projected area of a 1 in. optic when viewed at 45º. The mirror surface is enhanced aluminum rather than a dielectric coating in order to preserve polarization.

To create this mirror we drilled a hole through a NT32-132 enhanced aluminum elliptical mirror from Edmund Optics. The prospect of drilling a hole through the mirror raised serious concerns about preserving the surface flatness and quality of the optic. We were able to successfully develop a procedure for drilling holes in mirrors without degrading the reflectivity and flatness as described below.

Quality and Specifications

The trapping and cooling of the atoms depends on the shape of the laser's wavefronts as well as the intensity and detuning of the laser. Curved wavefronts represent a laser beam that is either focusing or expanding depending on the concavity of the curvature. If the laser beams forming the MOT have curved wavefronts then the beams will be expanding or converging as they cross each other leading to larger or smaller intensities on opposite sides of the trap. In this situation instead of a symmetric force around the trap center, uneven forces exist that can misshape or destroy the trap. It is important that the MOT beams are well collimated and the wavefronts are flat as the laser traverses the trapping region.

The wavefront curvature is not dependent solely on the collimation of the laser beams. It is also affected by the optics that direct and shape the beam. When laser light interacts with a surface the optic's shape and surface qualities affect the wavefront of the light. Optics such as lenses are specifically designed to take advantage of this property to manipulate beam size and shape. However, irregularities in optics have the same effect with an unintended negative consequence.

Flat mirrors whose surface deviates from that of a perfect plane can introduce wavefront distortions, and surface defects such as scratches or bumps can be large enough to introduce unwanted wavefront distortions as well.

For this reason optics are often specified by flatness and a scratch/dig parameter. The flatness of an optic is usually given as a fraction of a wavelength. A mirror that is specified as "$\lambda/10$" deviates by no more than a tenth of a reference wavelength, which is typically specified as one of the sodium D lines (589, 589.6 nm) or the wavelength of a He-Ne laser (630 nm). The scratch/dig specification is given as two numbers. The first number specifies the largest scratch width in microns and the second number specifies the largest scratch depth in tens of microns. High quality optics are employed in directing and forming the beams of the MOT. Most of the optics purchased for the experiment have a surface flatness of $\lambda/10$ and a scratch dig of $10-5$.

Elliptical Mirror Preparation and Drilling

The elliptical mirror is specified to have a flatness of $\lambda/8$, and a scratch/dig of $60-40$. It was important to develop a process which would preserve these properties as well as produced a clean edge on the hole. A jagged edge could create a shadow in the trapping beam that does not correspond to a clear path through the mirror. This would allow atoms traveling out of the MOT to collide with, stick to, and build up on exposed substrate around the hole.

In order to drill through the front surface of the mirror without chipping it, the mirror surface was secured to a plate of Pyrex 7740. This is the same material that the mirror's substrate is composed of. To protect the mirror surface from rubbing against and being scratched by this contact we experimented with coating the mirror surface using different nail polishes.

Clear nail polish seemed to be a good choice because the composition of it does not contain any dyes or compounds known to attack the protected aluminum surface. Nail polish is also easily removed from the mirror with spectroscopically pure acetone. An even coat of clear nail polish

was applied to the mirror and allowed to dry. Afterward, the mirror was "glued" to the Pyrex plate by applying a new coat of polish and then pressing the two pieces together.

The mirror and plate were held in a specially made jig and the hole was then drilled using abrasive diamond bits. Due to the thickness of the mirror substrate and the small diameter of the bits, the hole was drilled in three steps. The majority of the hole was drilled with a 1.5 mm bit, then a short length was drilled with a 1 mm bit, and finally the last portion through the mirror surface and into the plate was drilled with a .75 mm bit.

After the mirror was drilled, the mirror and plate were placed in a shallow bath of spectroscopic grade acetone until the mirror freely came off the plate. The mirror was then thoroughly cleaned with fresh spectroscopic grade acetone and methanol.

Effects on the Mirror Surface

To test the reflectivity of the mirrors a white LED was filtered with a pinhole and re-collimated to fill most of the mirror's area. The reflected light was then sent to an Ocean optics spectrometer. Figure 3.5 shows the reflectivity plots of a new mirror and one that had been coated with and cleaned of nail polish. We found that coating and removing clear nail polish from the mirror did not noticeably decrease the reflectivity of the mirror.

To determine whether or not the drilling process affected the flatness of the mirror we set up an optical interferometer. A laser beam from a He-Ne laser was focused through a pinhole, expanded, and split into two beams. One of the beams from the beam-splitter was reflected off of the elliptical mirror being tested. The other beam was reflected off of a 2 in. diameter E02 dielectric mirror from Thorlabs. A matching beam splitter to the first was then used to recombine the beams. Finally, a pair of 2 in. achromatic lenses were used to image the plane of the mirror onto the CCD chip of a camera. The use of the achromatic lenses ensured that any distortions of the fringes were due to the difference in flatness between the elliptical mirror and the dielectric mirror rather than spherical

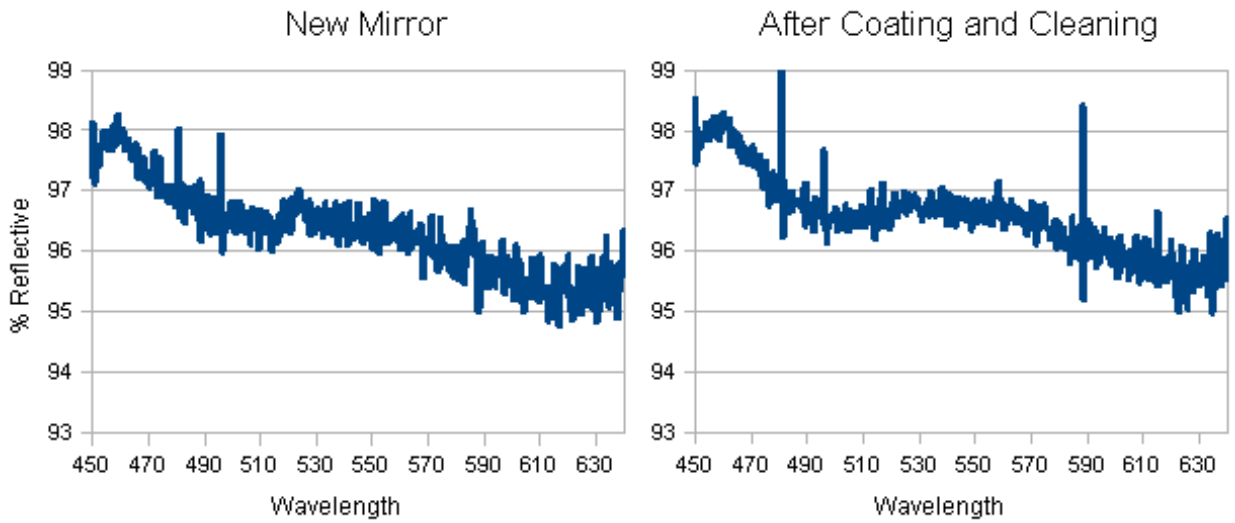

Figure 3.5 Reflectivity plots of the elliptical mirror. The plot on the left is the reflectivity of a new mirror. The plot on the right is of a mirror after it has been coated with nail polish and then cleaned.

aberration in the lenses.

The fringes produced in the interferometer with a new mirror are compared with the fringes produced with the drilled mirror in Figure 3.6. The image of the hole seen in the fringe pattern indicates that there was no noticeable distortion or chipping of the mirror surface by the drill bit. Computational analysis of the fringes from both mirrors showed that the mirror flatness was better than the specified $\lambda/8$ before and after the process.

3.4.2 LVIS Characteristics

Exit Velocity

As atoms leave the trap and travel toward the mirror, the same magnetic field used to Zeeman slow atoms approaching the trap serves to Zeeman accelerate the atoms leaving it. Unlike the slowing case, the field gradient does not need to be adjusted in order to accelerate the atoms. This is because, while the effect is the exact opposite of the slowing case, the detuning of the laser is

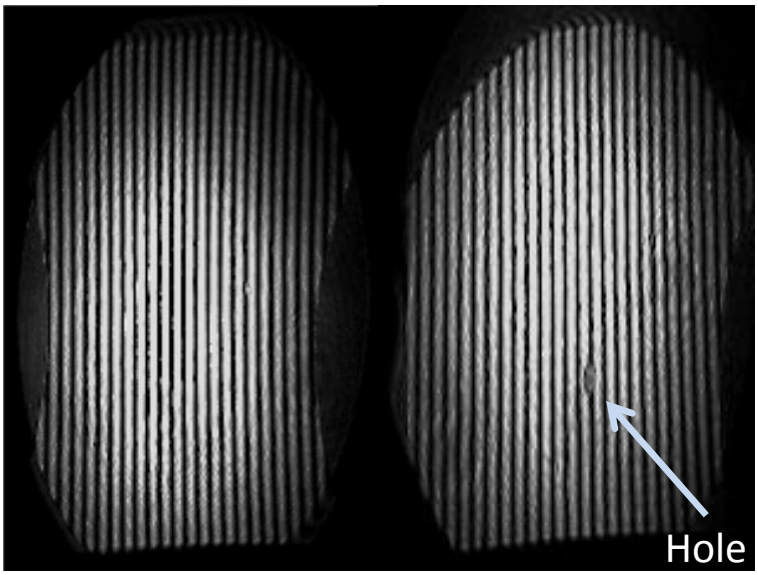

Figure 3.6 Interferograms of a new elliptical mirror compared to an identical mirror after drilling a .75 mm hole through the surface at 45°. The fringes were computationally analyzed to show that the $\lambda/8$ flatness of the mirror was preserved.

much smaller at the magnetic field zero.

Figure 3.7 shows the predicted exit velocity of the atoms as they leave the trap due to this effect. While the atoms would typically leave the trap traveling at the Doppler velocity of ~ 10.7 cm/s, they now get accelerated up to ~ 148 m/s when the cooling beam has a saturation parameter of 1.

The trajectory of the atoms in the LVIS is shown in Figure 3.8. These trajectories take into account the acceleration of the atom as it leaves the trap and are calculated for various values of S_0. Since there is ~ 14 cm between the the trap center and the mirror and another ~ 46 cm from the mirror to the front of the interferometer chamber, atoms leaving the trap at v_D could fall meters before getting to the ionization region, which is on the same horizontal plane as the trap. Atoms traveling 100 m/s will only fall < 0.1 mm on the path to the ionizing beam making the Zeeman accelerator a critical part of our experiment.

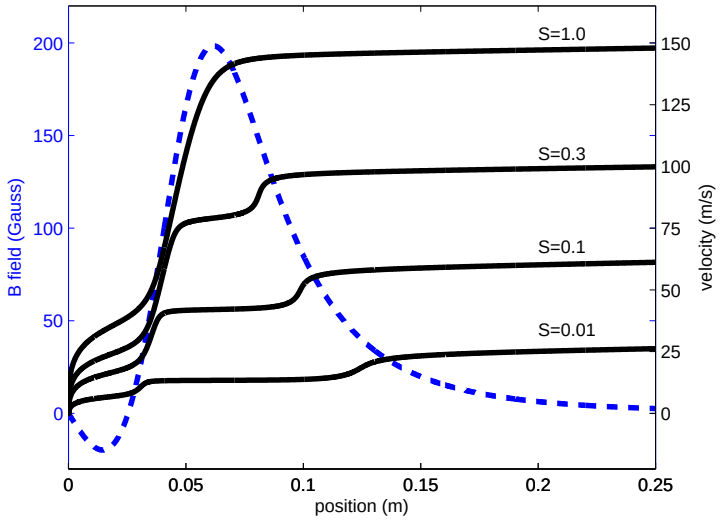

Figure 3.7 Exit velocity of the atoms as they leave the trap for different values of S_0. Even for very weak cooling beams the atoms are accelerated fast enough to span the relevant lengths of the apparatus.

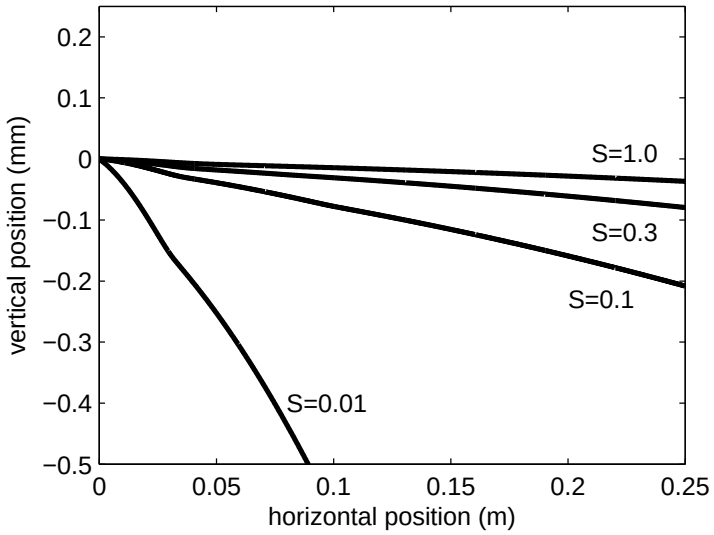

Figure 3.8 Vertical trajectory of atoms exiting the MOT for different values of S_0. These trajectories take into account the acceleration of the atom as they leave the trap.

Flux

The atom density of the exiting beam is principally determined by the capture rate of the MOT and the collisional loss rate of the trap. The capture rate for a MOT can be defined by [78, 79]

$$R_c = \frac{N_{MOT}}{\tau_{trap}}, \tag{3.2}$$

where N_{MOT} is the number of atoms in the MOT.

The lifetime of an atom in a MOT whose lifetime is principally limited by collisions with background gasses is given by $\tau_{trap} = 1/(\sigma n_0 v_{rms}) = 1/R_t$ where σ is the cross section of collisions with the background gasses, n_0 is the density of background gasses, v_{rms} is the root mean squared velocity of the atoms colliding with the MOT. The inverse of the trap lifetime is the rate of collisions with background gasses R_t. For a pressure of 10^{-7} Torr at room temperature and a typical kinetic collisional cross section of $\sigma = \sim 4 \times 10^{-19}$ m^2 $\tau_{trap} = 2.6$ s.

Background collisions are not the only loss mechanism for the case of strontium. Atoms can also be lost by spontaneous emission into the 5s4d ^{1}D$_2$ state as seen in Figure 1.2. The effective branching ratio for this splitting is $1 : 160,000$ [60, 61]. This amounts to a characteristic lifetime of $\sim .8$ ms in the trap before an atom is lost to the dark state.

If the LVIS were in the presence of perfect vacuum and the atom could could not decay to any intermediate energy states, then in steady state the flux of atoms in the beam would be equal to the capture rate of the trap. In reality each Sr atom can be knocked out of the trap with a collision, fall out of the trap into a dark state, or leave the trap by entering the LVIS beam. Taking into account all of the loss mechanisms Eq. 3.2 can be re-written as

$$R_c - F_B = N_{MOT} R_{lost}, \tag{3.3}$$

where F_B is the flux of atoms in the LVIS beam, and R_{lost} is the effective rate at which an atom leaves the trap due to collisions and the dark state. Additionally, the flux of atoms in the beam can

be written as

$$F_B = N_{MOT} R_{B,} \tag{3.4}$$

where R_B is the rate at which atoms enter the beam. Note that R_c and F_B are rates of atoms per second while R_B and R_{loss} are per atom rates. Equating these two expressions through N_{MOT} yields

$$\frac{R_c - F_B}{R_{lost}} = \frac{F_B}{R_B} \tag{3.5}$$

$$\Rightarrow F_B = \frac{R_c}{1 - \frac{R_{lost}}{R_B}}. \tag{3.6}$$

In the limit of $R_B \gg R_t$ the flux of the atom beam is equal to the capture rate and essentially all trapped atoms are transferred to the beam.

A rough estimate for R_B is given by the characteristic time it takes to cool an atom and load it into the LVIS beam. From Eq. 2.6

$$\mathbf{F} = -\alpha \mathbf{v} = M\ddot{x}$$

$$\Rightarrow \mathbf{v} = \mathbf{v}_0 e^{-\alpha t/M}$$

$$\tau_{cool} = \frac{M}{\alpha}, \tag{3.7}$$

which gives $\sim 7.4\ \mu$s for strontium. This is much shorter than the .8 ms it takes the atom to be lost to the 5s4d ^{1}D$_2$ state. Therefore, $F_B = R_c$ should be a decent approximation for our case. It should be noted, however, that atoms that are continuously recycled into the trap due to their transverse momentum in the LVIS beam could spend enough time in the trap to eventually be lost.

The capture rate R_c can be estimated from the flux and velocity distribution of the thermal atom beam used to load the trap. From [80] the total number of atoms per second exiting from a thermal source of our geometry (given in section 7.4.1) is given by

$$Q_{oven} = \frac{2r}{3l} n \bar{v} A_s, \tag{3.8}$$

where r is the radius of the source aperture, l is the distance between the collimating apertures of the oven, n is the density in the thermal source, $\bar{v}$ is the mean velocity of atoms in the thermal source, and A_s is the area of the source aperture. For our source the $Q_{oven} = 3.85 \times 10^{10}$ atoms/s including all the Sr isotopes. A rough estimate for the flux of atoms through the trapping region is given by comparing the cross section of the MOT to the area covered by the divergence of the thermal beam at the trap, giving $\sim 1.53 \times 10^8$ atoms/s.

Due to the limited range of the cooling force as described in section 2.2.1, only a fraction of this flux of atoms can actually be trapped. From the analysis of the Zeeman slower in section 3.3, when the slower is operating perfectly as many as $\sim 5.03 \times 10^6$ ^{87}Sr atoms/s can be loaded into the trap and by extension the LVIS beam. Without the slower, only ~ 79 ^{87}Sr atoms/s can be trapped and loaded into the beam.

3.5 Sr Ion Source

3.5.1 Generation of Ions

As seen in Figure 1.2, strontium has a bound state in the continuum that can be reached with one 460.8 nm photon (the cooling transition) and one 405.2 nm photon. The experimentally determined linewidth and peak cross-section of the auto-ionizing transition are ~ 1 nm FWHM and 5.6×10^{-15} cm^2 respectively [62, 81, 82].

We can deliver ~ 5 mW of 461 nm light and up to 100 mW of 405 nm light providing a minimum flux of 4.5×10^{17} pairs of ionizing photons per second. If the weaker cooling light beam is focused down to the diameter of the atom beam this would create a peak beam intensity of $I = 2.26$ W/cm^2. At this intensity an atom crossing the beam would spend approximately half of its time in the 5s5p $^1P_1$ excited state. For an atom traveling ~ 130 m/s this amounts to 2.9 μs during which the atom carves out a volume containing ~ 1.5 405 nm photons on average. The

low density of atoms in the beam constitutes a negligible loss of light intensity due to ionization, meaning nearly all the atoms will be ionized.

3.5.2 Velocity Tuning

Ionization of the atoms occurs between two high voltage electrodes in the vacuum. The electrical feedthrough allows for a potential of up to $\pm$20 kV to be placed across the electrodes. Therefore strontium beams of a few milli-eV up to 20 keV can be produced. Reversing the polarity of the field would also produce a beam of cold electrons over the same range. The design and function of the electrodes is given in section 7.5.1.

Chapter 4

Laser Systems

Several different laser frequencies are required for the experiment: 921.724 nm is needed for frequency doubling to 460.862 nm which is used for trapping, slowing and ionizing, 405.2 nm is used for ionizing, 407.886 nm is used in ion fluorescence probes and optical pumping, and two more 407.8 nm lasers whose relative frequency difference is locked to the hyperfine splitting of the $^{87}Sr^{+}$ ground state are used for driving the Raman transition the interferometer is based on. With the exception of the 405.2 nm laser, all of the laser systems contain a laser diode in an extended-cavity grating-stabilized configuration (ECDL) which is outlined in a basic tutorial on diode laser stabilization in Appendix A.

4.1 461 nm Laser System

The 461nm laser system is used for cooling neutral strontium into a collimated beam and to do the first step in the two-photon transition used to ionize the strontium. This wavelength was not available from commercial diode manufactures when the project began[1]. As such, the system is

[1]This wavelength is still not typically available, but we have since been able to purchase a wavelength selected diode at 460 nm. Even so, a diode in a grating stabilized ECDL would not put out as much power as our frequency

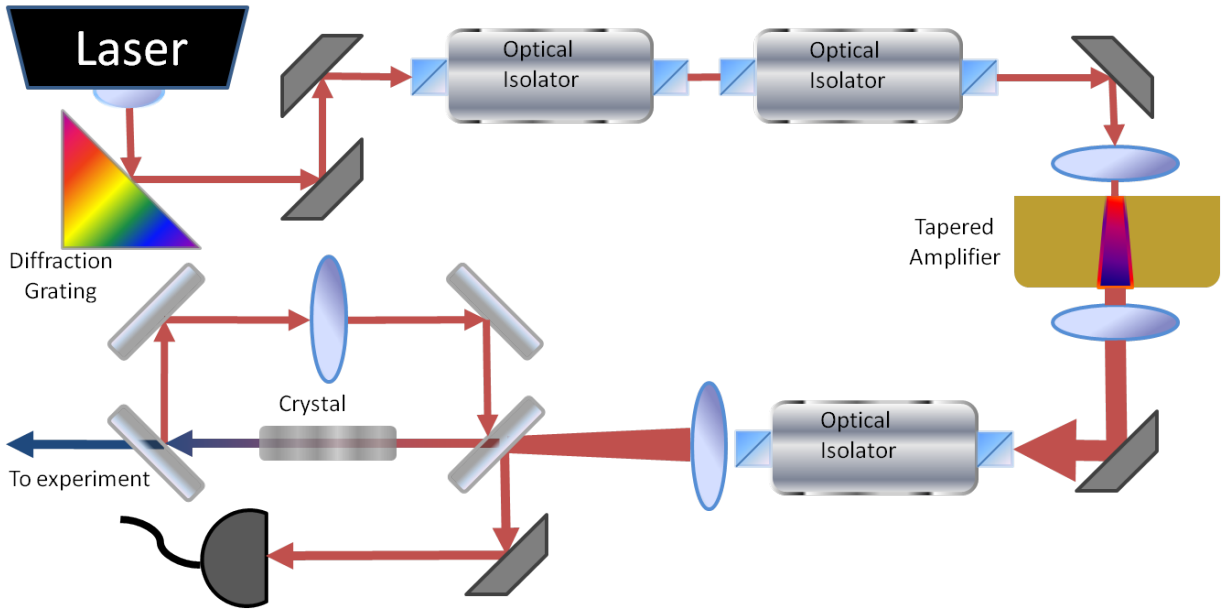

Figure 4.1 Functional diagram of the 461 nm laser system. A 922 nm master laser diode in an ECDL is coupled to a tapered amplifier (TA). The 1.5 W output of the TA is coupled to an optical buildup cavity housing a non-linear crystal. The crystal doubles the frequency of the light to produce 461 nm.

built around frequency doubling a longer wavelength with a non-linear crystal. The laser system is based off of a 922 nm diode in a grating stabilized extended cavity diode laser (ECDL). The ECDL puts out ~ 40 mW that is used to seed a 922 nm 1.5 W tapered amplifier (TA). The TA output is then coupled to an optical buildup cavity centered around a non-linear crystal. With the right conditions, the crystal produces laser light at twice the frequency of the input laser to provide the needed 461 nm. A diagram of the system is shown in Figure 4.1.

4.1.1 922 nm Master Diode Laser

As seen in the left hand part of Figure 4.2 a 922 nm grating stabilized extended cavity diode laser (ECDL) provides the basis for the entire 461 nm system. An anti-reflection coated (AR) laser diode from Toptica is stabilized with an 1800 lines/mm holographic diffraction grating. The laser

doubled system produces.

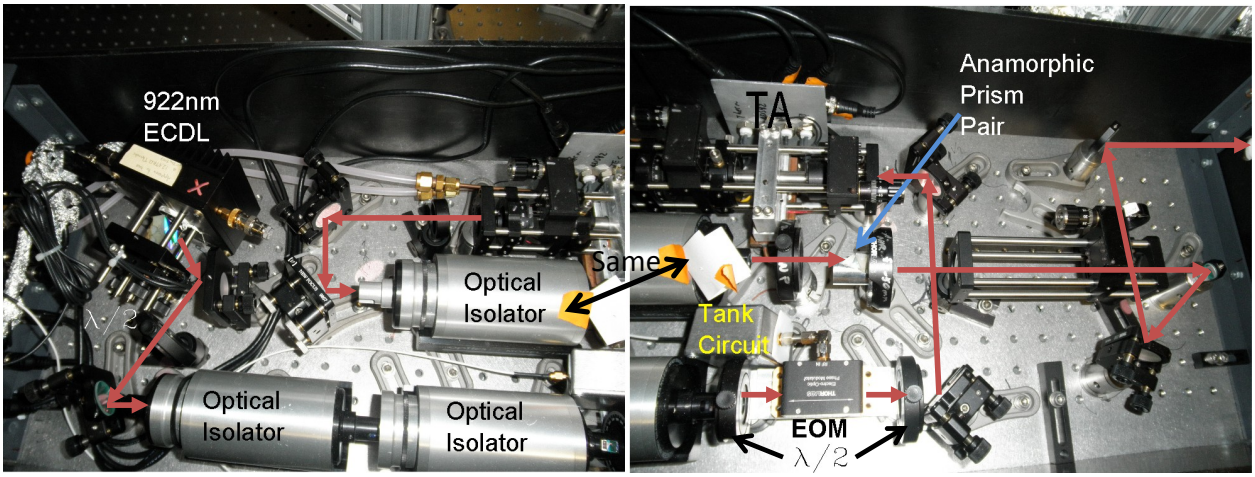

Figure 4.2 Picture of the 922 nm laser system. The system is split between the two pictures. The cage setup at the top of both pictures houses the tapered amplifier (TA). The 922 nm ECDL at far left side houses the master diode. After optical isolation the light passes through an electro-optic modulator at the bottom of the second picture. The light is coupled to the TA. The optical isolator in the middle of both pictures is the same piece and protects the TA from back reflection. An anamorphic prism pair and a series of lenses then reshapes the light for coupling to an optical frequency doubling cavity.

is isolated from downstream components by two optical isolators from Conoptics, which provide ~ 50 dB of isolation while allowing $\sim 85\%$ transmission. The light is then passed through a Thorlabs electro-optic modulator (EOM), which phase modulates the light to help us to lock the laser frequency to an optical doubling cavity. After the EOM, the laser is coupled to a Toptica 1.5 W TA. The output of the TA is shaped with an anamorphic prism pair and focused to match the mode of the optical doubling cavity.

The diode is rated for ~ 80 mW output power and its AR coating allows tuning with a diffraction grating over an extremely large range (915 nm to 980 nm). Wavelengths around 935 nm are favored by the gain medium with the gain falling off rather quickly towards the end of the frequency range. When grating stabilized at 922 nm output powers near $\sim 45 - 50$ mW are achieved at the maximum intra-cavity power (between the grating and the diode).

One of the characteristics of the ECDL is that it is extremely sensitive to the alignment of

the grating's feedback. When the optical feedback is not well enough aligned for the diode, the power in the principal mode can oscillate as other modes compete against it in the diode's gain medium while still coupling well to the TA. This results in a large amplitude modulation on the frequency doubled light and overall instability in the system. To help passively stabilize the ECDL, the diffraction grating's kinematic mount is coupled to the laser head with a 30 mm optics cage. A specially machined bronze slug holds the grating at the correct feedback angle to accommodate this configuration.

The entire laser system is passively stabilized to reduce technical noise from vibrations. Despite the fact that the optics table "floats" and that the floor of the lab is mechanically separated from the rest of the underground lab structure, recent on-campus construction has managed to produce large seismic vibrations in the lab. To over come these vibrations the entire 922 nm laser system is built on a breadboard with a high density, internal damping core. The breadboard sits on Sorbothane pads to further isolate the system from vibrations on the table.

4.1.2 922 nm Tapered Amplifier

The stabilized light from the 922 nm laser diode is used in a master-oscillator-power-amplifier (MOPA) setup with a Toptica 1.5 W tapered amplifier. The amplifier has an input aperture of 1 μm$\times$1 μm and an output of 1 μm$\times$400 μm. Light from the 922 nm master diode is coupled into the tapered amplifier through a microscope objective mounted in an x-y translation mount.

Since the dimensions of the input facet of the TA are small, the output power of the amplifier is very sensitive to overall alignment, drift, and vibrations. To stabilize the setup, all of the pieces are mounted to a 30 mm cage as seen in Figure 4.3. The TA as sold is built onto a standard C-mount that comes bolted and soldered to a proprietary mount made by Toptica, which protects the TA from damage when making electrical connections. In order to adapt the OEM mount for the cage a two-piece block was machined as shown in Figure 4.3.

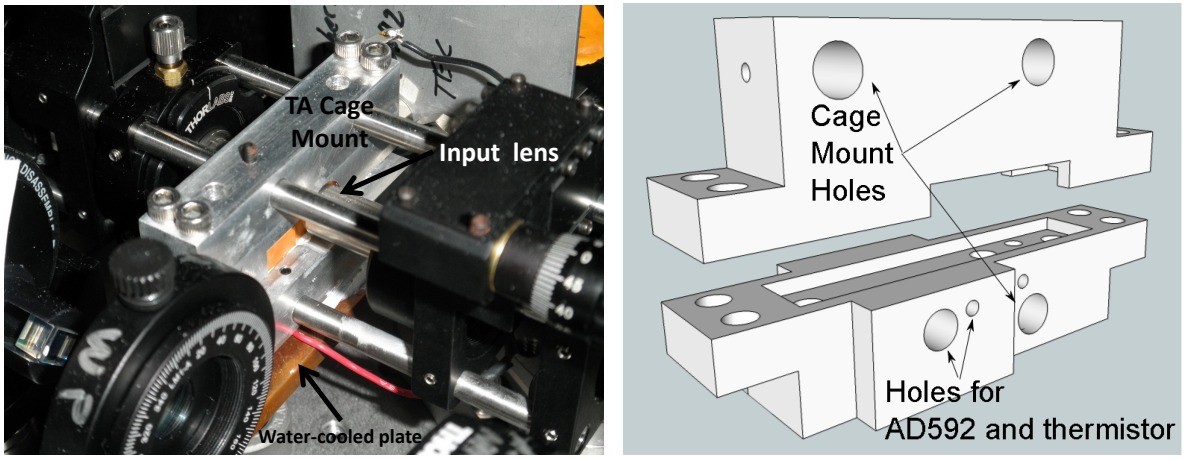

Figure 4.3 A picture of the tapered amplifier setup is shown on the left and a diagram of the mount for the TA is shown on the right. In the TA cage setup two microscope objectives, a cylindrical lens (not seen), and the TA are mounted on a 30 mm cage. The TA's OEM mount is sandwiched in a custom piece that adapts it for the cage. A water cooled copper plate is mounted to the bottom for thermal management.

The maximum current for the TA is 3 A corresponding to a potential drop of several volts. Since only 1.5 W of that power is potentially converted into light energy, thermal management of the TA is very important. To control the temperature of the TA a thermo-electric cooler (TEC) is mounted between the cage adapter piece and a water-cooled copper plate. The copper plate bolts to the adapter, holding the TEC in a seat machined for it on the adapter's bottom. A thermistor and an AD592 are embedded in the adapter for temperature control and monitoring.

The thermal assembly hangs from the aluminum block so that changing dimensions due to thermal drift and mechanical instability of the plate do not affect optical alignment. To create the water-cooled copper plate channels were cut in a piece of copper. A length of $1/8$ in. copper tubing was then threaded through the channels and soldered to the plate. Since the tube is one piece all the way through, its ends can be carefully bent and positioned to accommodate other components in the laser setup without the fear of producing a leak.

The vertical output of the TA is collimated by a matching objective to the input lens, and the remaining horizontal divergence is collimated by a cylindrical lens. This collimated beam

is then passed through a Conoptics isolator, which provides $\sim$ 35 dB of isolation and allows $\sim$ 83% transmission. Although the beam is collimated at this point, it is still cigar shaped and is unsuitable for coupling to an optical cavity. In order to correct the beam shape a mount was machined to custom space an anamorphic prism pair creating a nearly Gaussian beam. The beam is then telescoped and focused to match the TEM_{00} mode of the doubling cavity so that $\sim$ 67% of the light couples into the cavity.

4.1.3 PPKTP Doubling Cavity

Despite the output power of the TA, it still insufficient to produce the desired amount of 461 nm light in our non-linear crystal. The necessary intensity is created by placing the crystal in a low-loss optical ring cavity where constantly circulating light builds in power as seen in Figure 4.1.

Figure 4.4 shows the layout of the doubling cavity, which consists of four flat mirrors and a biconvex lens. This is opposed to the typical configuration of two curved mirrors and two flat mirrors to create the desired beam waists in the cavity. Since the non-linear crystal in our cavity creates no astigmatism of the beam, the presence of curved mirrors can only add astigmatism as opposed to correcting for it.

Typically, the non-linear crystal is in the smaller focus of a bow-tie cavity, which means one of the curved mirrors must be a custom dichroic mirror to let the doubled light out. Since all of the mirrors are flat in our setup, no custom parts need to be changed in order to change the beam waist. Only the lens might need to be changed as the cavity length is manipulated in order to keep the cavity stable. The use of four flat mirrors also makes short cavity lengths, and rectangular cavity configurations, possible because reflecting off of the mirrors at large angles does not produce astigmatism as it would with curved mirrors.

A picture of the actual doubling cavity is shown in Figure 4.5. The polarization of the pump beam is set by a $\lambda/2$ waveplate before it enters the cavity. The rejected pump light is sent to

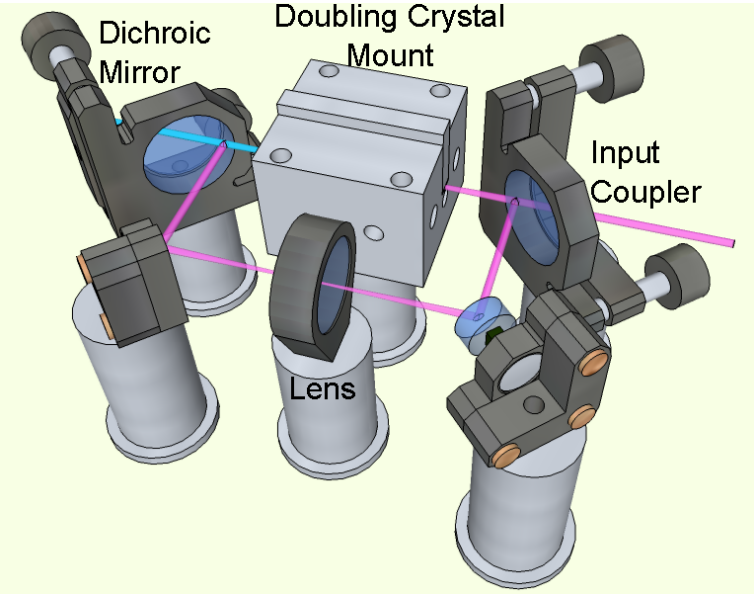

Figure 4.4 The basic layout of the doubling cavity for the 922 nm laser is shown. All the mirrors are flat and the cavity stability is dependent on the lens. The length of the cavity, and hence the size of the beam waist in the non-linear crystal, can be changed while keeping the same mirrors. Also, the cavity length can be made very short by reflecting off of the mirrors at large angles without introducing astigmatism as would happen with curved mirrors. One of the mirrors is mounted on a piezo-electric crystal for scanning the frequency of the doubled light.

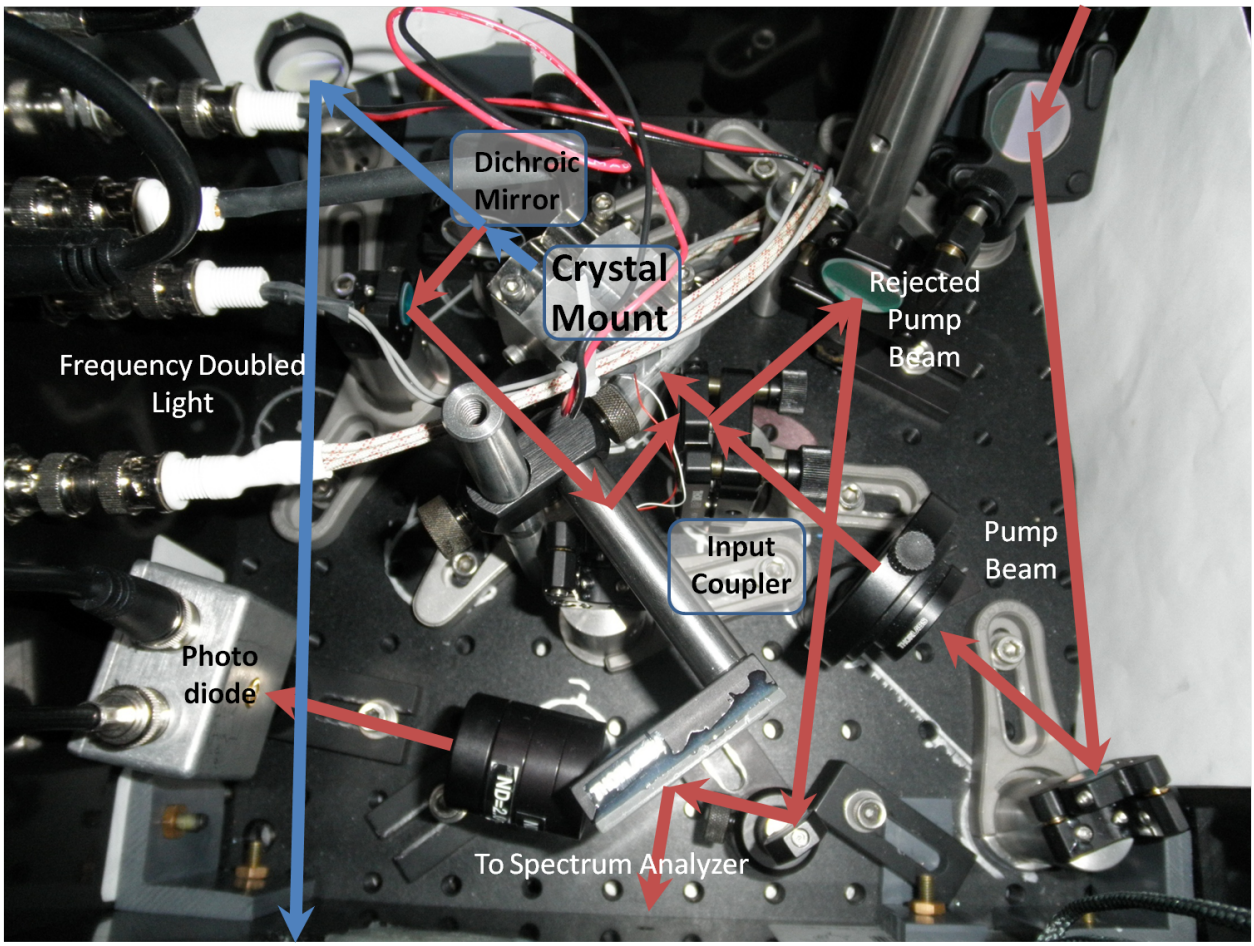

Figure 4.5 Picture of the doubling cavity. The pump beam enters the cavity from the right and frequency doubled light exits the cavity on the left before being directed to the bottom of the picture and out of the setup. The rejected pump light from the front of the cavity is sent to a photo-diode for locking the laser to the cavity and can be sent to a spectrum analyzer or wave-meter for diagnostics.

a photodiode and used to lock the 922 nm master laser system to the cavity. A small part of the rejected pump is also split off and can be sent to a spectrum analyzer or wave-meter for performing diagnostics on the laser system.

It is important to prevent light from circulating in both directions in the cavity as this can create a standing wave that decreases the power of the output light. While our setup increases the number of surfaces the light can reflect off of, these back reflections are minimal. The use of an AR coated bi-convex lens decreases the possibility of strong back reflections because the lens has no flat surfaces. Also, the doubling crystal's AR coating has a transmission of $\geqq 99.75\%$ for both wavelengths minimizing back reflections from its surfaces.

The cavity's quality factor (Q) likewise suffers little from our setup. Losses from the surface of the lens and crystal are negligible compared to the crystal's large single pass conversion efficiency. Since large losses in the crystal already result in a low Q, small losses from the optics are not as critical as in a high Q cavity.

The versatility of our cavity configuration was put to the test in optimizing the cavity beam waist for stability as well as second harmonic generation (SHG). Ultimately, the cavity was set up as a square. The lens has a focal length of 125 mm and the round trip cavity length is ~ 25 cm, producing a beam waist of ~ 192 μm in the crystal. A picture of the doubling cavity in operation is shown in Figure 4.6.

Doubling Crystal and Mount

The 922 nm light is converted to the 461 nm needed by second harmonic generation (SHG) in a periodically-poled potassium-titanyl-phosphate (PPKTP) crystal. PPKTP was chosen for several reasons. First, there are several advantages to using a periodically poled crystal. The crystal's structure is reversed at set intervals so that the phase difference between the fundamental and doubled light is canceled out across the length of the crystal. This technique is known as quasi

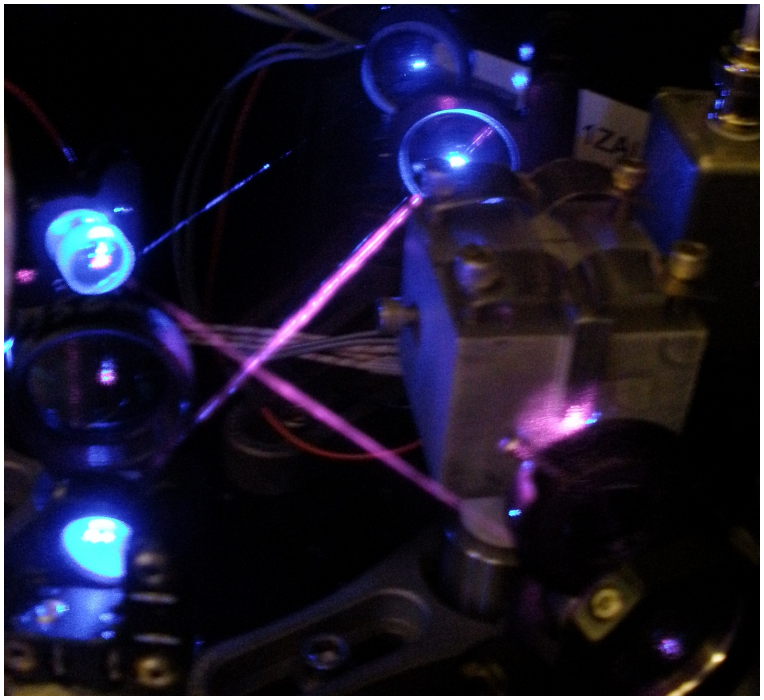

Figure 4.6 Picture of the doubling cavity in operation. Note that the cavity is shown in a bow-tie configuration here while the ultimate setup was a rectangular cavity. The input coupler is seen in the bottom right corner, and the dichroic mirror is second mirror from the top edge. The top mirror on the left side is mounted on a piezo crystal for manipulating the length of the cavity. The pink light is the camera's interpretation of the 922 nm pump light, which the camera is not very sensitive to, and the blue light is the frequency doubled output.

phase matching because the mismatch in the phase between the pump and frequency doubled light is periodically canceled.

In critical and non-critical phase matching, the second harmonic generation (SHG) light is orthogonally polarized to the pump light that it is phase matched to. In quasi phase matching, the pump light and the SHG light can be phase matched with the same polarization. The effective nonlinearities for parallel polarization modes tend to be significantly larger, resulting in greater conversion efficiency.

Quasi phase matching also allows the fundamental to travel down an optical axis of the crystal eliminating walk-off[2]. The absence of walk-off and the use of a nearly collimated beam in the crystal greatly simplified the calculations necessary to design the laser system and it can result in greater conversion efficiency.

Of the crystals suitable for frequency doubling 922 nm PPKTP boasted one of the highest accessible non-linear optical coefficients at $\sim$6 pm/V [83]. KTP can also be annealed if it sustains photo-refractive damage that is not too severe. KTP was also chosen in part because of the experience of other groups [84, 85] who had used KTP in similar applications and had published their results.

The only downside to using KTP is its relatively low damage threshold compared to other crystals. To avoid damaging the crystal we purchased the longest PPKTP crystal available (30 mm). By placing the crystal in a modest build-up cavity (finesse of $\sim$ 100) with a relatively large beam waist ($\sim$ 192 μm) enough 461 nm light is produced for our application. The intensity of the light in the crystal is much smaller in our setup than what other groups used when photo-refractive damage occurred in similar crystals.

The crystal sits in a mount designed to keep a uniform temperature across the entire length of the crystal. The mount can be seen in Figures 4.5 and 4.6 and a schematic of it is shown in

[2]There is no walk-off in non-critical phase matching as well.

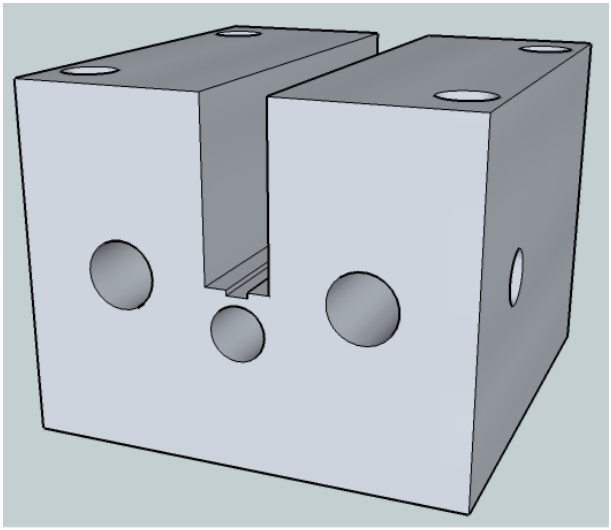

Figure 4.7 Mechanical drawing of the temperature controlled mount for the PPKTP doubling crystal. Cartridge heaters are used maintain the temperature of the crystal at $\sim 50^\circ$ C.

Figure 4.7. The mount consists of a deep slot machined lengthwise in an aluminum block. Along the bottom of the slot a groove is machined slightly wider and not quite as deep as the thickness of the doubling crystal. The crystal sits in the groove to prevent it from twisting or shifting in position. Two holes are drilled into the block on either side of this groove so that cartridge heaters can be inserted to control the temperature of the crystal. There are also holes for a thermistor and an AD590 for temperature control and monitoring. Soft springs made from shim stock are used to clamp down the crystal with a second piece of aluminum that fits in the slot.

4.1.4 Doubling Cavity Lock

The 461 nm laser system is locked to the doubling cavity using the Pound-Drever-Hall technique [86]. As phase modulated light is scanned across the resonance of an optical cavity, light leaking out of cavity will mix with reflected light from the incident laser. This signal can be detected with a photodiode and mixed down with the frequency of the phase modulation to create a DC error signal.

The EOM seen between the master diode laser and TA in Figure 4.2 modulates the phase of the 921 nm laser at $\sim$ 10 MHz, effectively placing sidebands on the light spaced on either side of the main carrier frequency by the modulation frequency. Light reflected off of the input mirror of the cavity interferes and beats with resonant light leaking out of the cavity. The amplitude and phase of the beating is detected with a reverse biased photodiode. This signal is then mixed with the drive frequency of the EOM to create an error signal.

The error signal is passed to a proportional-integral-differential (PID) lock circuit which feeds back to the 921 nm master diode laser. A description of the function and operation of the PID lock circuit is given in section 5.3.

Once the 921 nm master laser system is locked to the doubling cavity the frequency of the 461 nm output beam can be scanned by changing the length of the cavity. One of the cavity mirrors is mounted to a piezo-electric crystal whose size is dependent on the voltage across the crystal. As the voltage on the crystal is changed the length of the cavity changes accordingly and the frequency of the cavity's resonant mode is changed. The lock adjusts the frequency of the 921 nm master laser to stay resonant with the cavity when this happens, effectively scanning the frequency of the 461 nm output light.

4.1.5 461 nm Laser Specifications

The laser system consistently produces $\sim$ 80 mW of 461 nm light. The mode of the beam is the TEM_{00} mode of the doubling cavity, which is Gaussian and almost perfectly circular. When the system is peaked up and well aligned the 921 nm laser has been known to stay locked to the doubling cavity for the better part of an hour at a time with no significant amplitude modulation on the 461 nm light.

4.2 405 nm Diode Sources

The remaining laser systems are all based off of 405 nm diodes. We have found that laser diodes at this wavelength can be more difficult to work with than most red and infra-red laser diodes and can only be frequency pulled from about -1 nm to $+2$ nm from their free-running wavelength at room temperature. However, due to manufacturing errors the free-running wavelength of any given 405 nm diode may actually range from 400 nm to 410 nm. There is no guarantee that simply purchasing one will result in a useful diode for a specific application.

We initially purchased the only two wavelength-selected 405 nm diodes we could find that fit our needs. Overall, five diodes were needed in this wavelength range. To acquire the remaining diodes, we found it economical to purchase inexpensive bulk lots of 405 nm diodes extracted from consumer devices and wavelength-select from them in the lab. From this group of random laser diodes 5 were acquired in the right wavelength range making it a worthwhile endeavor.

4.3 405 nm Ionization Laser

As mentioned in section 1.2.3 and shown in Figure 1.2 strontium can be ionized in a two photon process driving the atom to an auto-ionizing electronic state in the continuum. The first part of this transition can be made by using light from the 461 nm laser system to drive the resonance transition. The second transition from the $5s5p\ ^1P_0$ excited state to the $5p^2\ ^1D_2$ auto-ionizing state is centered at 405.2 nm and is ~ 1 nm wide FWHM [87].

Since the width of this transition is much broader than the linewidth of a free running laser diode, this 405 nm laser is simply temperature and current tuned to the right wavelength with no diffraction grating or frequency lock. The laser system simply consists of a temperature stabilized diode protected by a pair of optical isolators.

4.4 407.8 nm $^{87}Sr^{+}$ Fluorescence Laser

This laser is required to drive strontium ions on the 5s $^{2}S_{1/2}$ to 5p $^{2}P_{3/2}$ resonance transition. The diode is frequency stabilized in an ECDL and is protected from back reflection with an optical isolator. It will be frequency locked to the same vapor cell as the cooling laser for neutral strontium, as described in section 4.6.

4.5 Design of Raman Laser System

While this laser system is now the subject of a Master's thesis for another student in the lab, I laid out the initial design and began the mock-up of the assembly before turning it over. It is similar to the laser system found in [88] and I present its basic theory of operation here. This system provides two output beams designed to drive strontium ions between the levels of the hyperfine ground state via a virtual state. In order to do this the relative frequency of the two beams must be stable and held equal to the hyperfine splitting.

As seen in Figure 4.8 the system is based off of a master 405 nm laser diode in a standard ECDL configuration protected by two optical isolators. The beam is then passed through a 2.5 GHz AOM, retro-reflected, and passed through again. The two diffraction orders created by passing through the AOM in opposite directions are separated from each other by twice the AOM frequency.

The required frequency between the two beams can then be tuned by adjusting the drive frequency of the AOM to match the hyperfine splitting of the ground state. An HP8648D frequency generator has been purchased to drive the AOM. The signal generator is stable to .01 Hz so that the frequency difference of the two beams will have a fractional frequency stability of $\sim 10^{-16}$. This results in a drift of the phase of the interferometer that is negligible over the time it takes an ion to pass through the interferometer.

The active aperture and damage threshold of the AOM will limit the two diffraction orders to

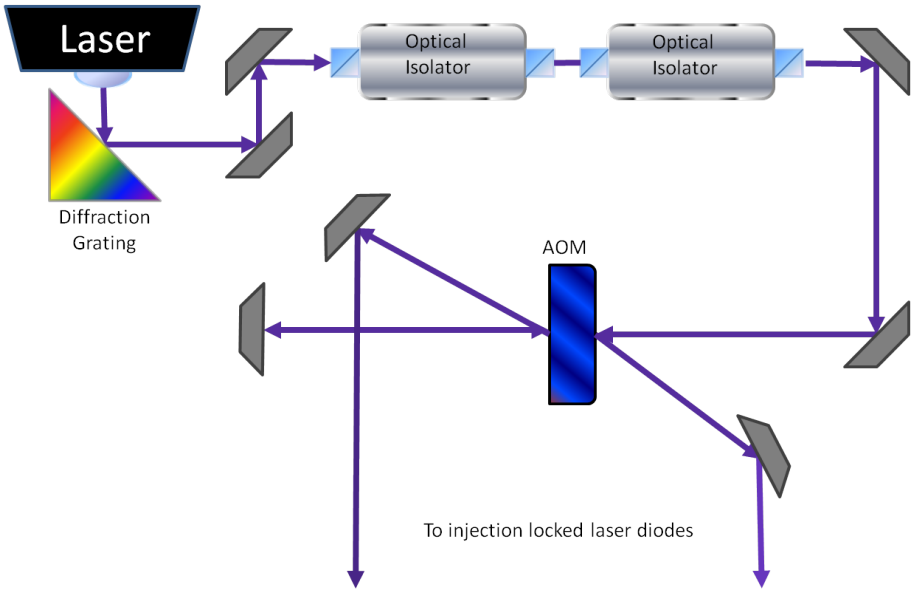

Figure 4.8 Overview of the Raman laser system. Two lasers beams differing in frequency by the hyperfine splitting of the $^{87}Sr^{+}$ ground state are created by sending a master beam through an AOM in both directions. The beams are then use to injection lock two more laser diodes not shown.

intensities too small to be of use in the interferometer. Therefore, these beams will be used to injection lock two more wavelength selected 405 nm diodes to provide a potential 100 mW in each beam.

The real benefit to this design is that both beams differ only in frequency. Because they are created by frequency shifting a single laser beam they are exact copies of each other in linewidth, and lineshape.

4.6 Frequency Locking

There exists several good references for understanding the theory and methods used to implement laser frequency locks [86, 89]. In atomic physics two different locks are commonly used. In one, the frequency of a laser is locked to the mode of an optical cavity as in the 461 nm laser system of

section 4.1.4. In this section we describe locking the frequency of a laser to an atomic transition using a classic Doppler-free saturated absorption technique.

Typically, lasers are locked to a particular atomic transition with the use of a vapor cell such as the one described in Chapter 6. In principle, an on resonance laser passing through the vapor is scattered by the atoms, which decreases the intensity of the beam. An absorption profile centered on the transition frequency can be generated by measuring the intensity as the laser's frequency is swept across the transition. However, atoms that are Doppler shifted into resonance with the beam absorb light far away from the atomic resonance. Whereas the natural linewidth of the transition is typically a Lorentzian shape of tens of MHz FWHM, the broad velocity distribution of the atoms usually creates an absorption profile that is a largely Gaussian lineshape of several GHz FWHM.

For most AMO applications the Doppler broadened absorption profile is too wide to create a precise enough lock. A narrow feature suitable for locking can be added to the absorption profile by splitting the laser into two beams that are aligned counter-propagating through the vapor cell. This technique is known as saturated absorption spectroscopy. As the laser frequency is scanned the two beams first come into resonance with opposite wings of the atom's velocity distribution. On resonance, both beams drive the transition in the same velocity class of atoms causing a decrease in the number density of ground state atoms.

In theory, when the laser is on resonance the more intense beam (the pump) drives a large population of atoms into the excited state. Depending on the intensity of the pump, as few as half of the atoms at the correct velocity are then available to absorb the lower intensity beam (the probe). This produces a feature in the absorption signal of the probe light known as a Lamb dip, whose linewidth is usually on the order of the natural linewidth. Figure 4.9 shows an absorption profile with the Lamb dip obtained in our setup.

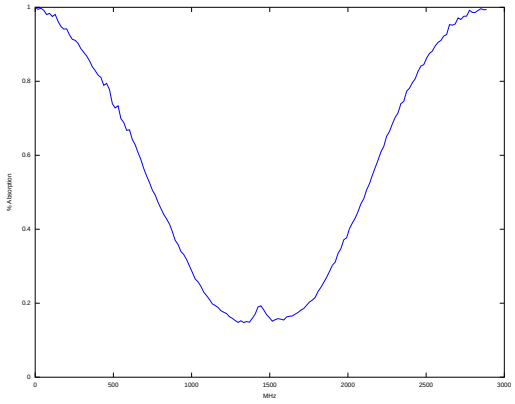

Figure 4.9 Absorption profile of the probe beam in the saturated absorption setup for strontium. Note that the entire Doppler broadened profile is not as wide as might be expected. This is due to the operation of the vapor cell as described in section 6.2. The Lamb dip at FWHM is $\sim 3\Gamma/2\pi$. While the Lamb dip as shown here is centered in the absorption profile, an AOM was subsequently used to shift it below the transition resonance for cooling and trapping Sr^{87}.

4.6.1 461 nm Lock

The lock setup for the 461 nm laser is shown in Figure 4.10. The laser is split into two beams by a $\lambda/2$-waveplate and a polarizing beam splitter, and the beams are overlapped counter-propagating through a strontium vapor. One of the beams passes through an AOM that shifts the position of the Lamb dip. This in turn shifts the lock frequency so that the laser is appropriately detuned for cooling and trapping.

The 461nm laser is frequency locked to ^{88}Sr absorption signal with shift. The shift is calculated so that the laser's frequency is red detuned from the ^{87}Sr cooling transition in order to trap and cool the atoms when locked. There are two main reasons to structure the lock this way. First, the natural abundance of ^{88}Sr is $\sim 82.5\%$ compared to 7% for ^{87}Sr. This means that ^{88}Sr provides a much larger absorption signal as seen in Figure 4.11. Also, the resonance frequencies of the isotopes are close enough that the ^{87}Sr absorption signal is in the wings of the ^{88}Sr signal, washing it out as seen in Figures 4.11 and 4.12.

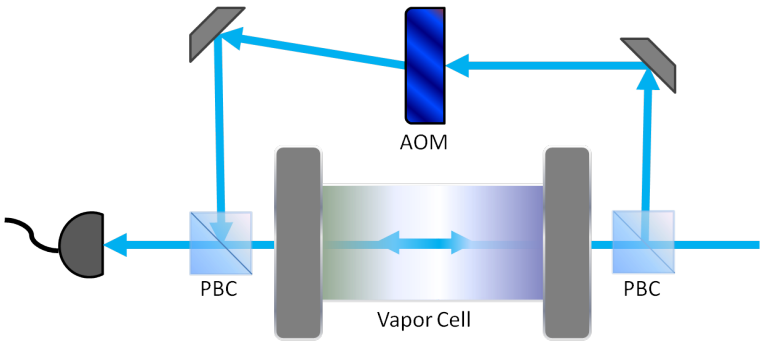

Figure 4.10 Layout of the saturated absorption lock for the 461 nm laser system. The laser is split into two beams and sent through a strontium vapor counter-propagating. When the laser frequency is such that both beams are resonant with the same velocity class of atoms a Lamb Dip is generated in the absorption scan. The lock-in electronics lock the laser frequency to the peak of the Lamb Dip. The use of an AOM in the pump beam shifts the location of the Lamb Dip to detune the laser for cooling and trapping.

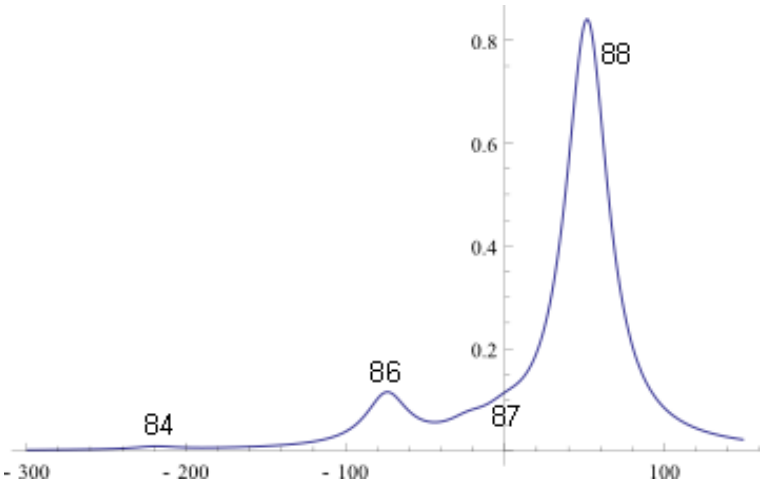

Figure 4.11 Calculated spectrum of the naturally occurring stable Sr isotopes. ^{88}Sr has the greatest abundance (82.58%), followed by ^{86}Sr (9.86%), then ^{87}Sr (7%), and finally ^{84}Sr (.56%).

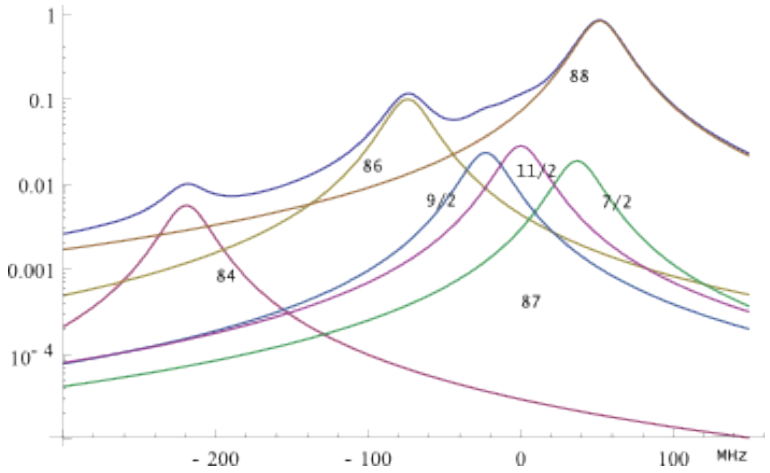

Figure 4.12 A logarithmic plot of the calculated absorption spectra of different Sr isotopes. The uppermost blue line is the overall absorption of a sample of strontium with natural isotopic abundances. The remaining lines correspond to the spectra of the individual isotopes. Brown is the ^{88}Sr line, light brown is ^{86}Sr, and the purple line on the left is ^{84}Sr. The hyperfine transition spectra for ^{87}Sr is given by the overlapping blue, purple, and green lines on the right.

Aside from that, ^{88}Sr provides a narrower resonance for locking. Strontium-87 has a nuclear spin $I = 9/2$ and a total spin of $F = I + J$. This means the cooling transition is split into three resonances from the $^1S_0$ $F = 9/2$ to the $^1P_1$ $F = 7/2,\ 9/2,\ 11/2$ states as seen in Figure 4.12. The overlapping hyperfine transitions of ^{87}Sr will produce a broader signal that is not as useful for locking as the single ^{88}Sr transition.

In order to lock the laser to a frequency which is $\Gamma/2$ below the ^{87}Sr F=$9/2$ resonance, a shift of 90.8 MHz from the ^{88}Sr peak absorption is required. This frequency difference is well within the reach of standard AOMs. As see in Figure 4.10 an AOM is placed in the path of the pump beam. This shifts the frequency relative to the probe beam, and effectively changes the position of Lamb dip in the absorption spectrum.

To understand how the AOM changes the frequency setpoint of the lock, consider the velocity of the atoms that the pump and the probe are resonant with as the laser's frequency is swept. As the the laser frequency ν_L is scanned the probe is resonant with atoms whose Doppler shift is equal

to the detuning from the resonance frequency ν_0. Mathematically we can write

$$\nu_L - \nu_0 = \nu_0 \frac{v_a}{c}, \tag{4.1}$$

where v_a is the velocity of the atoms, and c is the speed of light. The pump is similar except that its frequency is modified by the AOM frequency ν_{AOM} so that

$$\nu_L + \nu_{AOM} - \nu_0 = -\nu_0 \frac{v_b}{c}, \tag{4.2}$$

where the minus sign on the right-hand side comes from the fact that the laser is traveling in the opposite direction. When both beams are on resonance with the same velocity class of atoms $v_a = v_b$ so that

$$\nu_L + \nu_{AOM} - \nu_0 = -\nu_L + \nu_0 \tag{4.3}$$

$$\Rightarrow \nu_L = \nu_0 - \frac{1}{2}\nu_{AOM}. \tag{4.4}$$

The laser will lock at a frequency shifted from the atomic resonance by half of the AOM frequency shift. The frequency of the lock point is shifted down when the AOM shifts the pump frequency higher.

In order to lock the laser red detuned by $\Gamma/2$ for cooling ^{87}Sr atoms on the F$= 9/2$ cooling transition, an AOM with a bandwidth of 100 MHz FWHM centered at 200 MHz from Brimrose is used to frequency shift and modulate the pump beam. The AOM has $\sim 70\%$ diffraction efficiency and is set to blue-shift the pump beam frequency by $2 \cdot 90.8 = 181.6$ MHz.

The features of the Lamb Dip are ideal for locking to since the width of the peak is on the order of the transition linewidth as long as the laser intensities are not significantly larger than the saturation intensity [90]. In our setup a smaller and broader than expected Lamb Dip was initially observed. We think this is due to the fact that for most of the experiment the strontium sample used in our vapor cell was impregnated with mineral oil as discussed in Appendix D. Despite this undesirable effect, the FWHM of the Lamb Dip was still ~ 3 times the natural linewidth of the

transition but narrow and deep enough to produce a usable error signal. The characteristics of the Lamb dip have not been revisited since refurbishing the cell with a better strontium sample. An error signal is generated from the Lamb dip using a homodyne lock-in method as described in section 5.6.

4.6.2 407.8 nm Lock

Normally, for locking a laser to an ion's atomic transition the use of a simple vapor cell is insufficient and a discharge tube is used. The collisions in a standard vapor cell are typically not energetic enough to create an optically dense vapor of ions. In the discharge tube a high current arc is used to ablate neutral atoms off of a cathode. Collisions with electrons and other atoms in the high energy environment then creates a sufficient density of ions that absorption spectroscopy can be performed in the much the same way as with a standard vapor cell. However, Sr discharge tubes are expensive and the lifetime of the tubes are short as the strontium tends to absorb the buffer gasses used in the cell.

In order to use the same vapor cell as the 461 nm laser lock for this laser, a high-density column of ions will be created in the cell by photo-ionization. This setup is shown along with the 461 nm lock in Figure 4.13. The ions will be produced by taking the 461 nm pump beam and overlaying it with a beam from the 405.2 nm laser. The 461 nm pump beam from the 461 nm lock is rejected by a polarizing beam cube, and can be re-used in the 407.8 nm lock setup. A beam sampler will be used to overlay the pump beam and the 405.2 nm laser to create a dense column of strontium ions. The only requirements are that the beam sizes and intensities be large enough so that all passing atoms are likely to be ionized. A saturated absorption setup with the 407.8 nm laser can then be aligned along the column of ions.

A simple homodyne lock-in scheme will be used to implement the lock. The 407.8 nm probe beam will pass through an electro-optic modulator (EOM). The EOM phase modulates the light,

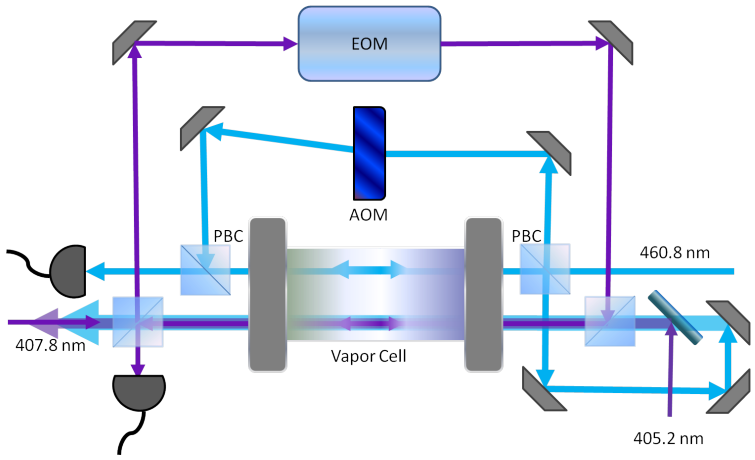

Figure 4.13 Diagram of the setup for locking to the resonance transition in $^{87}Sr^{+}$ using a standard vapor cell. Light is recycled from the 461 nm laser lock and overlaid with the 405.2 nm laser to produce a column of ions in the neutral atom vapor. The fluorescence laser can then be aligned along this column in a saturated absorption setup. Since no frequency offset from resonance transition is desired, this laser is modulated by an EOM to produce the error signal for locking.

effectively creating frequency sidebands about the main carrier frequency. As the laser frequency drifts, the modulation created on the probe beam signal by the frequency side bands can be mixed down with the EOM drive frequency to produce an error signal.

Chapter 5

Laser Control and Electronics

All the lasers involved with the apparatus are diode based lasers that require a certain amount of passive stabilization as well as relatively quiet electronics in order to operate stably in a single mode. The list of electronics devices developed for the laser systems includes a current driver, a temperature controller, a proportional-integral-derivative lock/scan controller, a homodyne detector, a piezo amplifier, a master processing unit to control the previous devices, and a homodyne lock-in system.

5.1 Diode Laser Current Controller

Our current driver builds upon the Hall-Libbrecht design [91], a standard precision current controller design which has been used extensively in atomic physics labs. The Hall-Libbrecht driver, like most home made current drivers, uses a trimpot to adjust the current set point. In order to lower the noise of our current driver our design implements a precision digital to analog converter (DAC) to program the current set-point. By eliminating the trimpot we have taken out a source of noise and instability.

We control the DAC with a micro-controller to allow for remote and accurate control of the

laser current. This makes it possible to place the current driver in close proximity to the laser diode, which significantly reduces current noise and lowers phase shifts when using the modulation input. We also gain more accurate control of the current since the DAC references each current set point with a 16-bit number, making it possible to return to a previous current set point with extremely high precision.

We also implemented several smaller design changes to make a versatile controller and updated other circuit components to improve the noise and stability. Most notably our design provides a single printed circuit board (PCB) that allows for either polarity of current, can drive high current devices, and uses surface mount technology to reduce stray capacitance and stray inductance. While many details are included here, please consult [92] for PCB files, tips, and specific instructions for constructing a driver.

5.1.1 Design and Construction

In Figure 5.1 we have broken down the schematic of the circuit into several sections. Each section performs an important function and we will address each separately. An image of the PCB designed for the controller is shown in Figure 5.2.

Current Regulation

The heart of the current driver is the "current regulation" section. This part of the circuit consists of an op-amp in a current source application. A regulated voltage, V_{reg}, is placed across a resistor, R_s, a MOSFET, and the laser diode in series. The op-amp controls the gate on the MOSFET by monitoring the voltage where the resistor R_s connects to the MOSFET. In this configuration, the op-amp forces the current through the resistor, MOSFET, and laser diode to remain fixed regardless of the load impedance. If V_{set} is the voltage on the positive input of the op-amp, then the current through the diode is $I = (\pm V_{reg} \mp V_{set})/R_s$ where the upper signs correspond to a cathode grounded

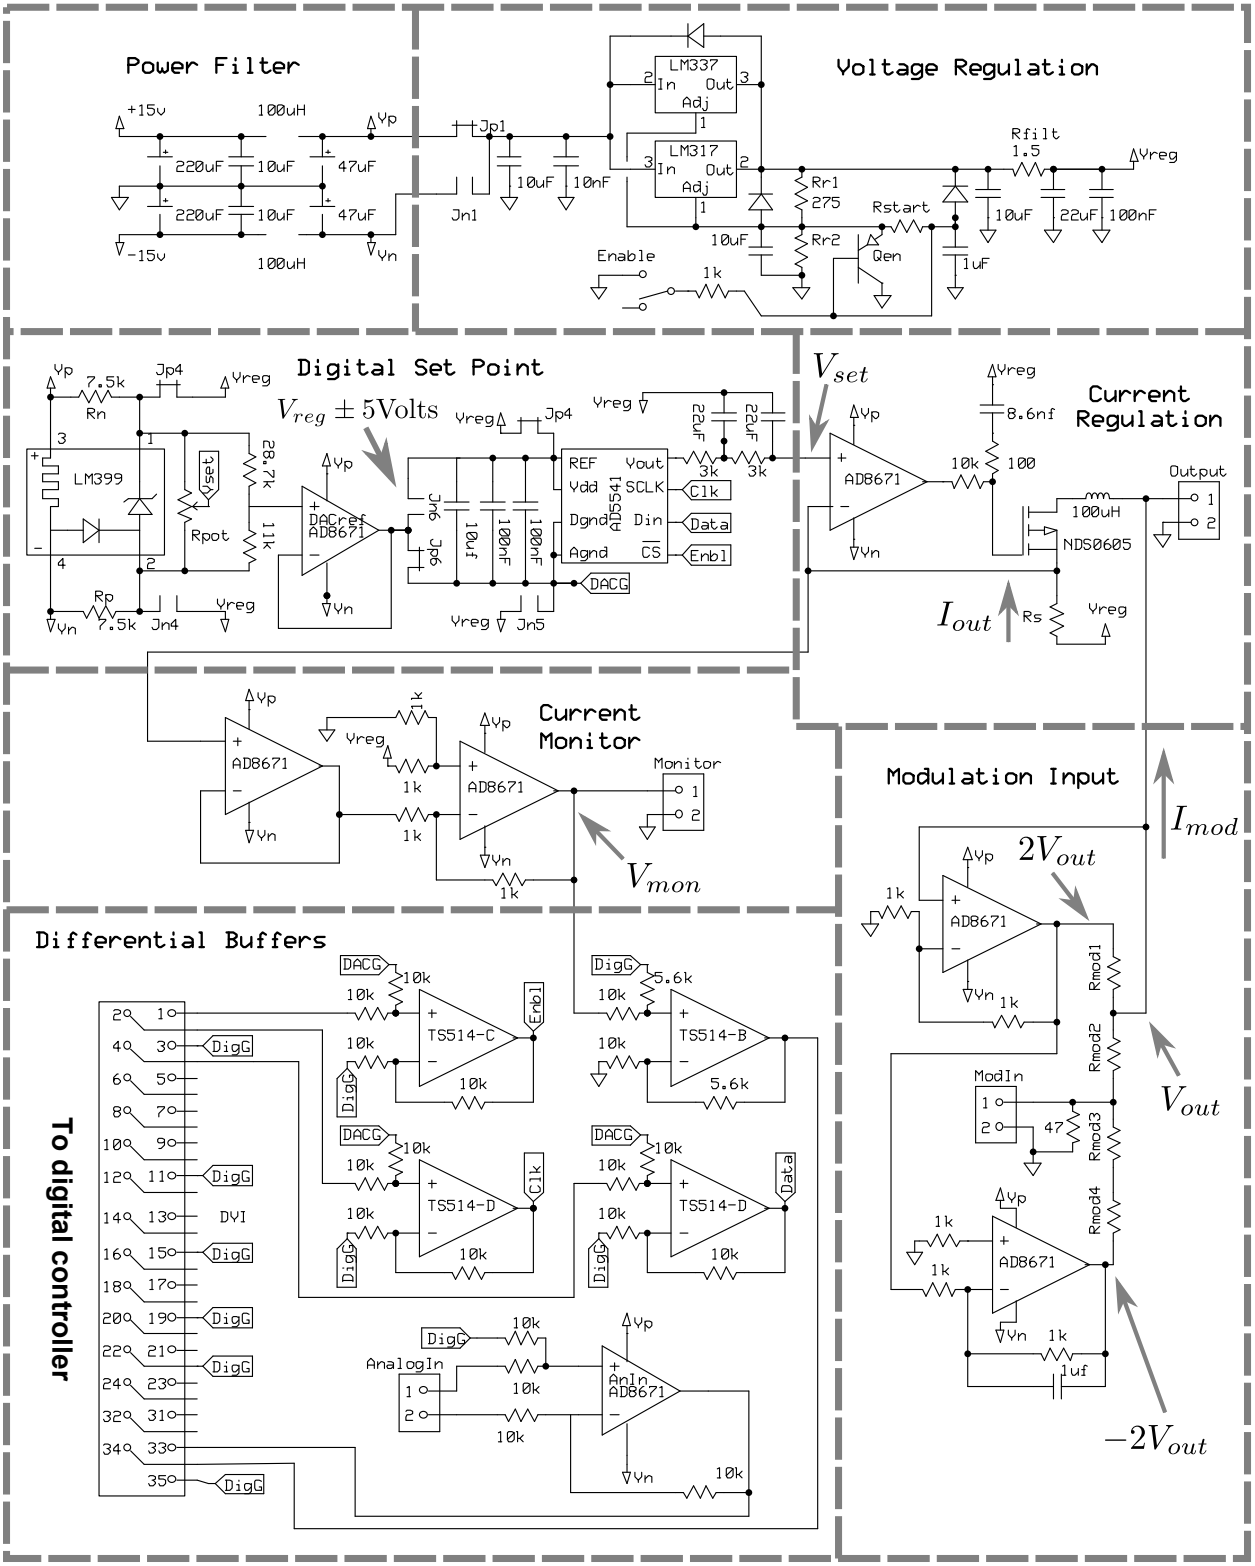

Figure 5.1 Schematic of the laser diode current driver circuit. The circuit is divided into separate functions for clarification and easy referencing.

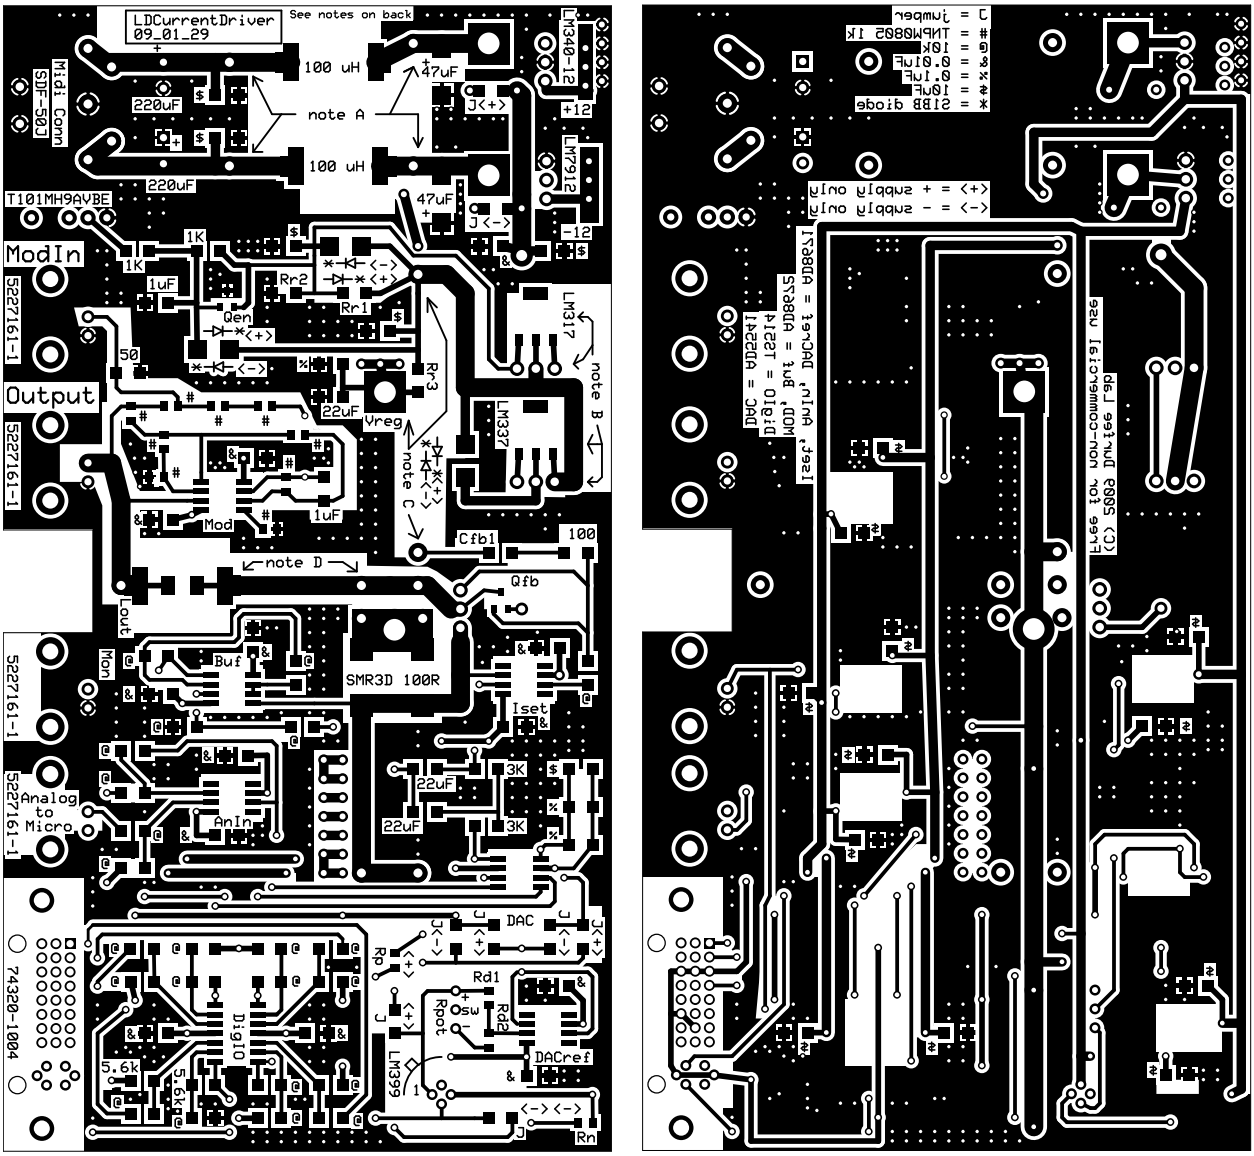

Figure 5.2 Images for the current driver PCB are shown with the top of the board on the left, and the bottom of the board on the right. The top section of the board is where the power enters and is filtered. Directly beneath is the section of the board labeled "Voltage Regulation" in Figure 5.1. The area just above and left of board center is the "Modulation Input"' section of Figure 5.1. The modulation signal enters through the connector labeled "ModIn". Beneath this lies the section of the board referred to as the "Digital Set Point" and "Current Regulation" sections of Figure5.1. The resistor pads labeled "SMR3D 100R" denote the resistor "R_s" in Figure 5.1. The op-amp "Iset" controls the voltage on the gate of the MOSFET "Qfb," which regulates the current through the resistors. An LM399 (whose pads are seen at the bottom of the board) and an op-amp (pads labeled "DACref") create the 5 V reference for the DAC. The op-amp labeled "Quad" buffers the digital signals to the DAC, and the two op-amps to the bottom right labeled "buffer" and "Dif" process the signal for the current monitor. Through holes are placed near surface mount pads where high current components may need to be placed instead.

driver configuration and the lower signs to an anode grounded configuration.

As in the design in [91], an inductor is placed after the MOSFET to limit high-frequency current noise and to de-couple the current regulation from any intentionally-applied current modulation. Also, to reduce high frequency noise and prevent the circuit from oscillating due to the MOSFET's gate capacitance, some resistors and a capacitor have been placed between the op-amp and the gate of the MOSFET. The board is configured so that a through-hole inductor may be used as well and the board can accommodate MOSFETs in either SOT-23 or TO-220 packages to accommodate driving a wide range of output currents.

The resistor R_s is a key component that should be considered carefully — stability of the current driver can be no better than the stability of this resistor. For low current applications we use Vishay SMR3D precision resistors. These parts are surface-mount components for reduced noise pick-up, and have a low temperature coefficient of just 2 ppm/°C. Through holes are provided on the PCB for installing low value power resistors for high current applications. Selecting the proper resistance is important in order to provide a large enough voltage drop for accurate current regulation at the target current while not reducing the compliance voltage of the driver below what is necessary to power a particular laser diode.

The other extremely critical component in this section of the circuit is the op-amp. The op-amp selected was an AD8671 from Analog Devices [93]. Similar to the LT1028 used in the original Hall design [94], this op-amp has voltage noise levels comparable to the Johnson noise of a 50 Ω resistor. However, the AD8671 has slightly better current noise and slightly poorer voltage noise than the LT1028.

In practice, as well as in numerical SPICE models, the AD8671 was less prone to oscillate in this circuit than the LT1028. Although the AD8671 has a considerably smaller gain-bandwidth product than the LT1028 used in the original Hall design, the filtering capacitors and resistors that would be needed to stabilize the circuit effectively eliminate any additional bandwidth that would

be gained by using the LT1028. Furthermore, the inherent low noise of our circuit and the use of the de-coupling inductor makes higher bandwidth in this part of the circuit unnecessary.

The AD8671 chip also has a much higher input impedance (both common and differential mode) than the LT1028, and comparable input capacitance. The higher input impedance limits changes in current leakage through the op-amps inputs due to thermal and other factors, which would create drift in the driver's output.

Voltage Filtering and Regulation

The "power filter" section of the circuit generates stable, filtered voltages for the rest of the circuit. Power lines are first filtered by a series of capacitors and inductors to provide voltages that are sufficiently stable for powering all of the chips on the board. This filtered power is further defined by an adjustable voltage regulator to generate the stable voltage, V_{reg}.

Regulators can be in either a TO-220 or SOT-223 package. For applications up to 1.5 A we use LM317 chips for positive current and LM337 chips for negative current. The regulated output level of these chips is determined by a voltage divider created by the resistors labeled R_{r1} and R_{r2} in Figure 5.1. For positive currents up to 5 A we use an LM338 and for high current negative applications we splice in a Darlington pair follower after the low current voltage regulator. An appropriate resistor should be chosen for R_{filt} based on the desired current output of the driver. It is also important to note that the necessary heat sink requirements are met for chips in higher current applications.

Soft Start Feature

The circuitry surrounding the adjustable voltage regulators in the "voltage regulation" section is part of a slow start mechanism shown in [95]. This soft-start feature helps protect the laser diode from damage. In the case of a power failure it allows the digital controller time to reset the current

set-point before the voltage turns fully on when power is restored.

When the "*Enable*" switch is grounded, the current driver is "disabled." In this case, the resistor R_{r2} is effectively reduced to the value of R_{start} and V_{reg} becomes too small to forward bias the laser diode.

When the switch is open the 1 μF capacitor prevents the current flowing through R_{start} from abruptly stopping. This causes the voltage regulator's output to slowly increase to its final value, providing a "soft start" effect. Due to the non-linear nature of this circuit, the start time is considerably larger than just the *RC* time constant of the resistor R_{start} and the 1 μF capacitor. Using a value of $R_{start} = 1$ kΩ and an MMBT2907A transistor in our setup, we measured a start-up time constant of 11.5 ms.

Current Set Point

The current set point is determined by the voltage applied to the non-inverting input of the op-amp in the "current regulation" section. It is generated by an AD5541 16-bit precision digital to analog converter (DAC). This device produces an output voltage that can swing from ground up to 5 V. Since the output current of the laser current controller is proportional to the voltage drop across R_s from V_{reg}, the DAC is biased such that it's maximum/minimum output voltage is V_{reg} for positive/negative supplies.

To bias the DAC's range, the chip's power and "ground" are provided by V_{reg} and a stabilized reference voltage that is 5 V lower or higher than V_{reg} depending on polarity. The voltage which serves as the DAC's "ground" is labeled V_{DACG} in Figure 5.1. To stabilize this voltage difference, a LM399 Zener diode was used to create an extremely stable 6.95 V difference from V_{reg}. Precision low temperature coefficient resistors, $Rset1$ and $Rset2$, are then used to make a voltage divider to produce the required 5 V difference.

Any noise between the voltage output of the DAC and V_{reg} will result in noise between the

voltage output of the op-amp in the "current regulation" section of Figure 5.1 and V_{reg}. This noise will subsequently show up on the current output of the driver. To eliminate this source of noise, low pass filters on the DAC's output are tied to V_{reg}. Therefore, any fluctuation in either V_{reg} or the output of the DAC will be common mode and the output of the driver will remain unaffected.

For applications where the noise, stability, and repeatability of the current driver are not a great concern, a potentiometer can be used in place of the DAC. In this configuration the voltage divider, op-amp, and DAC in the "digital set point" section of Figure 5.1 would not be installed. The sweeper of the potentiometer would be connected to the 3 kΩ resistors in series, and the posts would be connected across the 6.95 V drop created by the LM399 Zener diode.

Current Modulation

The modulation input uses an active circuit to add a modulation signal on the DC output of the driver. In the "modulation input" section of Figure 5.1 the top op-amp creates a voltage that is twice the output voltage of the driver, which is labeled "2Vout" in the figure. This output is separated from the modulation input connector by two precision 1 kΩ resistors labeled "Rmod."

If a modulation signal is present on the connector "ModIn" then the current added to the driver's output current can be easily calculated. Since the current labeled "Imod" in the "modulation input" section is the sum of the two currents entering the node labeled "Vout" between the two precision resistors, then

$$Imod = \frac{Vout}{Rmod} + \frac{Vmod - Vout}{Rmod} = \frac{Vmod}{Rmod}$$

where "Vmod" is the modulation signal voltage. This modulated current adds cleanly to the current driver's output.

There is a second op-amp shown in the "current modulation" section of Figure 5.1 below the first. Its output is labeled "$-2V$out." This op-amp produces a voltage that precisely balances the output of the first op-amp, which is also separated from the "ModIn" connector by two precision

1 kΩ resistors labeled "Rmod." This ensures that the "ModIn" connector voltage will be at ground when the modulation input is floating.

The modulation circuit allows modulation speeds greater than the bandwidth of the op-amps. In the limit of high modulation frequencies, instead of assuming that the output of the op-amps track with the modulated voltage we can assume that the op-amps "see" an average voltage from the modulated signal. In this case the modulated current *I*mod becomes

$$Imod_{hf} = \frac{Vmod}{Rmod}(1 - \frac{2Z}{Rmod}) \tag{5.1}$$

where Z is the impedance of the laser diode. Note that in this limit the modulation current $I\text{mod}_{hf}$ is not limited by the bandwidth of the op-amps.

Importance of Component Selection

Unlike the design in [91], we have used the same op-amp in the current regulation, modulation input, and monitoring sections of the circuit. The op-amp in the "current regulation" section is important because it's output voltage noise affects the gate of the MOSFET controlling the driver's current. Also, it is important that this op-amp has low bias and high input impedance for the current set point to be accurate. The op-amps in the "modulation" input section of the driver are directly connected to the driver's output and need high input impedance to avoid sinking current meant for the laser diode. The first op-amp in the "current monitor" section needs high input impedance as well.

For these reasons, we decided to use the AD8671 in any section connected to the driver's output. The input impedance for both common and differential mode of the AD8671 are in the giga-ohm and mega-ohm range respectively, and its 10 MHz gain bandwidth product is more than enough for this application. Plus, the AD8671 does not require any sort of trim or compensation to achieve low bias and low overshoot in our application.

Using the same op-amp also matches parameters such as the noise and input impedance between the different stages so that the controller is not limited by a "weakest link", which behaves considerably worse than the rest of the circuit. This makes good sense, especially considering the relatively low cost of the precision op-amp which we have chosen.

Current Monitor

The "current monitor" section of the Figure 5.1 contains circuitry for monitoring the driver's current output on an oscilloscope or voltmeter. The voltage drop across the resistor R_s is effectively buffered and passed off the board to an external device. Since R_s is a precision resistor, this signal accurately provides a voltage directly proportional to the output current of the driver.

Buffering I/O Signals

The use of a microprocessor to set the current, and the desire to monitor the voltage drop across R_s require that the board be connected to other remote boards or external equipment such as a microprocessor, a scope, or a multi-meter. The use of differential amplifiers effectively isolate the current driver from the ground noise of other equipment, thereby protecting the circuit from the noise of other instruments as well as ground loops.

The op-amps shown in the "differential buffers" section of Figure 5.1 re-reference the three digital lines passed to the DAC from the microprocessor to control the current set point and the analog line passed from the current driver to the microprocessor to monitor the voltage drop across R_s to the appropriate ground. Additionally, the input labeled "analog-to-micro" allows any off-board analog signal between 0 and 5 V to be buffered and sent to the microprocessor. This signal can be used to digitally implement a laser stabilization technique such as the one discussed in [96], or for processing any other analog signal the user may desire.

The op-amps that are used for buffering are not directly connected to the driver's output and,

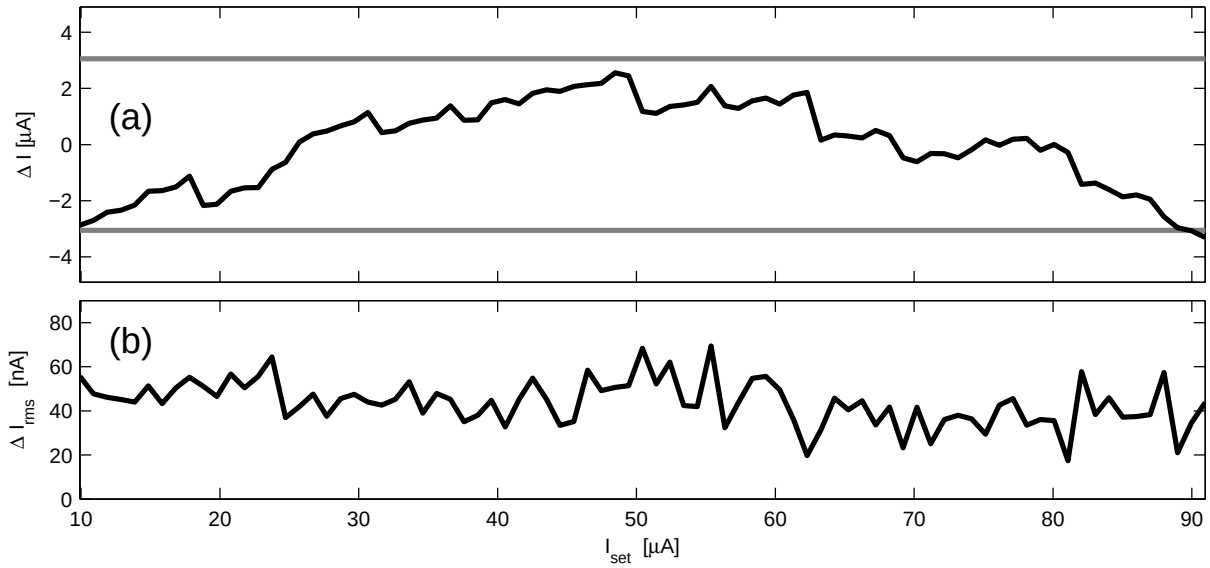

Figure 5.3 The top graph is the accuracy of the DAC while the bottom graph is the repeatability. The ability of the DAC to return to a given value is much better than the DAC's overall accuracy. While the repeatability is well within the specifications, the accuracy is slightly worse than specified.

therefore, do not have to be as carefully selected. However, we found that certain op-amps, such as the AD8674 and TL074, were unable to drive the capacitance of the DVI-I cable used to interface with the microprocessor circuit and required an additional 10 kΩ on their output in order to provide the correct signal to the processor. The TS514 chip, however, was able to drive the cable's capacitance.

5.1.2 Accuracy and Repeatability

The accuracy and repeatability of setting the current digitally is shown in Figure 5.3. The accuracy of the DAC is rated such that for a given setpoint the output of the driver is within 2 μA. As seen in the upper graph of Figure 5.3, the response of our circuit mainly falls within these bounds with greater error at the ends of the DAC's range.

The repeatability of setting the current to a specific point is excellent. To test this, several points

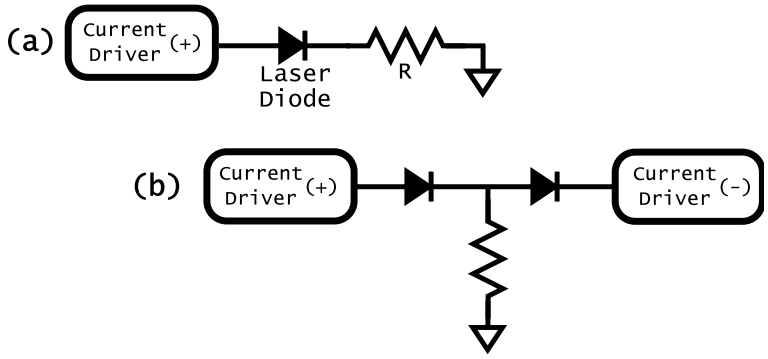

Figure 5.4 Noise measurement setups. The measurement setup used in [3] is shown in the upper diagram (a). The lower diagram (b) is the improved setup used to further lower the noise floor [4].

spanning the range of the driver were returned to 10 times each from a specified point. As shown in Figure 5.3, the current can be reset to the same point from a different value to within a few nA.

Noise and Stability

To characterize both the long term stability and the high frequency noise of the current driver initially we placed a 2.5 Ω precision resistor in series with the laser diode and measured the voltage drop across the resistor as shown in Figure 5.4 (a). From this data we extracted the current spectral noise density shown in Figure 5.4. The voltage across the resistor was measured with a Lecroy Waverunner LT 354 oscilloscope and Fluke 8505 multimeter, and the noise spectral density was calculated from these measurements. However, the measured noise of the driver was at the measurement noise floor for most of the frequency range measured, and therefore only represented an upper limit on the current driver noise and drift. The noise floor was mainly limited by the fact that the small current noise and drifts where on top of a large bias signal.

In order to lower the noise floor and get a more accurate measure of the noise spectral density, we used a revised setup with two current drivers, shown in Figure 5.4 (b). Each current driver was configured using lower current components and Vishay SMR3D precision resistors for R_s. One

current driver was set to +50 mA while the other was set to -50 mA. In this setup, almost all of the current flows from one driver to the other through the two laser diodes, while uncorrelated current noise and drift induces a voltage across a 1 MΩ resistor, which is measured to determine the noise spectral density. This setup greatly reduces the bias in our noise/drift measurement, and allows us to use a larger resistor to increase the size of the signal, allowing for a more accurate measurement. Unfortunately, the larger resistor also increased the time constant for our measurements, limiting us to measurements below a few kHz.

The results of our measurements, as well as data from [1], are shown in Figure 5.5. These curves represent the composite of several independent measurements, each spanning a small portion of the frequency range displayed. The part of the spectrum below 2 Hz was taken using a Fluke 8505A digital multimeter, and the higher-frequency data was taken using a Lecroy Waverunner LT 354 oscilloscope. To prevent aliasing, the noise spectra in different frequency ranges were calculated from different data sets taken at different sample rates, and for each data set the signal was conditioned with a SR560 preamp configured with a 12 dB/octave low-pass filter. The preamp was set to unity gain for all data sets, but the filter cutoff frequency was changed for each data set such that it was at least a factor of 20 below the sample rate of the multimeter or oscilloscope. After calculating the noise spectral density from each data set, noise at frequencies equal to or higher than one tenth of the filter cutoff frequency were considered invalid and thrown out. The fact that the different segments of data align seamlessly is an indication that significant aliasing effects are not present. The discontinuities in the noise floor measurement are due to the fact that the noise floor depends on which instrument is being used and the instrument settings. The most notable jump is at 2 Hz, where we switched from the voltmeter to the oscilloscope. Since at all frequencies the measurement noise floor is well below the measured noise of the current driver, this does not affect our measurement.

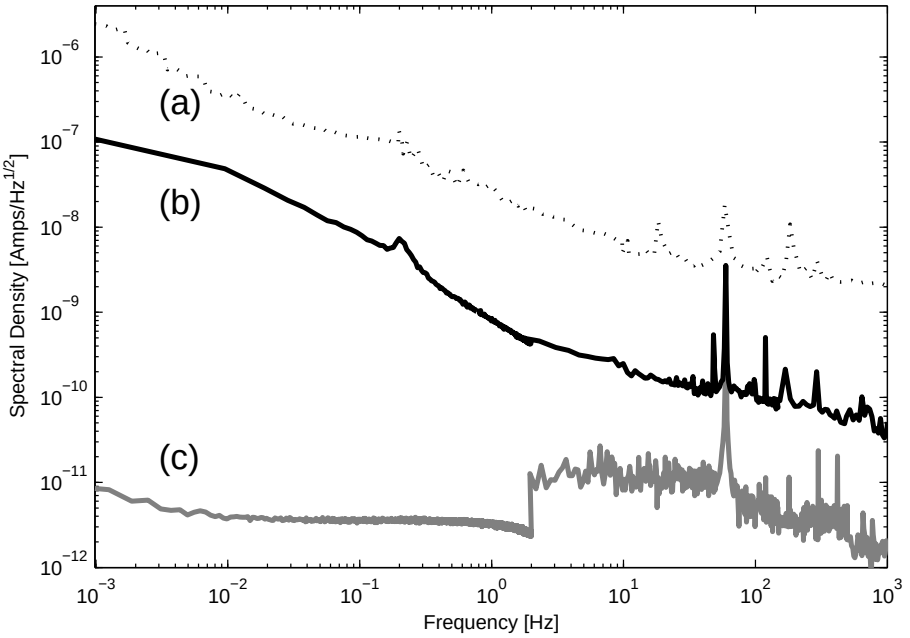

Figure 5.5 Noise current spectral density. The data from [1] is shown as a dotted line (a). The new noise measurements are shown as a black line (b) and the noise floor of the new measurement setup is shown as a gray line (c) [4].

Temperature Coefficient

In order to characterize the temperature dependence of the driver, thermistors were placed at different parts of the PCB and monitored along with the current. The current was determined by measuring the voltage drop over a 10 Ω precision resistor.

Figure 5.6 shows long term current drift (A) compared to temperature measurements taken at R_s (the resistor determining the output current) and the DAC (B), and the box housing the driver (C). From the correlation of these graphs we infer a temperature coefficient of a few parts per million (1.7 ppm/C). Given that the set resistor R_s in Figure 5.1 is rated with a temperature coefficient of 2 ppm/C, this was expected. However, we also found an unexplained drift uncorrelated with the temperature of the same order.

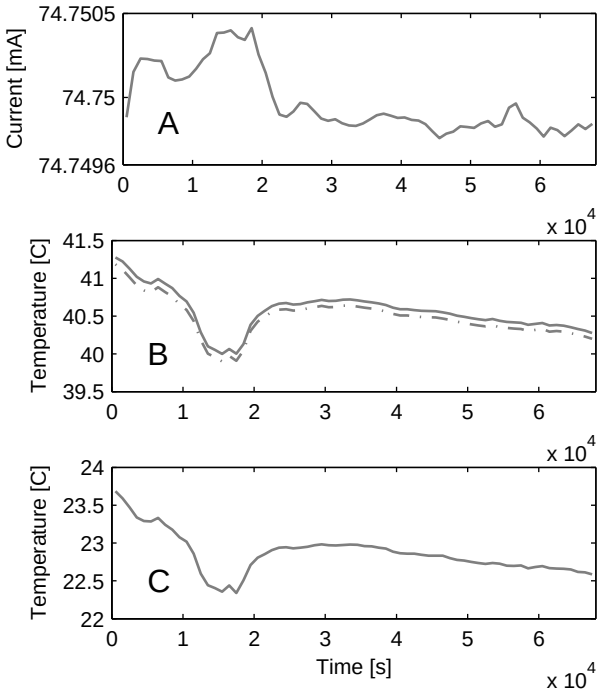

Figure 5.6 Temperature dependence of current driver output.(A) The current output of the driver through a laser diode. (B) Corresponding temperature measurements taken at the current set resistor Rs (solid line) and the DAC (dot-dash line). (C) Corresponding temperature measurements taken at the box housing the driver.

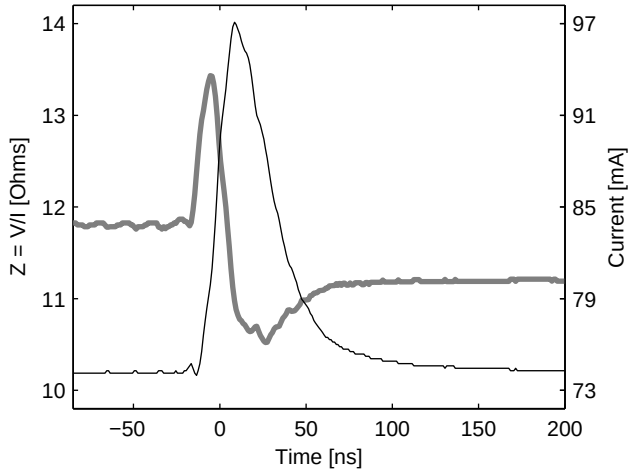

Figure 5.7 The transient response of the current driver to a total load change of roughly 3%. The black line is the current through a 10 Ω resistor, and the gray line is the impedance of the load.

Transient Response to Load Impedance Variation

The transient response of the current driver was evaluated using a low current positive driver. A 10 Ω resistor was put in series with a VZN1206 MOSFET while controlling the transistor gate to vary the overall impedance by ∼3% with a step function. The black line in Figure 5.7 is the current through the resistor determined by measuring the voltage across it, and the thick gray line is the impedance of the entire load calculated by dividing the voltage drop across the transistor-resistor pair by the current found in the resistor. The distance between the peaks in the two traces is just 14 ns. We extracted a time constant of 70 ns by fitting the current signal to an exponential decay after 100 ns where the impedance has settled to a constant value. This indicates excellent response time for a current driver.

Modulation Performance

We measured the modulation response of the current driver by terminating the output with a 50 Ω resistor and comparing the amplitude and phase of the modulation signal on the current driver

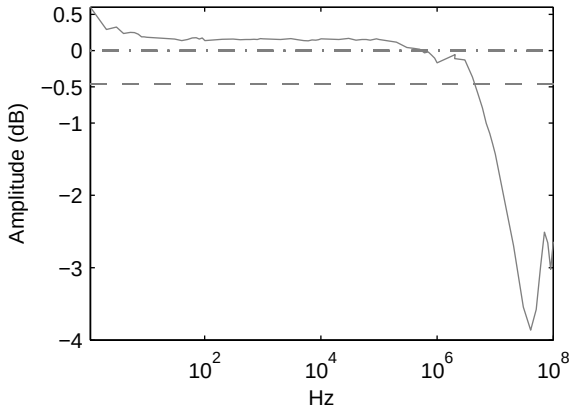

Figure 5.8 Bode plot of the response of the modulation circuit. The solid line is the amplitude in dB of the response of the driver relative to the theoretical low frequency response. The dot-dash line is the theoretical low frequency amplitude response, and the dashed line is the theoretical high frequency response.

output to the driving signal on a high impedance oscilloscope. The amplitude of the driver's output signal is compared to the original signal in the Bode plot in Figure 5.8. The response of the circuit hits a 3 dB point, where the amplitude of the output signal has rolled off by a factor of $\sqrt{2}$, near 50 MHz.

5.1.3 Conclusion

Our design for a current driver not only provides stability and low noise, but it also allows for repeatability, accuracy, and remote location through digital control. It features high bandwidth modulation, outperforms commercial systems, and it was easily implemented as a lab standard. This performance is mainly due to the use of a DAC with extremely low drift as well as advances in the technology of surface mount ICs and in their availability in small quantities. A technique for surface mount soldering custom PCBs in a laboratory environment is described in Appendix B.

5.2 Temperature Controller

Controlling the temperature of a laser diode is important not only for protecting the diode from overheating, but also in stabilizing the wavelength of the diode[1]. The schematic for our temperature controller is shown in Figure 5.9. The heart of the controller is a simple Wheatstone bridge where one pair of the resistors in the bridge is simulated by a DAC which selects the temperature setpoint. On an opposite leg of the bridge the resistor is comprised of a thermistor which is off the PCB mounted near the laser diode. When a voltage difference is created across the bridge due to the imbalance of the DAC voltage and thermistor value a simple proportional and integral stage feedback to the voltage on a thermo-electric cooler mounted under the laser diode.

5.3 Lock-circuit PID Controller

Many common laser locking-circuits consist of only an integral stage to create feedback to the laser head. For a laser who's linewidth is already smaller than the linewidth of the cavity or the absorption profile it is being locked to, this works well enough. However, to accommodate a large range of applications and provide the tightest lock possible in simple ones, we have also included a proportional and a differential stage in our controller.

At its heart, the lock-circuit is no different than any PID controller as seen in the left hand box of Figure 5.10. It has a proportional, integral, and differential stage whose gains are tuned to optimize the response of the servo system. The bandwidth of the circuit can be maximized by the use of high-speed low-noise surface mount chips and good PCB design. As seen in the image of our PCB for the circuit in Figure 5.11 the feedback loop on each amplifier has been minimized to reduce stray capacitance and inductance. To further minimize inductance, the gains on each stage

[1]Typically, red and infra-red laser diode outputs change by $\sim$ .3 nm per degree Celsius for a given current. The 405 nm diodes used in our apparatus change by $\sim$ .1 nm per degree Celsius.

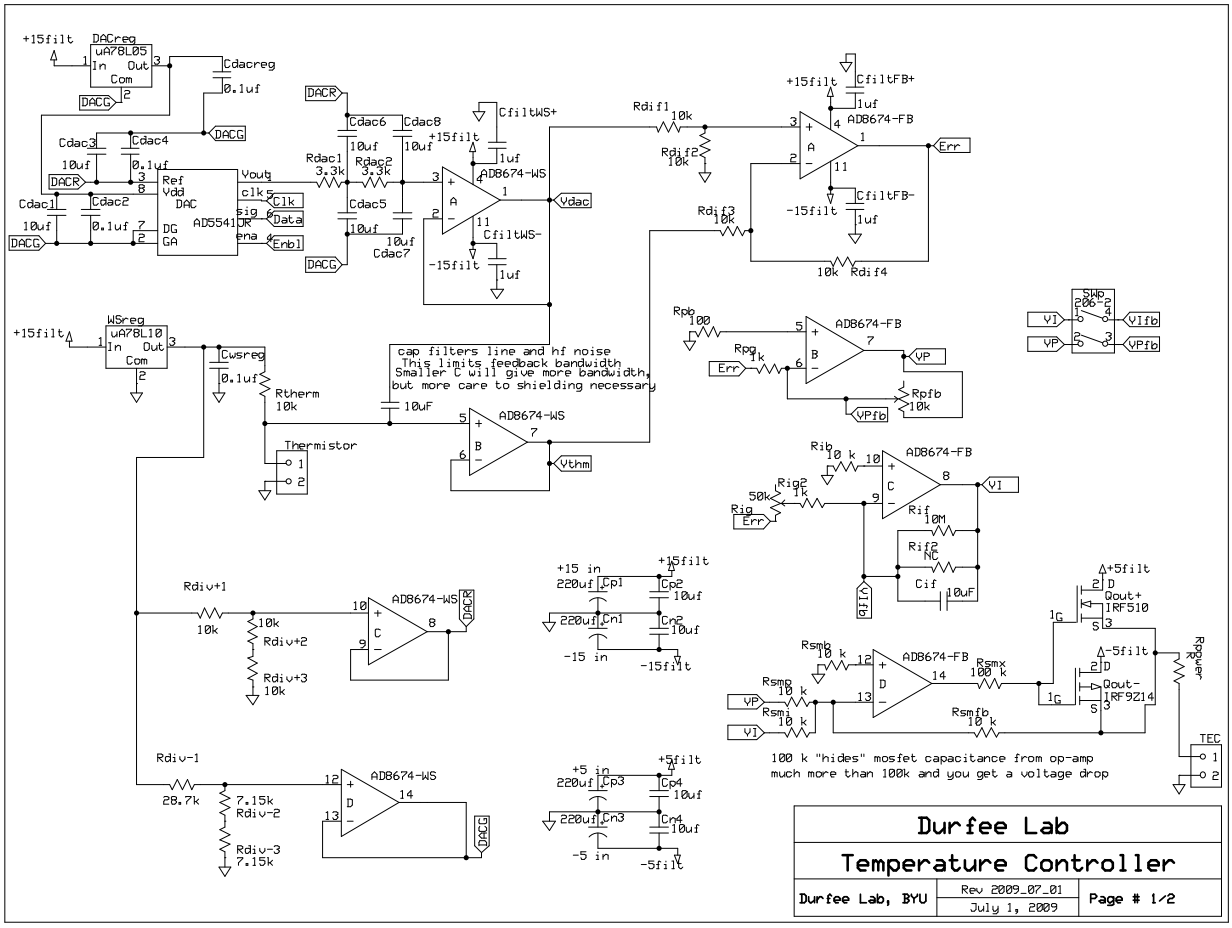

Figure 5.9 Schematic for the temperature controller.

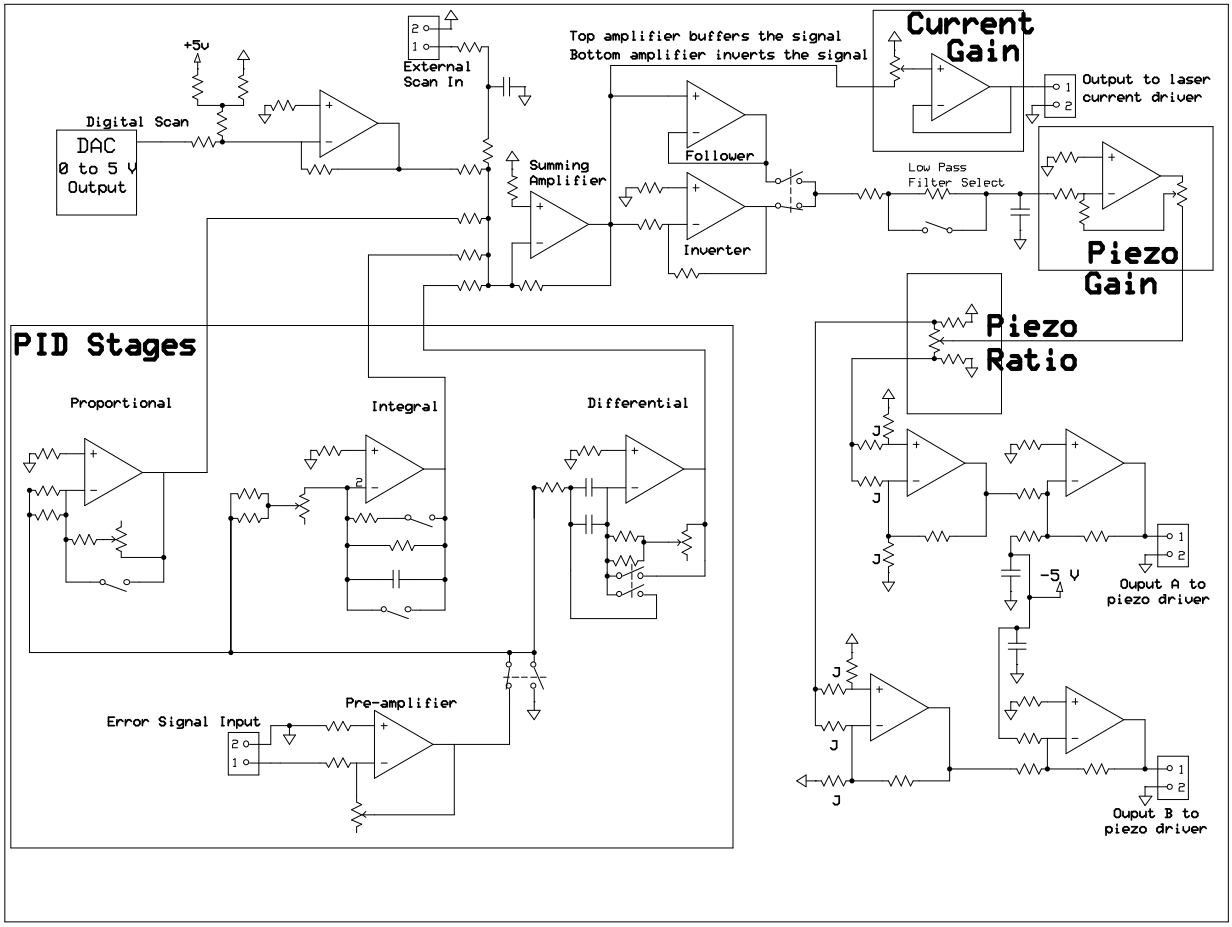

Figure 5.10 Schematic of the scan-balancing lock-circuit used in the experiment. The area enclosed by the large box on the left represents the proportional, integral, and differential gain stages. The outputs of these stages are summed together along with external and on-board scan inputs at the "Summing Amplifier" op-amp. The summed signal is then split. One signal path goes through a gain stage to the modulation input on a laser diode current driver. The other goes to the piezo amplifiers for the grating mount on the ECDL. The box on the right had side labeled "Piezo Gain" set the overall gain for the piezo signals and the box labeled "Piezo Ratio" sets the voltage difference between the two piezo signals for the grating mount. The circuit can be scanned by use of a digital to analog converter or through the input of an external function generator.

are set with 1 turn trimpots. Short trace lengths between chips also helps to improve the circuit performance.

Similar to the current and temperature controllers, the PID circuit can be digitally controlled. A separate unit consisting of a programmed micro-controller can be used to manipulate an on-board digital to analog converter (DAC) and four digital switches. The DAC is used to set and scan the laser frequency while the switches control filters and various aspects of the PID stages. This remote control of on-board functions allows the controller to be housed in the same box as the laser current driver keeping servo delay time to a minimum. Additionally, the micro controller could potentially be programmed to scan the laser and lock to an error signal on its own, as well as monitor and re-lock the laser if the lock is lost.

Figure 5.11 shows the actual PCB board for the PID controller. Values of individual components will vary depending on the integrated chips that are used and the laser system that the board is used in. In our application we used AD8671 and AD8672 op-amps for all of the gain stages. The DAC is an AD5541 and all the digital switches are contained within an ADG453 IC.

5.3.1 PID Processing

The box enclosing the lower left side of Figure 5.10 represents the PID control portion of the circuit. The error signal from the homodyne detector enters the board and goes through a pre-amplifier that passes the signal to a proportional, an integral, and a differential amplifier. The inputs to these stages are connected to a digital switch that either connects them to the error signal or to ground depending on whether the circuit is locking or scanning the laser respectively. The gain on each stage is set by a one turn trimpot. The signals from each stage are then sent through separate resistors to a summing amplifier. The ratio of the respective input resistor after each stage to the feedback resistor for the summing amplifier along with the local gain for each amplifier determines the overall gain for each stage.

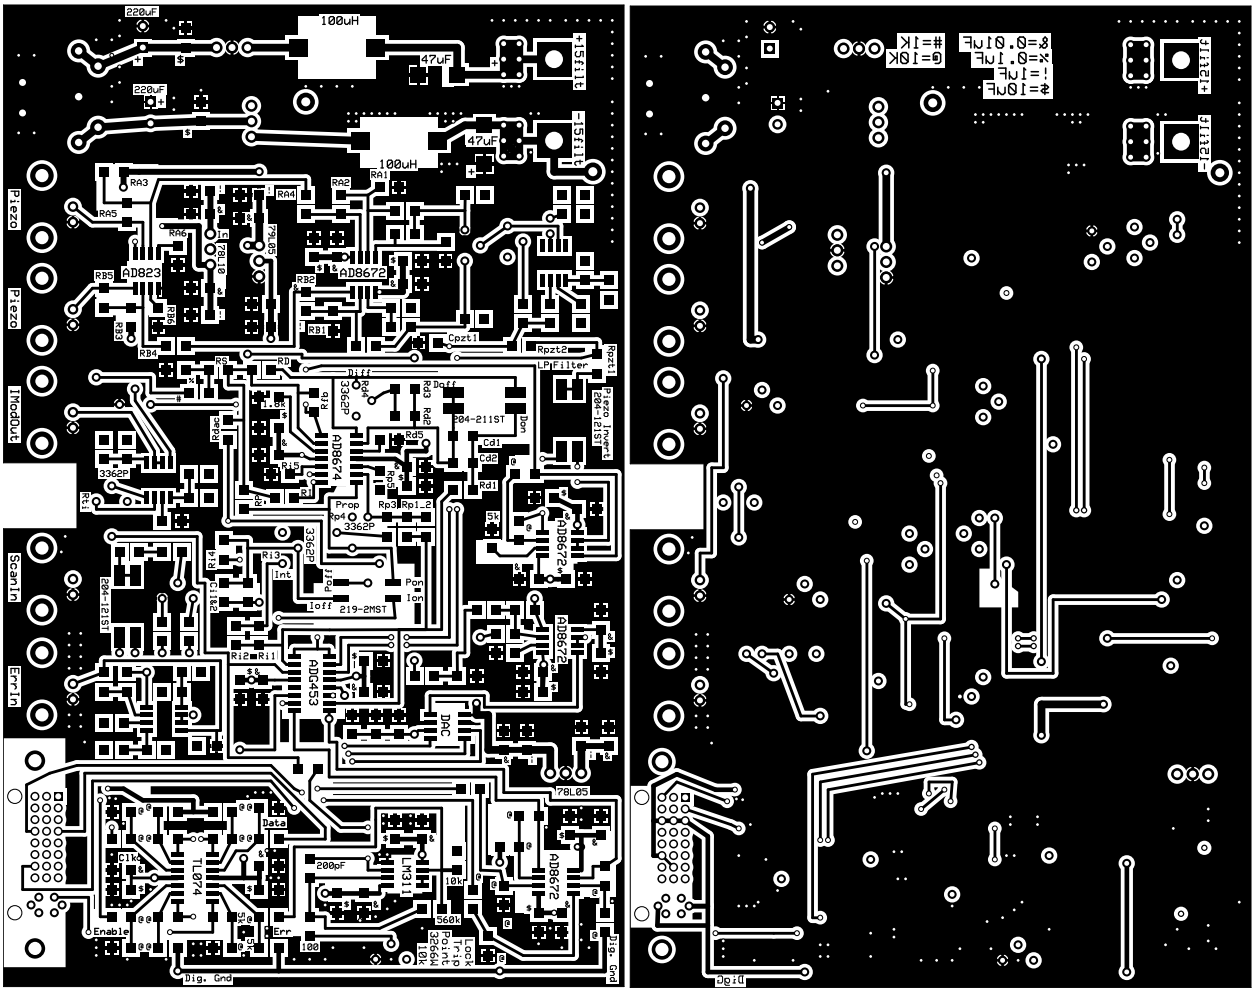

Figure 5.11 Image of the PCB used for the scan-balancing lock-circuit. The top of the board is the left image and the bottom of the board is the image on the right. The error signal enters through the connector labeled "ErrIn". It is then amplified by the chip labeled "Trans" and passed to the proportional, integral, and differential gain stages (labeled "Prop", "Int", and "Dif" respectively). These stages along with signals from the DAC and the connector "ScanIn" are then combined by the summing amplifier "Sum". The output of this amplifier goes to a gain stage for laser current feedback and exits the board through the connector "ModOut". The summed signal is also buffered or inverted by an op-amp before passing through one of two low pass filters, which are selected by a digital switch. The filtered signal then passes through a gain stage and is split into two signals of fixed ratio for controlling a piezo-electric kinematic mount. The piezo signals are further processed by op-amps before leaving through the connectors labeled "Piezo" to a piezo driver. Three manual switches are used to individually toggle the integral, proportional, and differential stages and one manual switch is used to select the polarity of the piezo signals. All of the digital switches in Figure 5.11 are contained in the chip labeled "ADG453". Op-amps buffer the digital signals to the DAC and the digital switch which enter through the DVI-I connector at the bottom left hand corner of the board.

5.3.2 Laser Frequency Scan-balancing Function

The laser can be scanned by the DAC or through an external function generator whose signal is brought in at the "External Scan In" labeled in Figure 5.10. The DAC puts out a value between 0 and 5 V, which is shifted to -2.5 V to 2.5 V by an op-amp. The external scan signal enters the board and goes through a low-pass filter set appropriately for the maximum desired scan speed. Both signals then also pass through the summing amplifier in similar manner to the PID circuitry with the overall gain for each determined by a resistor on the input of the summing amplifier.

The summing amplifier adds all of the incoming signals together according to their respective gains and the signal is sent to a scan-balancing circuit shown on the right hand side of Figure 5.10. This part of the circuit helps to keep the laser diode from mode-hopping while it is being scanned or locked. As the laser frequency scans, the length of the laser diode cavity and the angle of the diffraction grating with respect to the diode must change slightly to keep the cavity resonant with the laser frequency. Slightly different gains are needed for the current feedback and for controlling each side of the piezo mount to produce the right combination of rotation and translation of the grating in an ECDL simultaneously.

The laser current feedback is set by the gain of the op-amp labeled "Current Gain" at the top of Figure 5.10. The piezo feedback can be buffered or inverted as necessary before its overall gain is set at the op-amp labeled "Piezo Gain". The ratio of the two piezo signals sent to the grating mount of an ECDL is then set at the box labeled "Piezo Ratio". The two outputs of the "Piezo Ratio" box can then be buffered or inverted as necessary at the following op-amps. Since piezo electric crystals cannot take negative voltages the output signals are shifted to so that the output voltage ranges from $0-10$ V. This signal is then sent to an amplifier circuit for the piezo crystals.

The scan-balancing circuit also contains a very selective notch filter that functions only when the circuit is locking the laser to prevent a mechanical resonance in the diffraction grating's mount from destabilizing the servo. Using this scan balancing circuit allows mode-hop free tuning of a

laser diode over many GHz depending on the diode. While typical scans with the circuit tend to be $2-10$ GHz, the longest mode-hop free scan achieved was ~ 30 GHz with the AR coated 921 nm master laser diode.

5.4 Photodiode Homodyne Detector

A homodyne detector is a device that is used to detect a modulated signal. This is done by mixing the signal with the original modulation frequency to produce a simple DC signal.

This technique can be used for locking lasers. The laser's phase or frequency can be modulated at radio frequencies (RF) and detected with a photodiode. The photo-diode's signal can then be mixed with the modulation frequency to produce and error signal for locking the laser. This detection can be done by stringing together several modular components with coaxial cable. Our device condenses the detection into one unit, minimizing noise and time delay in the laser servo.

Figure 5.12 shows the schematic for the homodyne detector. On the left of the figure, a photodiode is reverse biased to decrease its capacitance and increase its frequency response. Then the signal from the photodiode is converted to a voltage by a transimpedance amplifier. The signal is further boosted by three amplification stages. A local oscillator signal from off the board is sent to the mixer along with the amplified signal, and the mixer output is sent to a lock circuit.

The physical layout of the printed circuit board (PCB) for the homodyne detector is shown in Figure 5.13. It is entirely packaged on a 3.5×5 inch PCB. The signal originates at the photodiode's pad in the center of the right side of the board. It is then passed through a transimpedance amplifier consisting of a single AD8015 op-amp. The amplification stages use OPA657 op-amps and the mixer is an SYPD-1 from Minicircuits. The connector in the top left hand corner is for power, and the pads along the top of the board are for filtering and regulating the power.

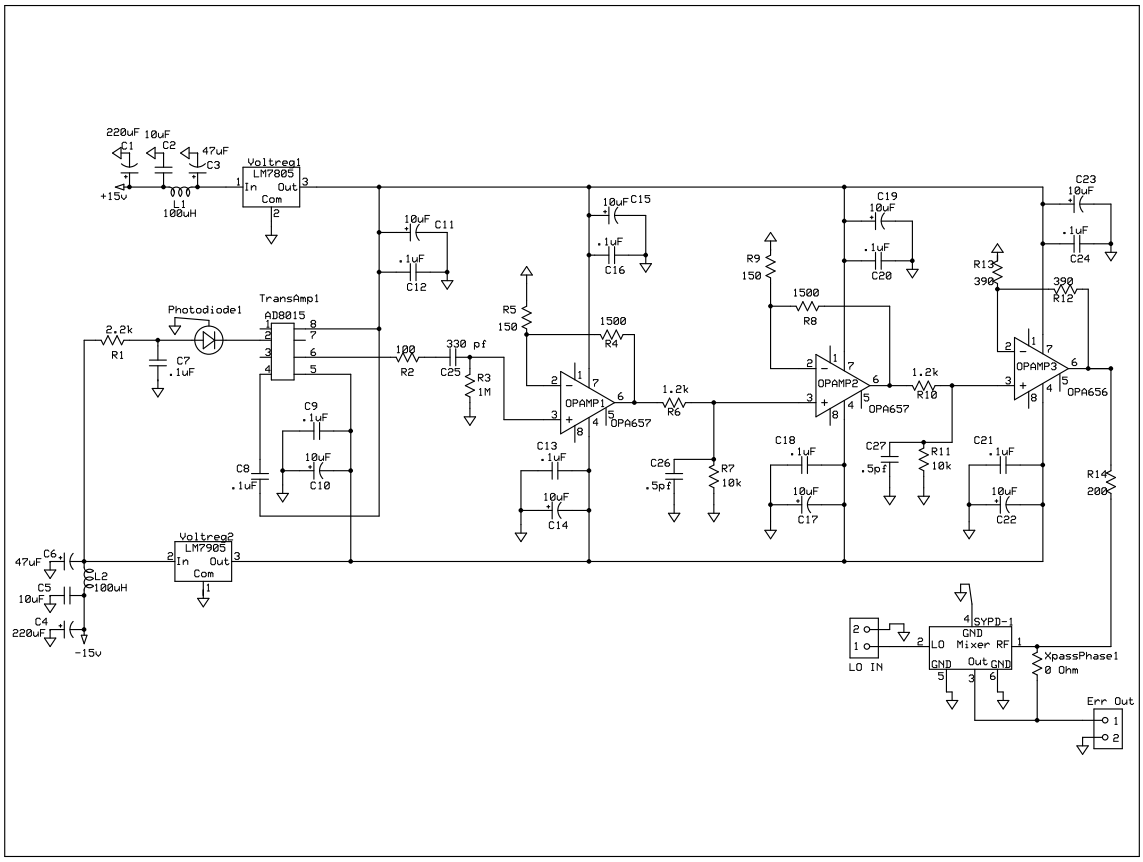

Figure 5.12 Schematic of the photodiode homodyne detector. A transimpedance amplifier converts the current to a voltage signal which is further amplified before it is sent to a mixer. The unit is designed for use as a simple photodiode, an amplified photodiode, or a homodyne detector which adds to its value as a lab standard.

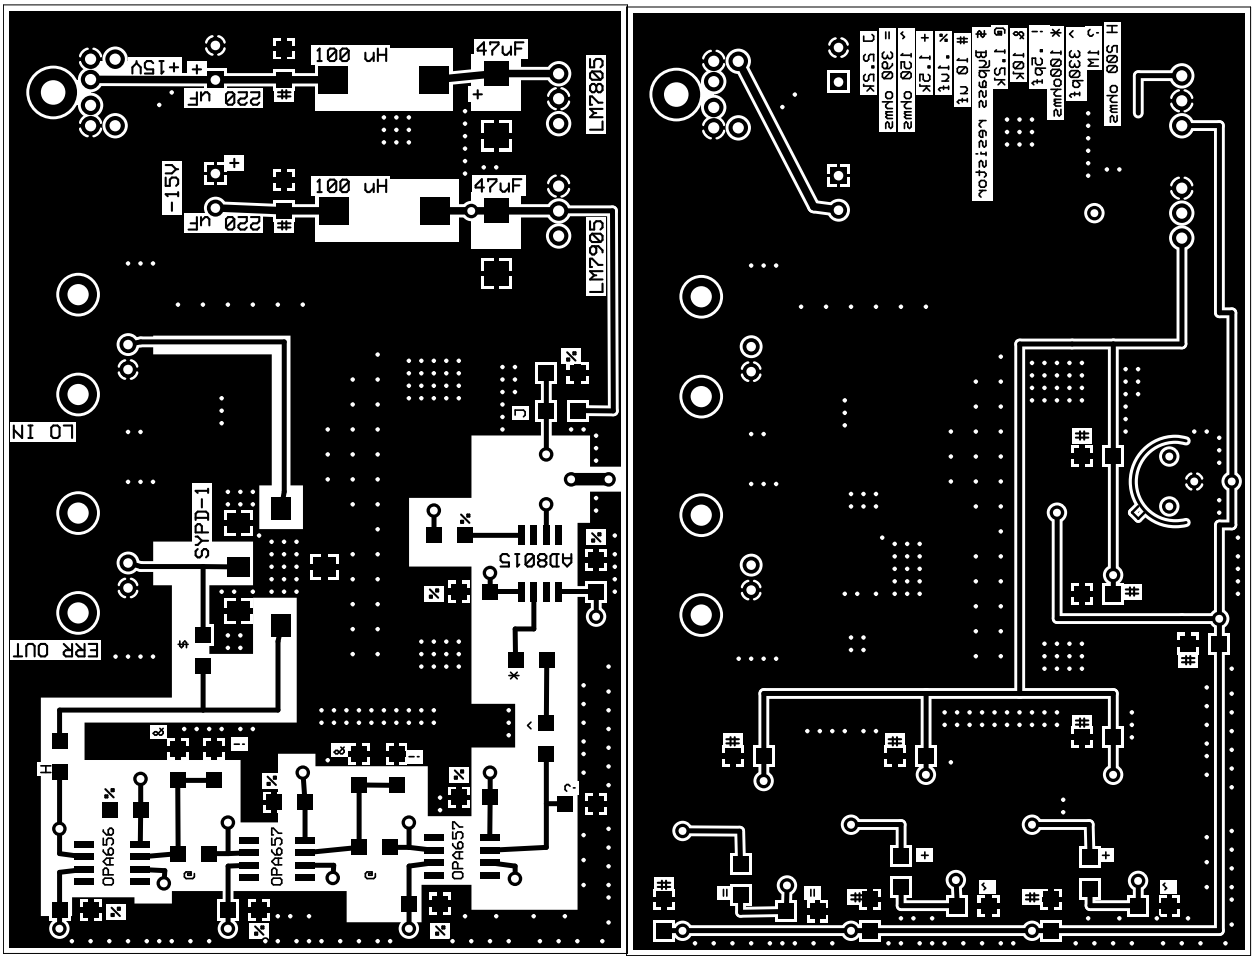

Figure 5.13 Homodyne detector PCB layout. The left side is the top of the board and the right hand side is the bottom of the board. The pad for the photo-diode is on the right hand side at about the middle of the board. The signal is immediately sent through a transimpedance amplifier before passing through up to three additional amplifiers on its way to an RF mixer. The local oscillator signal for the mixer is the modulation frequency signal driving the EOM. The mixer demodulates the light signal and then sends the resulting DC signal off the board.

5.4.1 Performance of the Homodyne Detector

The board layout is designed so that the distance between components and also between the mixer and the input/output connectors is never more than a few millimeters. This reduces stray capacitance and inductance, to which the bandwidth of the high-speed op-amps is very sensitive. Noise spectral density of the circuit is less than 10^{-4} V/$\sqrt{\text{Hz}}$ below 1 kHz and less than 10^{-5} V/$\sqrt{\text{Hz}}$ from 1 kHz out to 100 MHz. The total circuit response has a 3 dB point of about 50 MHz, which sets the limit for the frequency of the incoming modulated signal and the local oscillator as detailed in the senior thesis of Aaron Bennett [97].

5.5 Microprocessor

For the control of the diode laser current driver, temperature controllers, and scan-balancing lock-circuits we use a PIC18F4550 microcontroller from Microchip Technologies programmed to control all three of these devices simultaneously. The PIC18F4550 is a USB programmable microprocessor with 35 pins which can be used as inputs or outputs. It also includes built-in analog to digital converters (ADC) used to monitor the outputs of the devices. This chip is easy to update and can support many different functions, making it useful for remotely controlling the current driver, temperature controller, and the PID lock-circuit.

While the controller is USB programmable, that functionality only exists after a bootloader has been installed in the memory via a serial connection. To upload the bootloader, the chip is placed in the circuit shown in Figure 5.14. The circuit is connected to a standard communications port on a PC and a free program called WinPIC is used to upload the program. For more information on preparing and programming the PIC controller for use see [98] and Appendix F.

Once the bootloader is installed the PIC can be updated in situ via a USB connector on the board. The schematic for the microprocessing circuit is shown in Figure 5.15. The PIC uses 4

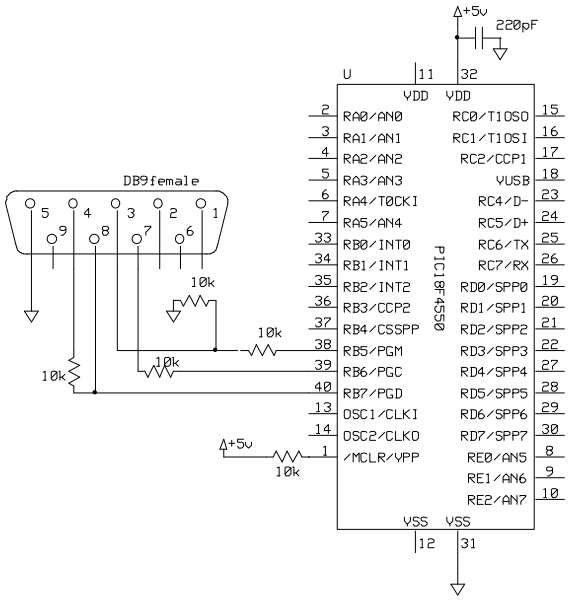

Figure 5.14 Bootloader circuit for the PIC18F4550 microcontroller.

ADC converters to read back analog signals: laser diode current, laser diode temperature, error signal voltage, and secondary signals for implementing additional locks on the laser diode. Eleven of the PIC's outputs are used for controlling an LCD screen, which can display the values read back by the ADCs as well as setpoints, functions, saved values, and other parameters. Five of the outputs are used for selecting and programming the DAC on one of 3 different satellite circuits (a current driver, PID controller, or a temperature controller). Two more of the outputs control digital switches on the PID controller. Eight pins are used as inputs to receive signals from the 5 buttons and one digital encoder used to manipulate the experiment through the PIC. These controls are also used to program setpoints and save important parameters into memory. The remaining pins are disabled by virtue of the chip's configuration.

In its primary function, the microcontroller reads in the signal from a digital encoder, and subsequently outputs a value to the DAC on a selected device. While the state of only one device can be changed at a time, the PIC can simultaneously monitor and display laser diode current,

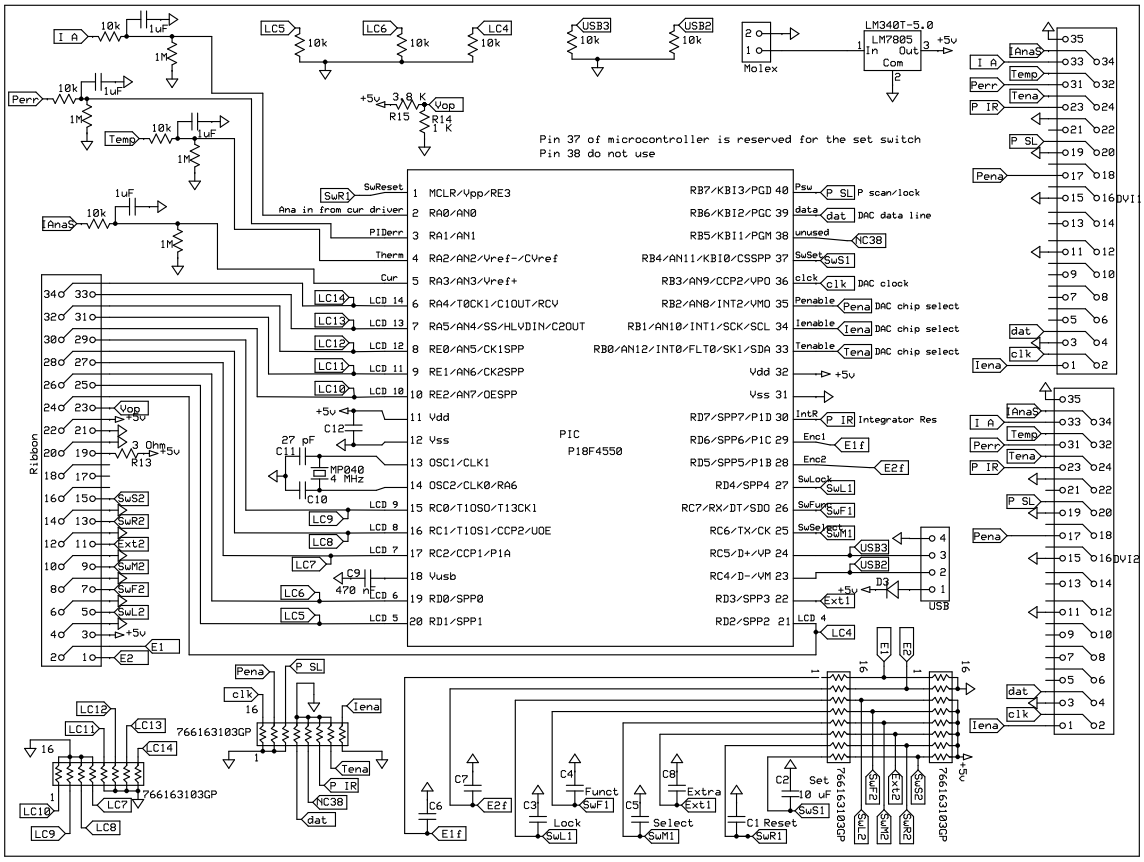

Figure 5.15 Schematic for the processing unit that controls the laser electronics. Once a "bootloader" is installed on a microprocessor, the USB connector on this board is used for uploading the program to run on the microprocessor. The multiple DVI-I connectors are used for putting signals out to and reading signals back from other electronics. The LCD is for displaying information from the devices being controlled.

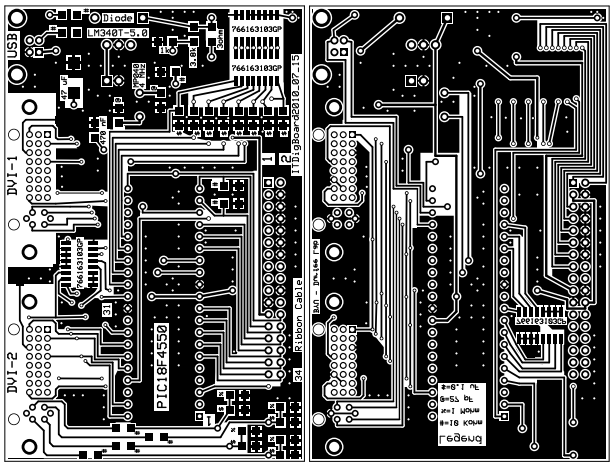

Figure 5.16 PCB of the processing unit for controlling laser drive electronics. The left-hand side is the top of the board and the right-hand side is the bottom of the board.

temperature, and the lock state of a laser system on the LCD screen. In order to create the most stable, versatile, and easy-to-use configuration possible much time has been spent redesigning the PCB for the controller (shown in Figure 5.16) as well as writing the PIC's software, which is given in Appendix F.

5.6 Photodiode Amplifier and Lock-in Circuit

The error signal for locking the 461 nm laser system to the strontium vapor cell is generated from the probe beam signal as discussed in section 4.6. The probe beam's signal is detected by the photo-diode amplifier circuit shown in Figure 5.17. An FDS010 small area photodiode is reverse biased with 15V and the signal is converted to a voltage by a transimpedance amplifier with a gain of $\sim 74,500\ \Omega$. Two AD8065 op-amps then amplify the signal by a total gain of $\sim +121$. The signal is then high-pass filtered before going to the lock-in circuit.

We originally tried to use the homodyne detector described in section 5.4 for this application, which would have incorporated the photo-diode amplifier with the circuitry that generates the error signal. Since the homodyne detector is designed for frequencies ≥ 1 MHz, the 461 nm

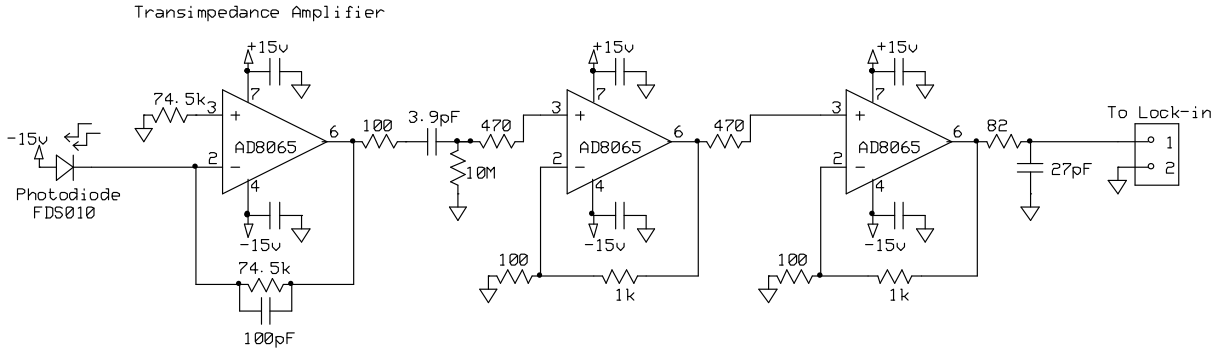

Figure 5.17 Schematic of the photo-detecting circuit. The signal from a reverse biased small area photodiode is sent to a transimpedance amplifier with a gain of $\sim 74,500\ \Omega$. A high-pass filter (~ 4 kHz) eliminates the DC bias that the modulated signal sits on top of before two more amplification stages of gain $+11$ each. A low-pass filter (~ 70 kHz) cleans up the signal before it is sent to the lock-in circuit of Figure 5.18.

pump beam was modulated by adding a 1 MHz signal to the AOM drive signal through a voltage controlled oscillator (VCO). However, no error signal was generated when the homodyne detector demodulated the probe signal. Unfortunately, the frequency response of the voltage controlled oscillator (VCO) at frequencies larger than hundreds of kHz is very poor. In order to produce an error signal for locking the 461 nm laser at lower frequencies a lock-in circuit was built.

The heart of the lock-in circuit is a 7 kHz square wave astable multivibrator shown in the upper left hand corner of Figure 5.18. The drive frequency of the AOM in section 4.6.1 is generated by providing the correct DC voltage to a VCO. In order to modulate the AOM about its main drive frequency the 7 kHz square wave is added to the VCO's DC signal in the "VCO Signal" section of the schematic. The addition of the square wave causes the AOM to switch the pump beam frequency across a 15 MHz span centered on the overall 180.6 MHz offset at a rate of 7 kHz.

The switching of the pump beam frequency also causes the probe beam to be modulated by 7 kHz in the vicinity of the Lamb dip. In the "Photodiode Signal Buffering" section of Figure 5.18 the probe beam signal from the photodiode detector is referenced to the lock-in circuit's ground. The signal is then sent to the "Digital Switch Signal Chopping Section" where a pair of switches

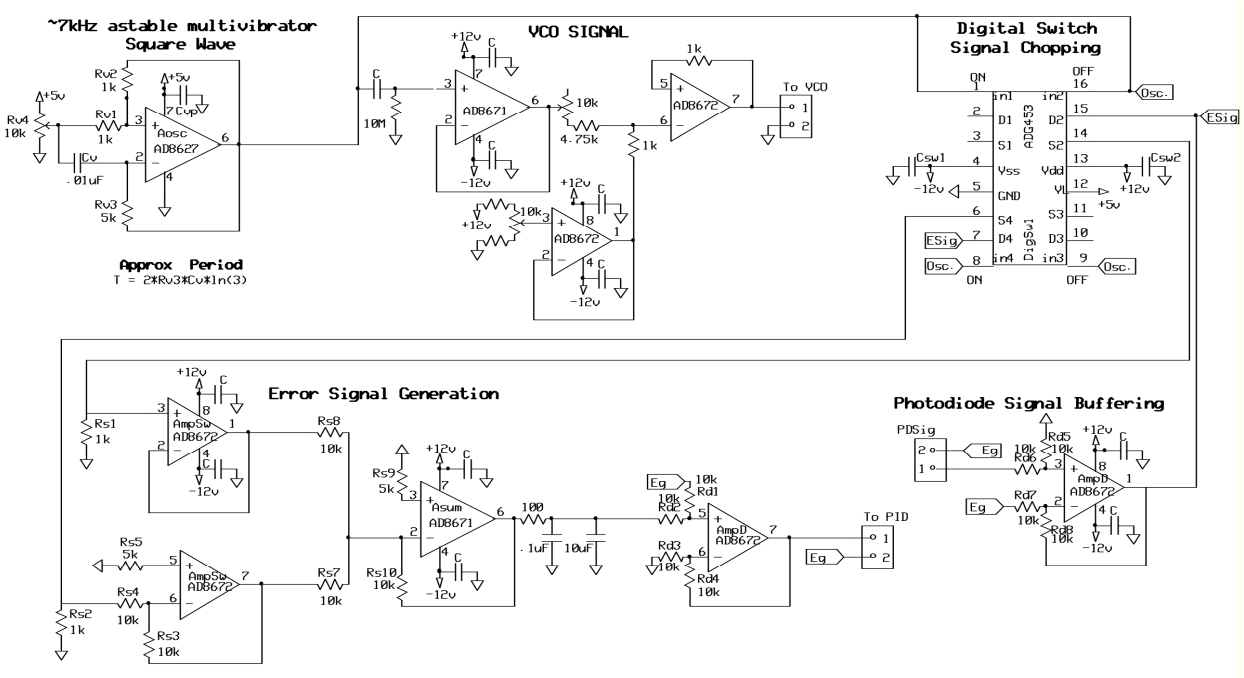

Figure 5.18 Schematic of the lock-in detection for the 461 nm laser lock to the vapor cell. An astable multivibrator generates a square wave used to modulate the pump beam frequency and mix down the resulting probe beam signal in section 4.6. The square wave is added to a DC offset for driving a VCO, which in turn drives the AOM in the optical setup. This signal switches the pump beam's frequency over a 15 MHz range at ~ 7 kHz centered on an overall frequency shift of 180.6 MHz. The signal from the probe beam is returned in the "Photodiode Signal Buffering" section and sent to a double pole, double throw switch which creates two signals: one corresponding to the high voltage of the square wave and the other to the low. One of the signals is then inverted and they are summed back together in the "Error Signal Generation Section". The summed signal is low-pass filtered to produce an error signal, which is sent to a PID controller.

divides the portions of the light signal into two signals corresponding to the duty cycle of the 7 kHz square wave. One of the signals is then inverted and they are summed back together in the "Error Signal Generation" section. A low pass filter is then used to eliminate the noise and reveal a DC error signal, which is sent to the PID controller for the lock.

Chapter 6

Sr Vapor Cell

Vapor cells are frequently used to perform spectroscopy on atomic and molecular species. The cell is usually a metal or glass chamber with a small sample of the species under vacuum. Usually the requirements on the background vapor pressure are minimal and in many cases the chamber is brought up to a specific pressure with a buffer gas.

The temperatures required to create an optically dense vapor of an element usually dominate the design of a particular vapor cell. Some species, such as alkalis, can produce significant vapors at room or moderate temperatures. Others, such as alkali-earths, often require serious temperature considerations especially if the spectroscopic lines of interest have narrow linewidths [1, 2]. This chapter details the construction and operational parameters of a custom built Sr vapor cell.

6.1 Requirements

6.1.1 Vapor density

In order to scatter a significant amount of an on resonance laser beam, an optically dense vapor of strontium must be created from the solid mass in vacuum. As a photon passes through a vapor of

length l the likelihood of it being absorbed is related the absorption cross section σ and the number density of atoms per unit volume n. In the low intensity limit, the relationship is given by Beer's Law

$$I = I_0 e^{-\sigma n l}, \tag{6.1}$$

where l is the length of the vapor, I_0 is the incident light intensity, and I is the final light intensity.

In a simple case l is the length of the strontium vapor that the laser passes through, and the typical absorption cross section is $\sigma = 3\lambda^2/2\pi$. For example, in order for a strontium vapor to absorb 50% of an on resonance 461 nm laser beam in a length of vapor $\sim$ 10 cm long, $\sim 6.8 \cdot 10^{13}$ atoms/m^3 are required.

6.1.2 Temperature

For alkali and alkali-earth elements the required density can usually be produced by heating the sample. The number density n is given by the Ideal Gas Law when the temperature of the sample and the vacuum chamber is in equilibrium:

$$n = \frac{P}{k_B T}, \tag{6.2}$$

where k_B is Boltzman's constant, P is the vapor pressure in Pascals, and the equilibrium temperature T is in Kelvin.

The vapor pressure P is also dependent on the equilibrium temperature. It can be calculated from the slope of the discontinuous boundary that separates the gas and solid phases of an element as a function of temperature. To calculate strontium's vapor pressure as a function of temperature, the vapor pressure data presented by Veeco [99] was fit with a Clausius-Clayperon equation. The vapor pressure of strontium up to $\sim 500^\circ$ C is well approximated by

$$\ln(P) = -\frac{19229}{T} + 18.834 \tag{6.3}$$

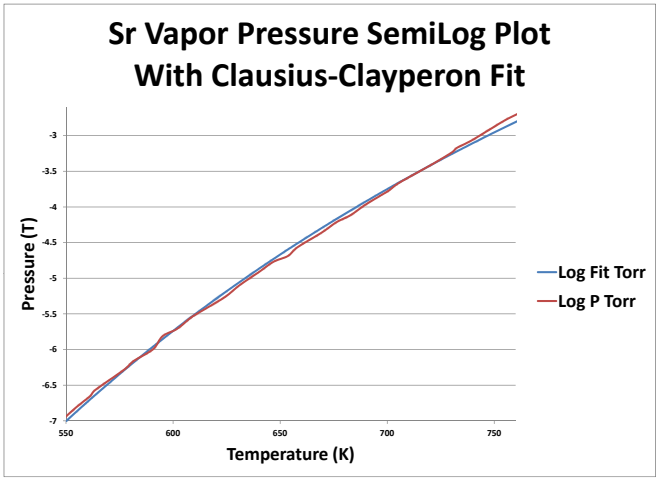

Figure 6.1 Clausius-Clayperon fit for Sr vapor pressure. The blue line is the fit and the red line is taken from vapor pressure data for strontium.

where the pressure is in Torr and the temperature is in Kelvin. The data and fit are shown in Figure 6.1.

Based on Eq. 6.3, to produce a strontium vapor that would scatter $\sim 50\%$ of an on resonance laser in a vacuum chamber 10 cm long the cell and strontium would have to be heated to $\sim 305°$ C. This seemingly allows for a simple vapor cell design since the temperatures required to produce very significant densities of Sr vapor are well below the maximum temperature for standard Conflat (CF) vacuum fittings. However, most strontium samples develop an oxide layer when transferred to the vacuum chamber. This coat prevents strontium from forming a vapor and must itself be vaporized in order for a strontium vapor to form.

Strontium oxide can be evaporated off the sample in vacuum at temperatures of $\sim 600°$ C or greater [100]. Therefore, while the vapor cell can easily be run long term with standard CF parts, its design must also allow temporary heating of the strontium well above the maximum temperature for standard vacuum parts (450° C).

6.1.3 Materials

As an alkali-earth, strontium has the tendency to react with many other materials. Fortunately, most standard vacuum materials do not present any issue. Typical vacuum chambers are constructed out of stainless steel and are sealed with copper gaskets. The stainless steel is nonreactive with strontium and strontium copper compounds do not pose a problem at temperatures less than 506° C [101].

Strontium does pose a problem for glass fittings. At the moderate temperatures of several hundred degrees Celsius strontium quickly chemisorbs into standard 7056 Kovar glass and other silicates. While strontium can also attack sapphire, this does not occur at temperatures below $\sim 600°$ C [100], making sapphire viewports a good option for optically accessing the vapor cell.

6.2 Sr Vapor Cell Design and Operation

A functional diagram of the vapor cell is shown in Figure 6.2. The cell is constructed from two CF tees that are joined at their center ports. One of the tees has been modified with a stainless steel cup welded opposite the central port. Standard CF sapphire viewports are mounted on the ends of the unmodified tee while a set of high temperature Swagelok valves and a CF blank cap the end flanges on the modified tee.

The cell is loaded by placing a Sr sample in the cup before evacuating the chamber through the valves. The cell is loaded with a 99% pure .5 gram sample of Sr sealed in an argon filled glass ampoule from Metallium [102]. With the cell ready to be evacuated, the ampoule is broken open under argon and the sample is cut into several small pieces to create surfaces with a minimum amount of oxidation. The pieces are then placed in the cup while flowing argon through the chamber.

The oxide layer is evaporated off of the strontium by heating the reservoir cup to $> 600°$ C. A

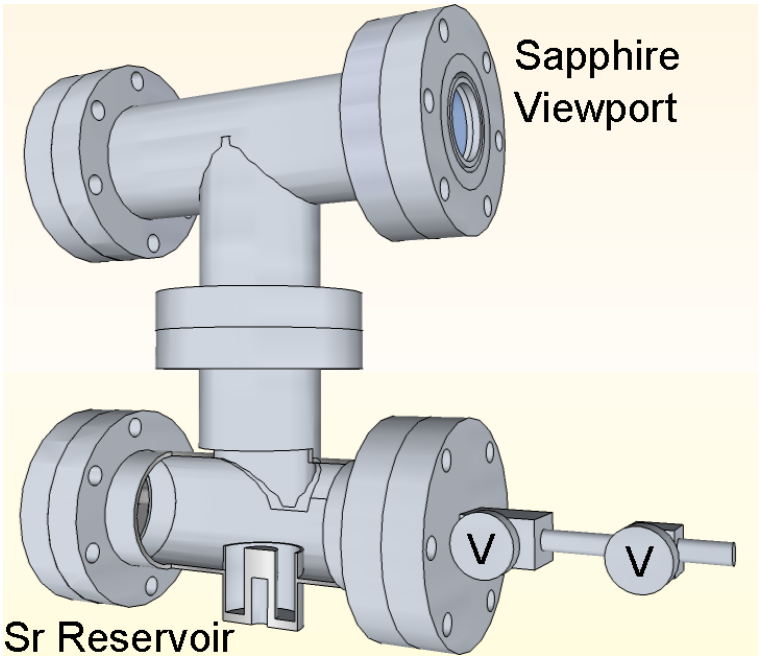

Figure 6.2 Drawing of the Sr vapor cell. The vapor cell is constructed from two CF tees. Part of the bottom tee has been cutaway in the drawing to show the cup that holds the strontium. The cup can be heated external to the vacuum with a cartridge heater. This allows a strontium sample to be heated well above the maximum temperature for the CF flanges. The top tee is capped with two sapphire viewports for optical access to the Sr vapor. Two valves are used to seal the cell. The valve directly on the cell is heated to prevent Sr from condensing on it. The second valve is not heated and serves to prevent atmosphere from corroding the first valve when hot. The cell can be run as a vapor cell or used to produce a broad thermal beam of strontium.

cutaway view of the cup is provided in Figure 6.2. It is 1.1 in. diameter, 1.5 in. long and in the center of the cup hollow there is a .5 in. diameter post. This post is hollowed out from the opposite side (external to the vacuum) to allow for the insertion of a cartridge heater.

The cup is heated with a .25 in. diameter 2.25 in. long, high-temperature cartridge heater from Omega. The heater has a maximum sheath temperature of 760° C and a maximum lead end temperature of 250° C. The heater extends out of the cup $\sim$ 1 in. and the lead end is clamped to a 1.5×4 in. aluminum heat sink in order to maintain the required temperature differential across the heater at its maximum temperature.

After loading a new strontium sample and evacuating the cell, bringing the cup temperature up to $> 600°$ C causes the cell pressure to rise to $\sim 10^{-4}$ Torr as the oxide layer evaporates. A turbo pump is used to continuously evacuate the cell while evaporating the oxide layer until the pressure returns to the level it was at before heating the cup.

After the cell has been evacuated and the oxide layer evaporated, the valves are closed. The valve mounted directly to the chamber is heated with the chamber to prevent condensation of strontium on the valve. Strontium vapor does not reach the second valve, which is not heated and is used to prevent atmosphere from corroding the heated valve.

The rest of the cell is heated with 2.75 in. mica band heaters placed around all of the CF flanges. The cup is insulated with a ceramic tape as its temperature can exceed the melting point of aluminum, while the rest of the modified tee is insulated with aluminum foil. The opposite tee is only insulated by aluminum foil around the sapphire viewports. The center of this tee is left open to the air to purposefully create a cold point that prevents condensation of the Sr on the viewports.

6.2.1 Results

When the viewports are brought to $\sim 450°$ C the center of the sapphire window is $\sim 360°$ C. An optically dense vapor resulting in $\sim 70\%$ absorption of on resonance laser light can be produced

by heating the CF flanges on the modified tee to $\sim 350°$ C and the cup to $\sim 600°$ C. At these settings the uninsulated length of tube comes to equilibrium at $\sim 250°$ C. By observation of the fluorescence, the density of atoms in this setting is very sensitive to the temperature of the cup and can be well controlled by the voltage driving the cartridge heater.

This indicates that the atoms are streaming from the back end of the chamber and condensing on the middle length of tube on the unmodified tee as expected. The thermal beam created in the cell results in a Doppler broadened absorption profile that is a little more than half of that expected for a vapor in equilibrium at $\sim 310°$ C. A seemingly almost uniform fluorescence is witnessed from window pane to window pane in the cell at equilibrium.

The cell is currently run as a thermal beam out of concern that strontium may still attack the viewports at these temperatures. Eventually, when all the strontium has been transferred out of the cup, the cell will necessarily be run as a true vapor cell. In this case, the coldest point in the cell will determine the density of the vapor.

Chapter 7

Vacuum Chamber Design and Contruction

The apparatus uses two main vacuum chambers: a chamber for cooling and trapping the atoms, and a chamber for ionizing and accelerating the ions through the interferometer and detecting the interference. The MOT and interferometer chambers are shown connected together in the functional diagram of Figure 7.1. In reality, they are separated by a small section of vacuum fittings that house a low-conductance tube to provide a large pressure differential between the two chambers.

Neutral strontium atoms are heated in the MOT chamber to create a thermal beam that is laser cooled, trapped, and formed into an LVIS as described in Chapter 3. This beam of atoms travels through the differential pumping section and into the interferometer chamber. Upon entering this chamber the atoms pass between high voltage electrodes where there are photo-ionized and accelerated. Windows along the length of the interferometer chamber allow optical access to the ions from both sides.

In order to achieve the best vacuum levels each chamber was prepared, cleaned, and baked as described in Appendix E. In practice, careful measures are taken to ensure the necessary background pressure is reached in a particular vacuum chamber. This is because collisions with the background gasses in the vacuum can significantly affect the trapping of atoms to create the LVIS

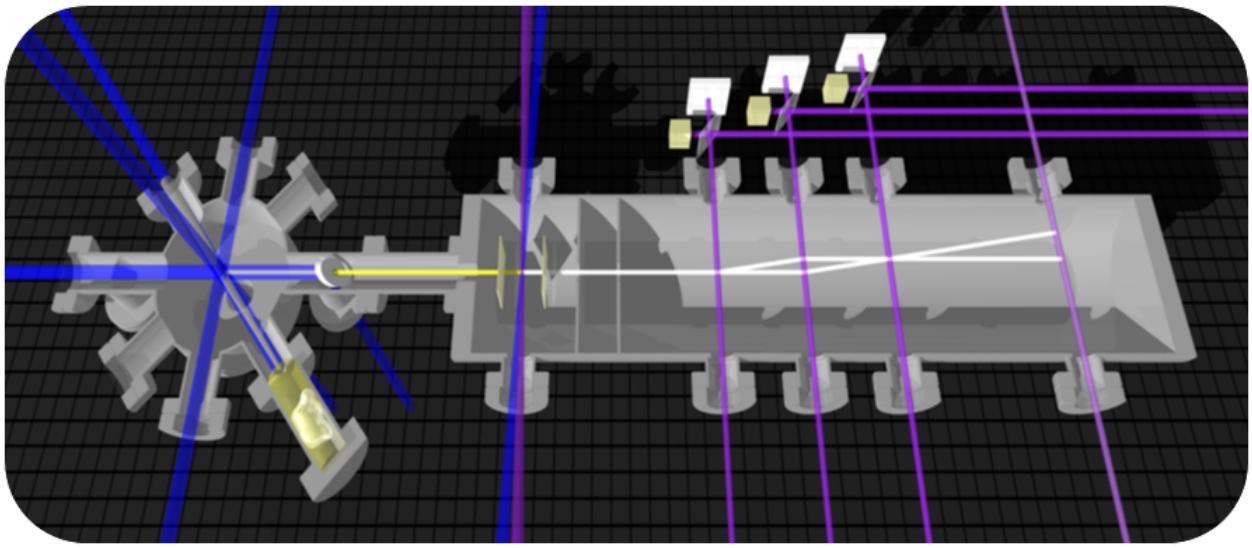

Figure 7.1 A simple diagram of the vacuum chambers. The chambers are not to scale and not all of the ports on each chamber are represented. In the experiment strontium atoms are laser cooled to form an LVIS in the MOT chamber on the left hand side. The atoms are then ionized, accelerated, and pass through the interferometer in the right hand chamber. The chambers are separated by a differential pressure section since the vacuum requirements for each chamber are different.

in the MOT chamber as well as the contrast ratio of the interferometer in the interferometer chamber.

7.1 Background Pressure and Mean Free Path

The likelihood of a strontium atom or ion colliding with background atoms and molecules in the vacuum can be determined by the pressure in that chamber and the assumption that residual gasses are in thermal equilibrium with the chamber. This allows us to make a few simple arguments in regards to determining the ultimate pressure required in each chamber for the apparatus to be functional.

In the case of the MOT chamber, the time that it takes for a background particle to collide with an atom loaded into the slow atom beam should be long compared to the time it takes for the atom to be trapped and accelerated through the differential pumping section. For the case of the interferometer chamber, this time scale should be long compared to the time it would take the slowest ions to cross the length of the chamber.

The mean free path is given from kinetic theory by

$$\lambda_{mp} = \frac{k_B T}{\sqrt{2}\pi d^2 P} \tag{7.1}$$

and is defined as the average distance a particle travels before colliding with another particle in a gas at pressure P, made up of particles of diameter d, in thermal equilibrium at temperature T. The factor of $\sqrt{2}$ corresponds to the assumption that the particle we are considering as well as all of the particles that can collide with it are moving. For the cases we will consider in this chapter, one set of particles will be approximately motionless, and this factor will be dropped. The temperature also gives the characteristic thermal velocity of the particles where the rms velocity is $v_{therm} = \sqrt{3k_B T/M}$, where M is the mass of the particle.

The time it takes a background particle to collide with the MOT or atom beam needs to be

much longer than the time it takes to cool an atom, load it into the LVIS beam, and travel to the interferometer chamber. From Chapter 3, the time it takes to cool the atom is ~ 7.4 μs, the time it takes to travel across the MOT at the Doppler temperature is ~ 2.7 ms, and the time it takes to travel to the interferometer chamber at ~ 100 m/s is roughly 5 ms. For the time between collisions to be just 160 ms vacuum level of $\sim 10^{-7}$ Torr is required.

For the atom to cross the interferometer chamber at the slowest useful velocity (~ 3.8 km/s), the time would be again roughly 5 ms. However, once the atoms are ionized their sensitivity to background gasses increases as discussed in section 8.3.2. Given the analysis in that section pressures in the interferometer chamber need to be on the order of $\sim 10^{-10}$ Torr.

7.2 Differential Vacuum Systems

It is possible for the MOT and the interferometer chambers to have an open connection between them and yet maintain vastly different pressures. This is achieved by having the atoms cross from one chamber to the other via a region of extremely low conductance as seen in Figure 7.2, which is also labeled with rates and pressures corresponding to different parts of the system. To find the conductance that would produce the desired pressure differential we can set up a simple equation. In steady state the rate of mass being pumped out of the chamber is equal to the rate of mass entering the chamber

$$QP_{Int} = (P_{MOT} - P_{Int})C + P_{Atm}R_{Atm} \tag{7.2}$$

$$C = \frac{P_{Int}(Q - R_{Atm})}{\Delta P}, \tag{7.3}$$

where Q is the rate gas is pumped out of the interferometer chamber in L/s, C is the conductance of the tube separating the chambers, and R_{Atm} is the rate at which atmosphere enters the interferometer chamber by permeation through elastomer seals (the main gas load in the chamber as discussed

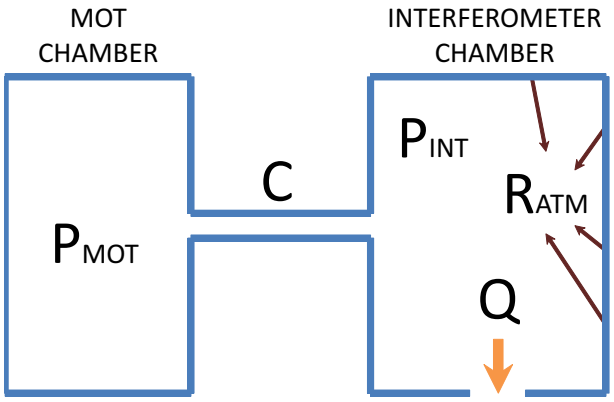

Figure 7.2 Diagram of the pressure system of the vacuum chambers. In order to provide a pressure differential of three orders of magnitude in our system the upper limit on the low conductance region is $\sim .14$ L/s.

in section 7.6). Given that the rate of gas entering the chamber through the elastomer seals is $R_{Atm} \sim 1.16 \times 10^{-3}$ L/s (see section 7.6), and the pump rate is 140 L/s, a region of minimum conductance $C = .14$ L/s is required to keep the interferometer chamber at 10^{-10} Torr with the MOT chamber at 10^{-7} Torr.

In the molecular flow regime, the conductance in L/s through a tube that is much longer than its diameter is given by

$$C = \frac{12D^3}{L}, \tag{7.4}$$

where D is the diameter of tube and L is the length of tube given in centimeters [103]. In order to provide the necessary conductance a .25 in. diameter tube would need to be ~ 8.5 in. long.

7.3 Differential Pumping Chamber

The differential pumping section connecting the two large chamber is shown in Figure 7.3. This section consists principally of a 2.75 in. CF cross that is mounted to a gate valve leading to the MOT chamber. The horizontal ports on the cross are windowed in the same manner as described below for the the MOT and interferometer chamber to provide optical access to the atoms between

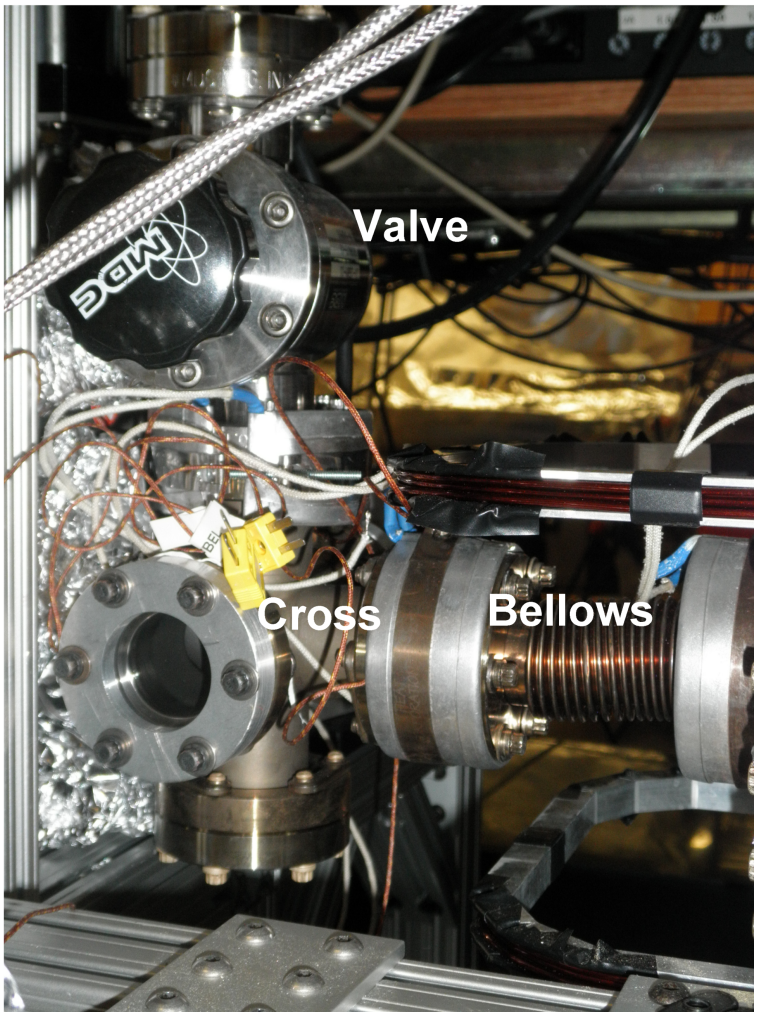

Figure 7.3 Picture of the differential pumping section. The cross on the left hand side is mounted to a pneumatic gate valve on the MOT chamber. The low conductance tube is mounted to a flange on the interferometer chamber and extends through the bellows seen on the right hand side. A valve has been mounted on the cross so that this section can be pumped out independent of the two main chambers.

the two chambers. A bellows valve is mounted on the top port of the cross for evacuating this section independent of the other chambers, and the bottom port is sealed with a CF blank.

Opposite the gate valve port on the cross is a CF bellows. The bellows is mounted to a double sided flange on the interferometer chamber. In the center of the double sided flange a .25 in. diameter, 8.5 in. long stainless steel tube is welded place. This tube provides the only free pathway

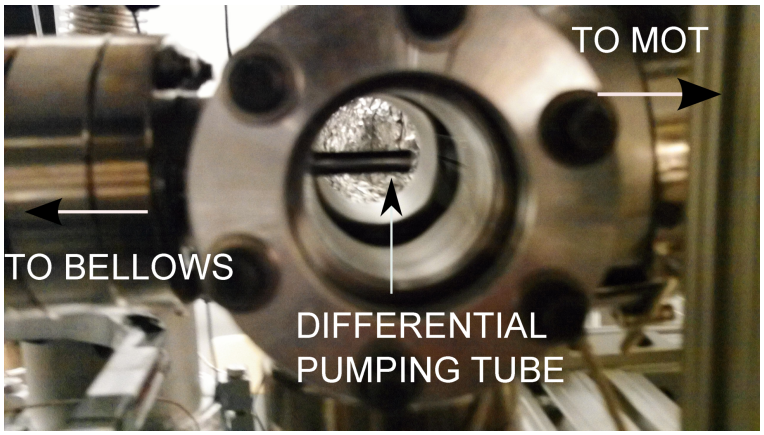

Figure 7.4 Picture showing the differential pumping tube extending into the cross.

between the interferometer chamber and the MOT chamber. The tube end can be seen in the center of the cross in Figure 7.4. The tube extends from the cross center through the bellows and into the interferometer chamber, stopping several inches from the electrodes.

7.4 MOT chamber

The MOT chamber is somewhat pancake shaped with four pairs of opposite facing CF ports extending out of its horizontal axis as can be seen by the cutaway view in Figure 7.1. A picture of the MOT chamber in situ is shown Figure 7.5. There are four CF ports that extend out of the chamber vertically with one on the top of the chamber and three on the bottom.

The elliptical mirror described in section 3.4.1 is mounted into the chamber through one of the bottom facing ports. One downward port sits opposite the top port and allows optical access to the atoms in the vertical direction. The last downward port branches off to an ion pump and valve.

Most of the horizontal ports are windowed. One is mounted with a pneumatic gate valve (leading to the differential pumping section), another is fitted with an Sr oven, and a third is mounted with a tee for branching off to an ion gauge while allowing the slowing laser beam to be directed at the oven.

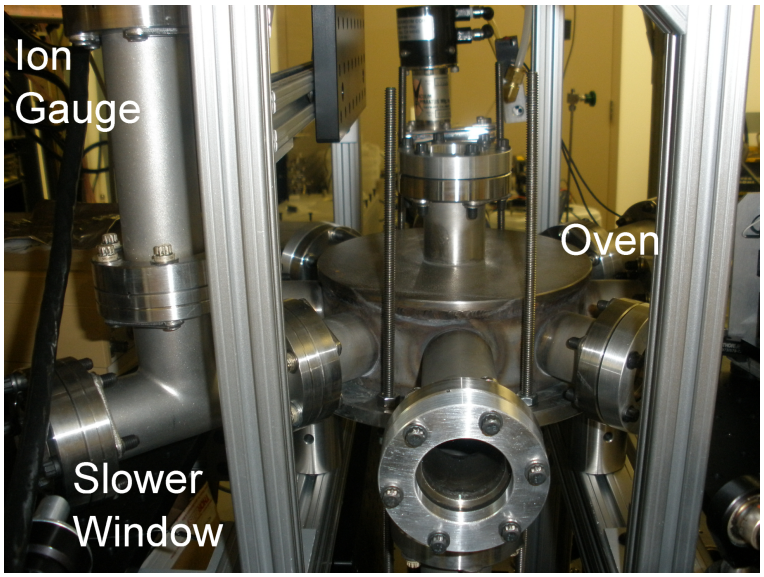

Figure 7.5 Picture of the MOT chamber. The slower window sits opposite the oven, and the other ports in the horizontal plane provide optical access to the chamber. Breadboards are mounted above and below the chamber for optics directing the beam vertically.

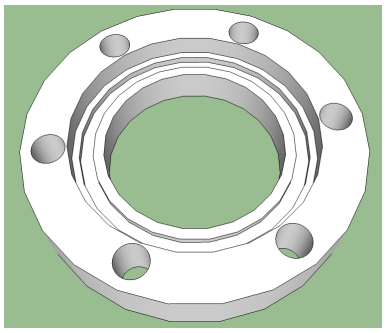

Figure 7.6 Drawing of the window clamps used on the apparatus.

Eight of the twelve ports on the chamber are windowed with anti-reflection coated optically flat glass. The vacuum seal is created at each window with Viton gaskets. The window is held in place by a custom made ring shown in Figure 7.6. These rings were machined out of aluminum. The clamp has a groove for a rubber gasket to soften the contact between it and the glass.

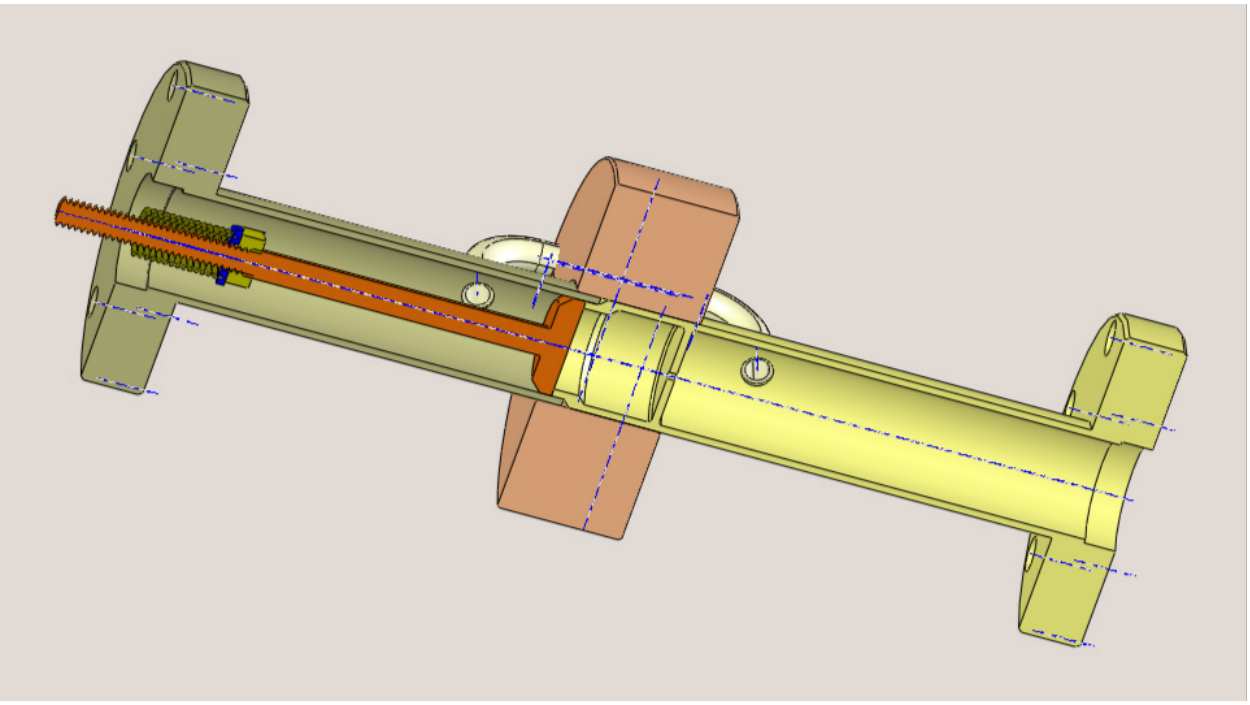

Figure 7.7 Cutaway drawing of the Sr thermal source.

7.4.1 Sr Oven and Beam Collimation

A cutaway view of the Sr oven is given in Figure 7.7. The assembly consists simply of a long tube with a 2 mm aperture disk-like nozzle welded in the middle of it. A spring loaded rod and plunger then extend from the end flange up to 1 in. away from the nozzle. A short tube .25 in. in diameter connects the volume on both sides of the nozzle/plunger. The strontium sample is placed in the tube and pressed up against the nozzle with the rod and plunger before the back flange is sealed with a CF blank and the assembly is evacuated through the tube connecting to two parts of the tube.

The oven is heated by placing a 3 in. diameter 1 in. thick brass slug around the nozzle/plunger area where the strontium is housed. The purpose of the slug is to allow the use of a standard band heater despite the presence of the tube, as well as to evenly distribute heat to the area housing the Sr.

This assembly is connected to a port on the main MOT chamber via a double-sided flange fitted with a second matching 2 mm aperture skimmer which sits $\sim$ 12 cm ahead of the oven's nozzle. The relative position of the apertures limits the full angle of divergence of the exiting atoms to $\sim 2^\circ$ and is enough to keep the atoms from being deposited on all the windows (except for the slower window) by preventing a direct line of sight to the source.

The oven is placed several inches away from any CF flange as it must be heated well beyond their maximum temperature to initially evaporate an oxide layer off of the strontium. Once that is done enough strontium vapor is produced at $\sim 270^\circ$ C to create a visibly bright MOT.

The port opposite the oven is the only one to support a standard CF viewport which is heated in order to prevent condensation of the strontium. This viewport provides for slowing light to directly access atoms leaving the oven in order to boost the capture rate of the trap.

7.4.2 Elliptical Mirror Mounting

The elliptical mirror used to form the LVIS is mounted an a custom made piece through one of the bottom ports on the MOT chamber. The mount is shown in Figure 7.8. To mount the mirror a 1 in. diameter stainless steel tube is cut at a 45° angle and welded to a double sided flange. The tube's length is precisely cut so that the mirror is centered horizontally to the main body of the MOT chamber. A hole bigger than the cooling laser beam is drilled in the tube to allow the cooling light and atoms clear passage.

The mirror is secured to the tube with the use of TorrSeal, a low vapor pressure epoxy. The double sided flange is faced on the opposite side with an anti-reflective coated window and ring and the assembly is bolted to the chamber with Viton gaskets. Light is then sent through the chamber, reflected downward through the tube by the mirror, retro-reflected outside the cell and returned with a shadow in the beam cast by the hole in the mirror.

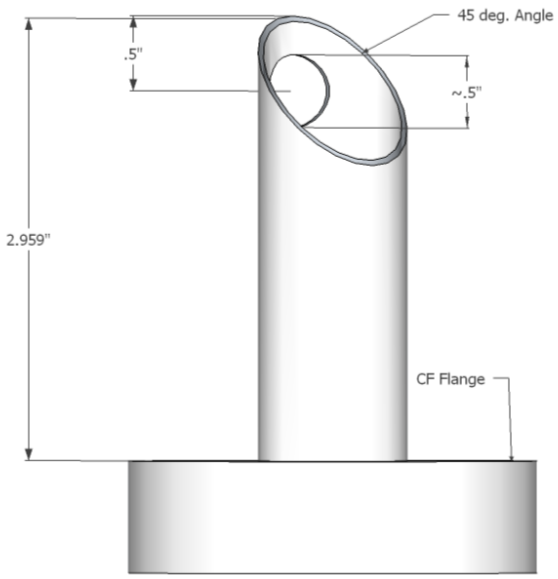

Figure 7.8 Drawing of the mount used for the elliptical mirror. The mirror was glued to the angled part of the tube with a vacuum epoxy. The hole in the tube aligns with the hole in the mirror to allow the atoms to pass through. Incident and retro-reflected beams pass through the hollow of the tube. The tube is welded to a double-sided CF flange.

7.4.3 MOT Chamber Pressure

The MOT chamber was cleaned and baked out and maintains a constant pressure of $\sim 10^{-7}$ Torr at room temperature with the help of a 30 L/s ion pump. When running the oven, pressure in the chamber rises quickly to $\sim 10^{-6}$ Torr before settling to $\sim 5 \times 10^{-7}$ Torr. Residual gasses in the MOT chamber most likely come from two major sources. When the nozzle and oven are cold, air diffusion through the Viton gasket seals on the ports of the chamber provides the bulk of the gas load. When the strontium is hot the increase in pressure is probably from increased outgassing in the oven assembly.

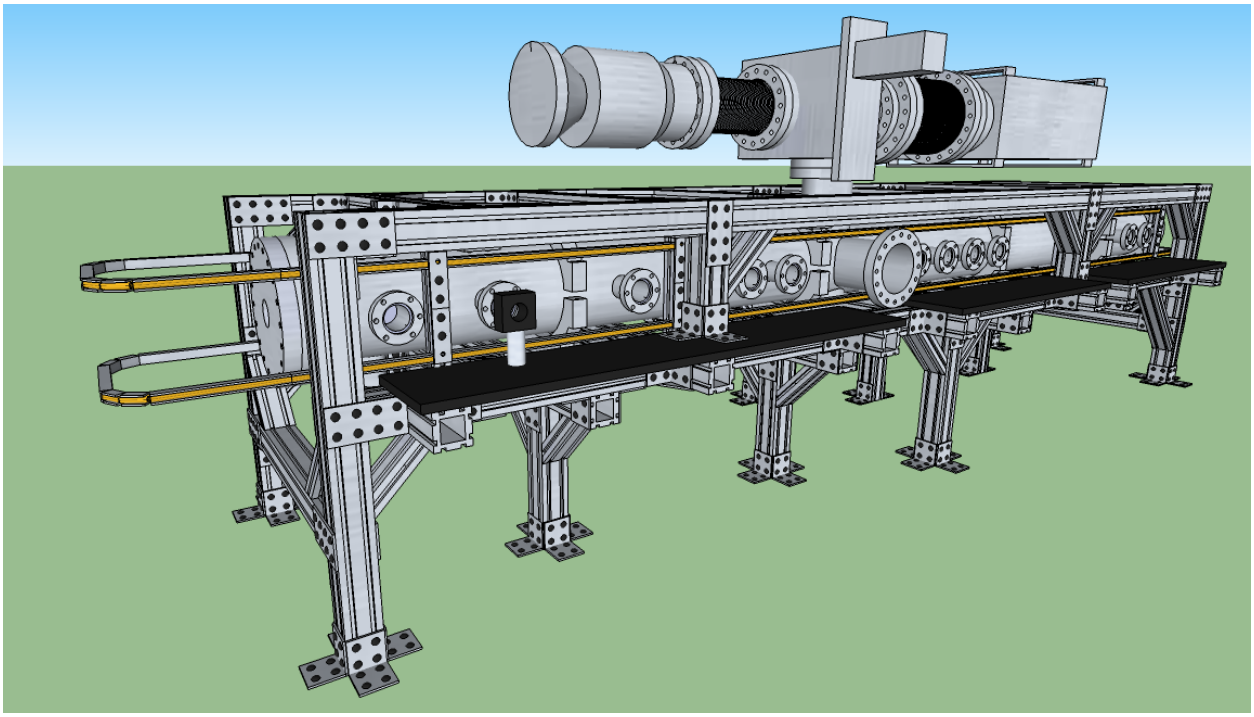

Figure 7.9 Illustration of the interferometer chamber and its support structure. The two wire loops encircling the chamber create a quantization axis for the experiment. An ion and a turbo pump are mounted to the top of the chamber for quickly reaching and maintaining a high vacuum environment.

7.5 Interferometer chamber

A diagram of the interferometer chamber is seen in Figure 7.9. The chamber was constructed out of an 8 ft. long 6 in. diameter 304 stainless steel tube. The chamber has 8 in. CF flanges welded on both ends, 11 pairs of 2.75 in. CF ports on its horizontal axis, 2 pairs of 2.75 in. CF ports on its vertical axis, and a set of three 6 in. CF ports at about the midpoint of the chamber. During the welding process, stresses on the chamber caused a slight bend about the 6 in. ports in the horizontal axis.

The ends are capped with 8 in. to 2.75 in. CF adapters, one of which was machined to mount a set of electrodes inside the chamber. A hole was also drilled through this modified adapter and a safe high voltage 20 kV (SHV-20) electrical feedthrough was welded in place. This adapter is

fitted with the flange holding the low conductance tube while the other is fitted with a 2.75 in. CF viewport.

All 22 of the 2.75 in. CF ports on the horizontal axis are fitted in the same manner as the Viton sealed MOT windows with custom made rings and optically flat anti-reflective coated windows. The downward vertical port near the modified 8 in. CF adapter and the upward port near the back of the chamber are both fitted with tees. The remaining vertical ports as well as the tees are window capped in the same manner. A nude ion gauge and a residual gas analyzer RGA (Extorr XT100) are mounted to CF nipples on the center ports of the tees.

Six inch CF tees are mounted to the vertically aligned 6 in. ports in the center of the chamber and the sole horizontal 6 in. port is for now capped off with a CF blank. The upper tee branches off to a 140 L/s Varian style ion pump and a pneumatic gate valve leading to a turbo pump. The lower tee houses a titanium sublimation pump, which is placed in such a way as to have direct line of sight to the upper tee in order to maximum the surface area coated with titanium and thereby the pump rate.

7.5.1 High-Voltage Electrodes

A pair of electrodes and an SHV-20 feedthrough are mounted to the 8 in. zero-length adapter flange on the chamber for accelerating strontium ions as they are created. A picture of the electrodes is shown Figure 7.10. The electrodes are machined out of oxygen free high-conductivity (OFHC) copper and consist of a cup shaped cathode, and a disk shaped anode. The cup is 4 inches in diameter with a wall thickness of .5 in. on the cylinder, a wall thickness of .25 in. on the base, and a depth of .75 inches. Eight through holes for $1/4-20$ screws are placed along the cylindrical axis of the cup through the walls for mounting to ceramic standoffs. Four standoffs are used to mount the anode cup to the flange, and four are used for mounting the cathode plate. Quarter inch wide slits have also been cut in the walls of the cup at orthogonal axis to allow for immediate optical

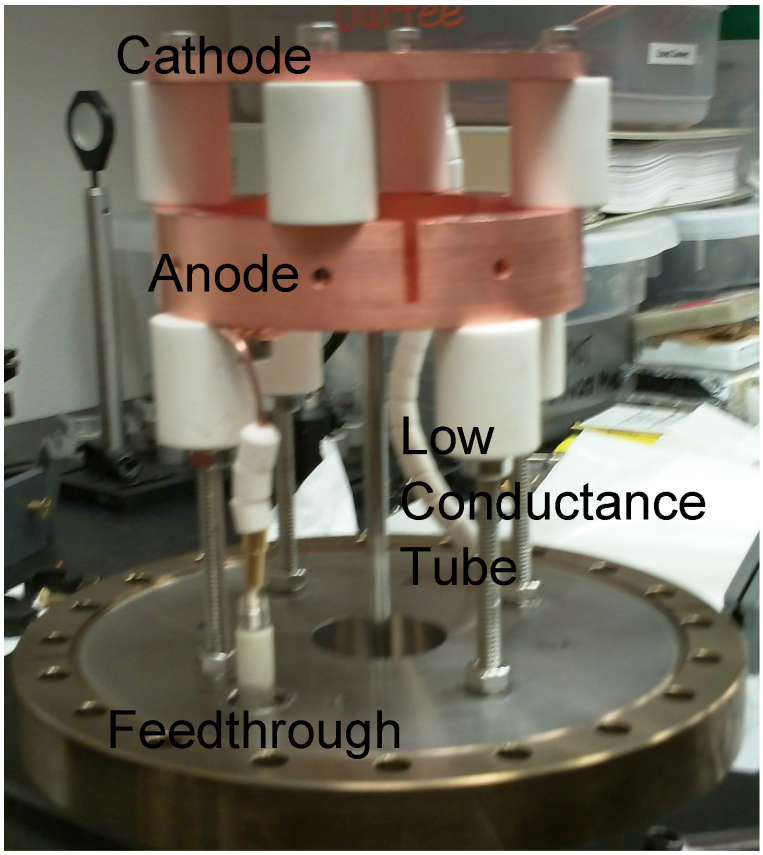

Figure 7.10 Picture of the high voltage electrodes. The low conductance tube is shown extending to the electrodes because it was used to align them with the flange. It is not positioned there in the experiment. A SHV-20 electric feedthrough mounted in the flange can provide up to 20 kV across the anode and cathode. The electrodes are centered on the low conductance tube, which is also shown. The ceramic spacers are 96% alumina and the ceramic fish-spine beads are 95% alumina. The electrodes are machined out of OFHC copper.

access to the atoms from two directions as the atoms enter the anode. The cathode is a .25 in. thick 4 in. diameter plate with four $1/4-20$ through holes for mounting to the anode. The aperture in the center of each electrode is .25 inches in diameter and the wall of the hole is slowly tapered through the thickness of each piece so that the metal comes to a thin point around the aperture.

Components of the electrode assembly are spaced by 1 in. diameter 1.25 in. long $1/4-20$ threaded ceramic standoffs. The standoffs are composed of 96% alumina, which has excellent

vacuum properties given that it is gas tight and does not absorb water. The remaining 4% of the ceramic is made up of magnesium oxide and silicon dioxide: the magnesia is used to inhibit the grain growth of the alumina crystals, and the silica is used because it forms a crystalline compound with the alumina [104]. The alumina is a good electrical insulator with a dielectric strength of 9.8 kV/mm and a volume resistivity of $> 10^{14}$ Ω·cm, which also makes it ideal for rigidly mounting the electrodes to the chamber and each other [105]. All of the electrodes and standoffs are bolted together using 316 stainless steel $1/4-20$ slotted screws to prevent virtual leaks. The electrodes are electrically connected to the chamber and the SHV-20 feedthrough with OFHC copper wire threaded through 95% alumina fish-spine beads.

The electric field lines produced by the electrodes are shown in Figure 7.11. The electrodes produce a field that accelerates the ions while introducing a half-angle divergence of only ~ 100 μrad in the ion beam.

High-Voltage Safety

Since the electrodes are designed to be charged up to 20 kV, we were concerned about risk of electric shock either from a mistake in the assembly of the electrodes or the possibility of a person mishandling lab equipment. In order to reduce the risk of personal injury from high voltage on the experiment we built a simple protection circuit for the high voltage power supply to pass through before going to any equipment.

A schematic of the circuit is shown in Figure 7.12. The circuit is housed in a grounded waterproof conducting enclosure fitted with SHV-20 bulkheads. Upon entering the box the signal sees a 1.1 GΩ power resistor to ground through an ammeter. The signal then passes through a low pass filter composed of another 1.1 GΩ power resistor and a 40 kV, 2 nF capacitor before passing through a third 1.1 GΩ power resistor and out of the box. Despite the high voltage, if the output were shorted a peak current of ~ 18 μA could be produced by the capacitor. A maximum

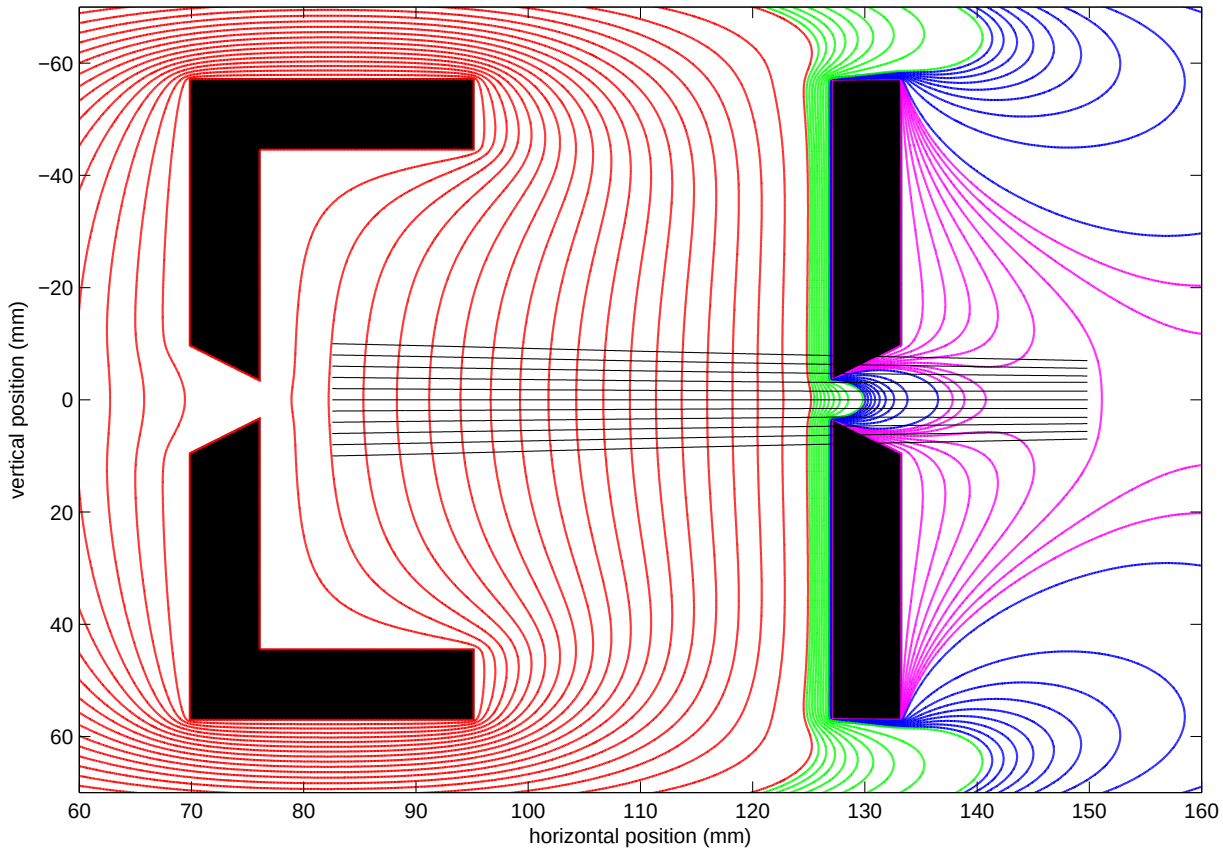

Figure 7.11 Electric field generated by the electrodes and the resulting ion paths. The thick black shapes represent the physical location of the copper. The electrodes have cylindrical symmetry and the figure shows only the cross section. The red, green, blue, and magenta lines are equipotential lines. The red lines span a potential range of 500 to 10,000 V and are spaced by 500 V. The green lines span a potential range of 50 to 500 V and are spaced by 50 V. The blue lines span a potential range of 5 to 50 V and are spaced by 5 V. The magenta lines span a potential range of 0 to 5 V and are spaced by 0.5 V. The thin black lines show the predicted trajectories of ions produced at different points off axis and then accelerated in this field. If a 1mm diameter beam is ionized, after leaving the acceleration region the outermost ions will be traveling at an angle of 91.28 μrad away from the axis.

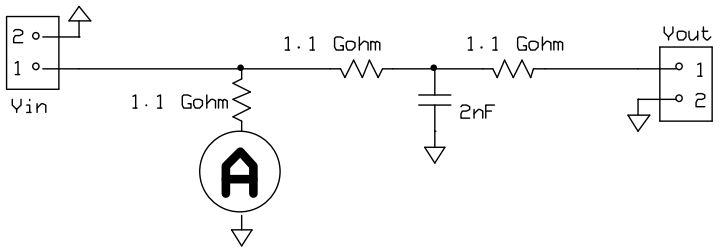

Figure 7.12 Schematic of the HV protection circuit. The voltage is measured by an ammeter through 1.1 GΩ. Another 1.1 GΩresistor and a 2 nF capacitor form a low-pass filter for the signal and a third 1.1 GΩ resistor prevents the capacitor from discharging quickly if the output is shorted.

continuous current of only $\sim 9\ \mu$A would be allowed to flow delivering ~ 180 mW.

Precautions have also been taken to prevent arcing inside the box and to help protect anyone who happens to open the box while the power supply is on. Each component in the circuit is physically separated from other circuit components and the enclosure by a large enough distance to prevent dielectric breakdown of the air in the box. Additionally, circuit components are housed in Teflon tubing and leads and connections are coated in Dow Corning 732 RTV sealant. The tubing has a dielectric strength of $25-50$ kV/mm and the sealant has a dielectric strength of 21.6 kV/mm [106, 107].

7.5.2 Interferometer Chamber Pressure

With the normal pumping action of the 140 L/s ion pump the interferometer chamber has an equilibrium pressure of $\sim 10^{-7}$ Torr, limited by the gas load described in the next section. This is not sufficient for ultimately running the ion interferometer. To increase the pump rate on the chamber the intermittent use of a titanium sublimation pump achieves pressures $\sim 10^{-10}$ Torr for up to a few hours after the pump is run.

7.6 Gas Load from Viton Seals

The vacuum pressure in the both the MOT and the interferometer chamber is limited by permeation through the elastomer sealed flanges. This gas load is particularly large on the interferometer chamber where 26, 2.75 in. CF ports are Viton sealed. A preferred option would have been to use a Kasevich seal [108]. However, the complications of implementing such a system on so many flanges has made it impractical so far. It is likely that in the future such seals might replace the current elastomer ones.

The permeation rate of atmosphere through a standard 3 mm thick Viton gasket is given as $2.5 \cdot 10^{-8}$ Torr·L/s per linear inch [109]. For a gasket of 1.68 inch inner diameter this translates into a gas load of $1.3 \cdot 10^{-7}$ Torr·L/s for every 2.75 in. flange that is elastomer sealed. An average sized chamber with no active pumping and just one elastomer seal can see the pressure raise from $\sim 10^{-7}$ Torr to $\sim 10^{-3}$ Torr over the course of several hours. The permeation rate is also temperature sensitive as indicated by the RGA on the interferometer chamber during bake-out.

In an attempt to lower the permeation rate of the elastomer seals the exposed surface of all the Viton gaskets on the chamber were coated with VacSeal. VacSeal is rated for sealing vacuum leaks of 2 μm·L/s on vacuum systems down to 10^{-12} Torr. However, the RGA showed no improvement in the gas load of atmosphere after the sealant had cured. No data on permeation through large surface areas of VacSeal is available.

Chapter 8

The Ion Interferometer

8.1 Charged Particle Interferometry

With the exception of only one, all existing and previous charge particle interferometers have been based off of electrons. The one ion interferometer was not published [33]. Without exception all previous charged particle interferometers have been based off of a physical grating or bi-prism scheme using a thermal source of particles.

For many applications, electron interferometers suffer from several disadvantages when compared to an ion interferometer in which the ions are manipulated with optical forces. First, electron sources are thermal and have less coherence than laser-cooled atom beams. As such, contrast is often an issue in electron interferometers.

Another difference is that the phase of the interferometer cannot easily be encoded in the internal state of an electron as it can with an atom. The phase of an ion interferometer can be read out by optically detecting the state of the ion. The fringes of an electron interferometer must be resolved spatially.

Since electrons are not easily manipulated with light, electron interferometers must use phys-

ical gratings. This means that the particles pass close to surfaces that can cause systematic shifts, decoherence, and noise in the phase of the interferometer. Also, each physical grating produces many diffraction orders. Only a small fraction of the total electrons contribute to the interference fringes, decreasing the signal to noise ratio.

Many of these limitations can be overcome in an optically driven ion interferometer. The ions will be produced from laser-cooled atoms and driven on an optical Raman transition to form the interferometer. Manipulating the internal state of the ions allows for high contrast, a high signal-to-noise-ratio, and the ability to optically read out the phase of the interferometer.

8.1.1 General Design of the Ion Interferometer

The interferometer will start with a slow beam of ions whose creation is described in section 3. High-voltage electrodes then accelerate the beam creating a mono-energetic, velocity-tunable, cold ion beam. At the left-hand side of Figure 8.1 the ion beam is optically pumped to be in one of the hyperfine ground states.

The ions then cross three counter-propagating sets of continuous-wave laser beams that serve the same function as the optics in a Mach-Zehnder optical interferometer. The laser intensity and time of flight of the atoms through the first set of cw laser beams is controlled in order to produce a $\pi/2$ pulse, which acts like a beam splitter. The ions are placed in a superposition of the two hyperfine ground states of $^{87}Sr^{+}$, and half of the quantum superposition recoils aways with a momentum of $2\hbar k$. The center set of beams delivers a π pulse that acts like the mirrors in the light interferometer. This pulse drives each half of the superposition to the other hyperfine state, directing the two arms of the interferometer back towards each other. Finally, the third set of beams delivers a second $\pi/2$ pulse to the ions as the two arms of the interferometer physically come back together. This pulse acts like the last beam splitter, mixing the quantum waves of the two arms. The quantum superposition of the atoms upon exiting the interferometer depends on the relative

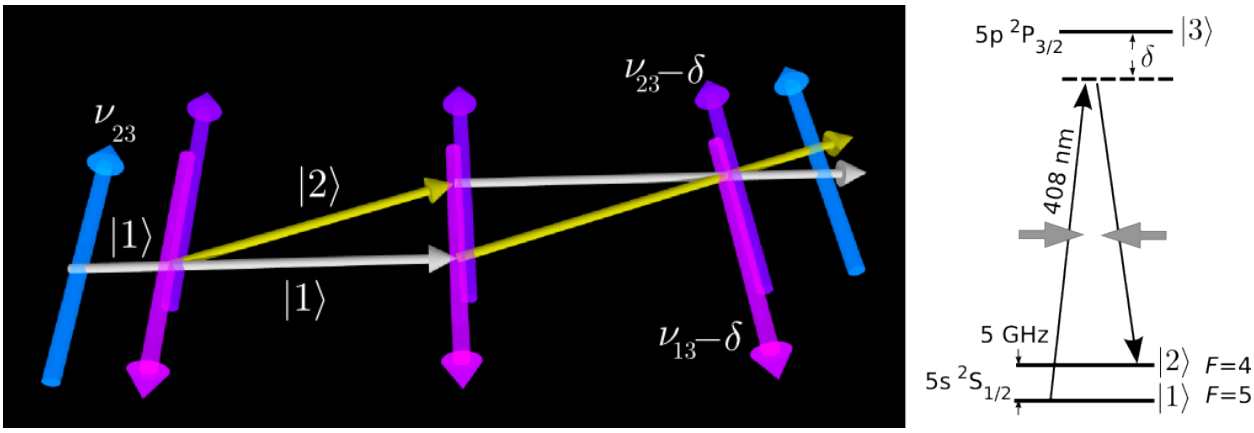

Figure 8.1 Ions enter the interferometer on the left hand side of the figure where they are optically pumped into one of the hyperfine ground states. A $\pi/2$ pulse splits the quantum state of the ions between the hyperfine ground states. The center pair of beams delivers a π pulse that swaps the momentum in each arm of the interferometer to bring them together. The ions then pass through a third pair of beams receiving a second $\pi/2$ pulse. At the end of the interferometer a fluorescence probe detects which hyperfine ground state the ions are in.

phase of the two arms of the interferometer. By "collapsing" the ion's wave-function through measurement with a fluorescence probe, the phase of the interferometer can be determined.

The Raman transition between the hyperfine ground states of $^{87}Sr^{+}$ is shown in the right hand side of Figure 8.1. Even though the hyperfine splitting corresponds to a microwave frequency, the transition can be driven with optical frequencies by the use of an intermediate energy state. By detuning two lasers from each 5s $^2S_{1/2}$ hyperfine ground state to the 5p $^2P_{3/2}$ state by the same amount a "virtual state" is created. An ion can be coherently transferred from one hyperfine ground state to the other by absorbing a photon from one beam and then by emitting the photon into the second laser beam via stimulated emission.

8.2 Laser Intensity and Transition Probability

8.2.1 Driving π and $\pi/2$ Pulses

The time an ion at a particular velocity must spend in each pair of laser beams to be driven to the π and $\pi/2$ positions on the Bloch sphere depends on the characteristic time it takes to complete the transition. This time depends on the Rabi frequency which is defined by

$$\Omega = \frac{-eE_0}{\hbar} < e|\mathbf{r}|g >= \frac{\mathbf{d}E_0}{\hbar}, \tag{8.1}$$

where $|e>$ denotes the excited state for the ion, $|g>$ denotes the ground state, and $\mathbf{d}$ is the transition dipole moment. While there are many definitions of Ω in the literature, Ω defined in this way gives the real oscillation frequency of the probability of finding the ion in the excited state when the ion is in a resonant light field. When the light field is detuned from the resonant frequency the effective Rabi frequency is

$$\Omega_{eff} = \sqrt{\Omega^2 + \delta^2}, \tag{8.2}$$

where δ is the detuning of the laser from the resonant transition frequency.

Other important parameters can be defined by solving the optical Bloch equations in the steady state. This not only yields the elements of the density matrix ρ for this two-state system but also provides motivation for defining the saturation parameter S

$$S = \frac{|\Omega|^2/2}{\delta^2 + \Gamma^2/4}. \tag{8.3}$$

From S the on-resonance saturation parameter used in Chapter 2 is

$$S_0 = \frac{2|\Omega|^2}{\Gamma^2} = \frac{I}{I_{sat}}, \tag{8.4}$$

where I_{sat} is as defined by Eq. 2.3 and the intensity of the laser is $I = \frac{1}{2}c\varepsilon_0|E|^2$. The Rabi frequency and effective Rabi frequency can now be written and calculated in terms of more easily accessible

parameters as

$$\Omega = \sqrt{\frac{\Gamma^2 S_0}{2}} \tag{8.5}$$

and

$$\Omega_{eff} = \sqrt{\left(\frac{\Gamma^2 S_0}{2}\right)^2 + \delta^2}. \tag{8.6}$$

However, this is only valid for a one photon, two state transition, i.e. calculating the Rabi frequency for the 5s $^2S_{1/2}$ to 5p $^2P_{3/2}$transition. It is not valid for the two photon process that drives the ion between the hyperfine ground states. As given in [110] the effective Rabi frequency for a stimulated Raman transition is given by

$$\Omega_{Raman} = \frac{2\Omega_1\Omega_2}{\Delta} \tag{8.7}$$

where $\Omega_{1,2}$ are the on-resonance Rabi frequencies of each state to the intermediate state and Δ is the overall detuning of the light fields from the intermediate state.

In order to effectively receive a π pulse, the ions must remain in the light field long enough for one half of a Rabi oscillation to occur at a given intensity. As seen from Eqs. 8.5 and 8.7, the frequency depends on the intensity of the light and the detuning of the virtual state from the 5p $^2P_{3/2}$ state.

Likely values for the intensity and size of the beam can be easily estimated. The diodes purchased for the Raman laser system are capable of putting out $\sim$ 100 mW of power free running. In a very conservative estimate, $\sim$ 5 mW of light might be delivered from each laser diode for the π pulse and half of that for each of the $\pi/2$ pulses.

Over the course of the interferometer the natural divergence of the laser should have a negligible effect on the ions. Therefore, a Rayleigh range of $>$ 1 m is desirable for the collimated beam. This could be satisfied by collimating the laser to a beam waist of $\sim$ .4 mm.

The peak intensity of a laser beam can be written as

$$I = \frac{2P}{\pi\omega_0^2}, \tag{8.8}$$

where P is the power of the laser and ω_0 is the beam waist of the laser beam. With the estimation that the intensity is uniform over the size of the beam, a laser with a .4 mm waist and 5 mW of power has an intensity of ~ 2 W/cm^2. If the detuning from the intermediate state is 2.5 GHz, the Rabi frequency for the Raman transition is $\Omega_{Raman} \backsimeq 120$ MHz. In this case, the characteristic time an atom needs to spend in the beam to receive a π pulse is $\tau_\pi \backsimeq 4$ ns.

At this intensity ions traveling at a velocity of $\sim 1.9 \times 10^5$ m/s would receive a π pulse. This corresponds to 16.7 keV ions, which is $3/4$ of the maximum energy the electrodes can impart to the ions. In order to accommodate faster ions or greater detuning, either more power would be required from the injection locked diodes, or the beam could be focused down in the vertical dimension in order to boost intensity while keeping the interaction time the same.

8.2.2 Spontaneous Decay

The overall detuning of the lasers from the 5s $^2S_{1/2}$ to 5p $^2P_{3/2}$ transition is very important as the lifetime of upper state is ~ 44 ns. The smaller the detuning of the lasers from the transition, the more likely it is that atoms will actually be driven to the excited state and spontaneously emit before stimulated emission into the second laser occurs. Atoms that spontaneously emit from the excited state can still be cycled into the other ground state but since their phase has been randomized their inclusion in the interferometer decreases the contrast.

The probability of decay out of the intermediate state is given by

$$p_{sp.decay} = \frac{\tau_\pi \Gamma \Omega}{2\Delta}, \tag{8.9}$$

where τ_π is the time it takes to receive a π pulse, Δ is the detuning of the virtual state, and Ω is the the Rabi frequency is for the single photon 5s$^2S_{1/2}$ to 5p $^2P_{3/2}$ transition. When the detuning $\Delta = 2.5$ GHz, $p_{sp.decay} \backsimeq .16$ for just the π pulse of the interferometer. While losing 16% of the ions to spontaneous decay will not wash out the interference signal, it is a significant loss in the

signal to noise ratio. This number can be improved by increasing the detuning[1], which in turn requires either slower atoms, or an increased laser intensity.

8.3 Interferometer Phase

The interferometer's output phase is sensitive to many different parameters including the physical setup of the experiment. For instance, in theory the phase difference due to different energy states of the ions cancels because the ion spends equal amounts of time in each state, but this assumes that the interferometer is well aligned. If not, a velocity dependent offset will be present. Such systematic phase shifts do not matter as long as they are constant and amount to a very stable overall shift. Less stable systematic phase shifts will necessarily need to be eliminated or well characterized.

As an ion interferometer, this apparatus will be sensitive to all of the same effects as neutral atom interferometers as well as electro-magnetic fields. Thus inertial and gravitational forces will also affect the phase of the interferometer as well as electric potentials and magnetic fields. The phase sensitivities that are common with neutral atom interferometers are well known and understood thanks to rigorous study of those devices. This treatment focuses only on the interferometer's sensitivity to fields and potentials particular to an ion. A "phase budget" due to other effects will have to be taken into account when collecting data with the interferometer in the future.

The phase difference between the arms of the interferometer due to an electric field is given by

$$\Delta\Phi = \Delta E \frac{\tau}{\hbar} = e\Delta\phi \frac{\tau}{\hbar} \tag{8.10}$$

where ΔE is the energy difference between the two paths of the interferometer, and τ is the time is takes for the ion to the travel a horizontal segment of the paths shown in Figure 8.1. The energy

[1] There is an exception in the case where Δ approaches the hyperfine splitting of the ground states. In this case, one of the lasers will be close to the resonance from one of the hyperfine states to the $5p\ ^2P_{3/2}$ state.

difference due to an electric field is the charge of the electron e times the electrostatic potential difference between the interferometer arms $\Delta\phi$. To first order, only the horizontal path segments should contribute to the phase difference since both of the diagonal paths cross the entire potential difference between the two arms, giving a similar phase shift to both arms.

8.3.1 Sensitivity to Electro-magnetic Fields

As can be seen from Eq. 8.10 the sensitivity of the interferometer depends on the interaction time of the ion with the electric field. This leads to wide range sensitivities based on the configuration of interferometer (see section 8.4.2) and the strength and scale of the field being measured.

The electric field sensitivity of the interferometer can be written as

$$E = \frac{2\hbar\Delta\Phi v_{ion}}{eL_{int}\beta}$$

$$= \frac{2\hbar\Delta\Phi}{eL_{int}\beta}\sqrt{\frac{2E_{ion}}{M}}, \tag{8.11}$$

where β is the separation of the interferometer arms, L_{int} is the length of the interferometer, E_{ion} is the energy of the ions, $\Delta\Phi$ is the phase shift in radians, and M is the mass of the ion. The electric field strength required to incur a 1 radian phase shift for the different interferometer configurations and ion velocities possible in the apparatus are given in Table 8.1. The calculations assume the field is present across the entire length of a given interferometer configuration.

While simple multimeters can read sub-mV potentials, it must be remembered that free space fields are much harder to measure. The predicted sensitivity of the interferometer is far higher than any other device measuring free space fields. In fact, most state of the art free-space measurement devices of DC electric fields, with the exception of electron interferometers, only have sensitivities of $100-1000$ V/m [111]. Even when running at the maximum possible velocity for our apparatus this device will be as sensitive as the most sensitive electron interferometers.

For $\Delta\Phi = 1$ rad.	$L_{int} = .4826$ m	$L_{int} = .9652$ m	$L_{int} = 1.168$ m	$L_{int} = 1.930$ m
$E_{ion} = 10$ eV	11.1 μV/m	2.79 μV/m	1.90 μV/m	697 pV/m
$E_{ion} = 5$ keV	5.57 mV/m	1.39 mV/m	952 μV/m	348 μV/m
$E_{ion} = 10$ keV	11.1 mV/m	2.79 mV/m	1.90 mV/m	697 μV/m
$E_{ion} = 20$ keV	22.3 mV/m	5.57 mV/m	3.81 mV/m	1.39 mV/m

Table 8.1 Electric field sensitivities possible with the different configurations of the interferometer. The sensitivities are calculated based on fields that are present across the entire interferometer and that are strong enough to create a phase shift of 1 radian. The highest sensitivity expected under these circumstances is ~ 697 pV/m.

8.3.2 Sensitivity to Background Vacuum Gasses

The interferometer's phase will also be sensitive to residual background gasses whose polarizability in the presence of a charged ion could create a potential difference between the arms of the interferometer. Such a potential can be written as

$$V_{dip} = \frac{k_e \mathbf{d} cos\theta}{r^2} = \frac{k_e \alpha \mathbf{E} cos\theta}{r^2} \tag{8.12}$$

where $\mathbf{E} = k_e q/r^2$ is the field of a point charge, k_e is the Coulomb constant, and α is the polarizability of the background particle. If we assume the worst case scenario that the dipole will always be oriented such as to create the strongest field possible the magnitude of the potential can be written as

$$V_{dip} = \frac{k^2 \alpha e}{r^4} \tag{8.13}$$

For a particle that passes within a distance b of the ion beam we can write the the total phase shift for the interferometer as

$$\Delta\Phi = \frac{k^2 \alpha e^2}{\hbar} \int_{-\infty}^{\infty} \left(\frac{1}{r_1^4} - \frac{1}{r_2^4} \right) dt = \frac{k^2 \alpha e^2 \tau \pi}{2\hbar v} \left(\frac{1}{b^3} - \frac{1}{(b+\beta)^3} \right) \tag{8.14}$$

where $r_1^2 = b^2 + v^2 t^2$, $r_2^2 = (b+\beta)^2 + v^2 t^2$, v is the relative velocity of the particle to the ion, and β is the separation of the interferometer arms. For 6 keV ions and a background gas at room

Molecule	Polarizability (a.u.)	$\Delta\Phi$ radians
H_2	.7	.0485
O_2	10.6	.7357
N_2	11.74	.8148
H_2O	9.6	.666

Table 8.2 The table shows the polarizability of background gasses in the vacuum and the relative phase shift they produce between the arms of the ion interferometer. The phase shift given is for 6 keV ions in a 1 meter long interferometer (arm separation of 97 nm) with one arm passing within 4 angstroms from the background molecule. The polarizabilities are from experimental data presented in [112–114] and are given in atomic units.

temperature ($\sim$ 300 m/s), the velocity v is essentially that of the ion's because the velocity of the background molecule is much smaller. In this case, the separation of the arms for a meter long interferometer is $\beta = 97$ nm. The phase shifts created by passing different background particles at an impact parameter of only 4 angstroms are given in Table 8.2.

Table 8.2 gives an idea of the vacuum level required in the interferometer chamber to prevent residual gasses from affecting the phase. N_2 has the largest polarizability of background gasses likely to be present in significant quantities in the chamber [2]. In order for a phase shift due to an N_2 molecule to equal one hundredth of a radian the impact parameter would be 1.7 nm. The collisional cross section based on this impact parameter is then $\sigma = 9.4 \times 10^{-18}$ m^2. For a mean free path much longer than the interferometer chamber ($\lambda \sim 100$ m) this implies a minimum vacuum pressure of 3.3×10^{-8} Torr. As the ion energy is decreased to increase sensitivity, the limits on the background pressure will necessarily be lower. The chamber can currently reach pressures as low

[2] Many polarizabilities are given in units of length cubed. Atomic units for polarizabilities are in r_0^3, where r_0 is the Bohr radius. By converting from atomic units to Gaussian units of cm^3 the polarizability can then be easily converted to S.I. units.

as 10^{-10} Torr.

8.4 Experimental Layout and Design

8.4.1 Interferometer Scheme and Vacuum Chamber Versatility

The apparatus has been designed to accommodate a large range of experimental parameters. The chamber itself has been constructed in such a way that the interferometer can be set up with four different total lengths: 19 in., 38 in., 46 in., and 76 inches. For the two interferometers longer than the mid-point of the chamber, the horizontal 6 in. CF port can be used to introduce materials in the apparatus so that the interferometer can probe their electromagnetic properties.

The electrodes at the beginning of the chamber can be used to speed ions up to 20 keV. Additionally, it might also be possible to conduct slow electron experiments by reversing the voltage on the electrodes. Such an electron source would be mono-energetic with a sub-thermal velocity distribution. A Magnum channeltron has been purchased for the chamber that could be used for charged particle detection in future experiments as well as troubleshooting and diagnosing the ion beam.

8.4.2 Beam Optics and Relative Phase

Managing the different physical properties of the beam is very important as shown above. Equally important is the relative phase of the laser beams as the ions traverse the interferometer. If the coherence length of the laser is too short the contrast of the interferometer will suffer as a result since the phase differences between the optical gratings can affect the phase of the interferometer.

In order to reduce phase noise written onto the lasers from the mechanical vibrations of optics the layout shown in Figure 8.2 is proposed. As seen in the figure, the two lasers driving the Raman transition are passed through the same optics. Fifty-fifty beam splitters are used to produce the

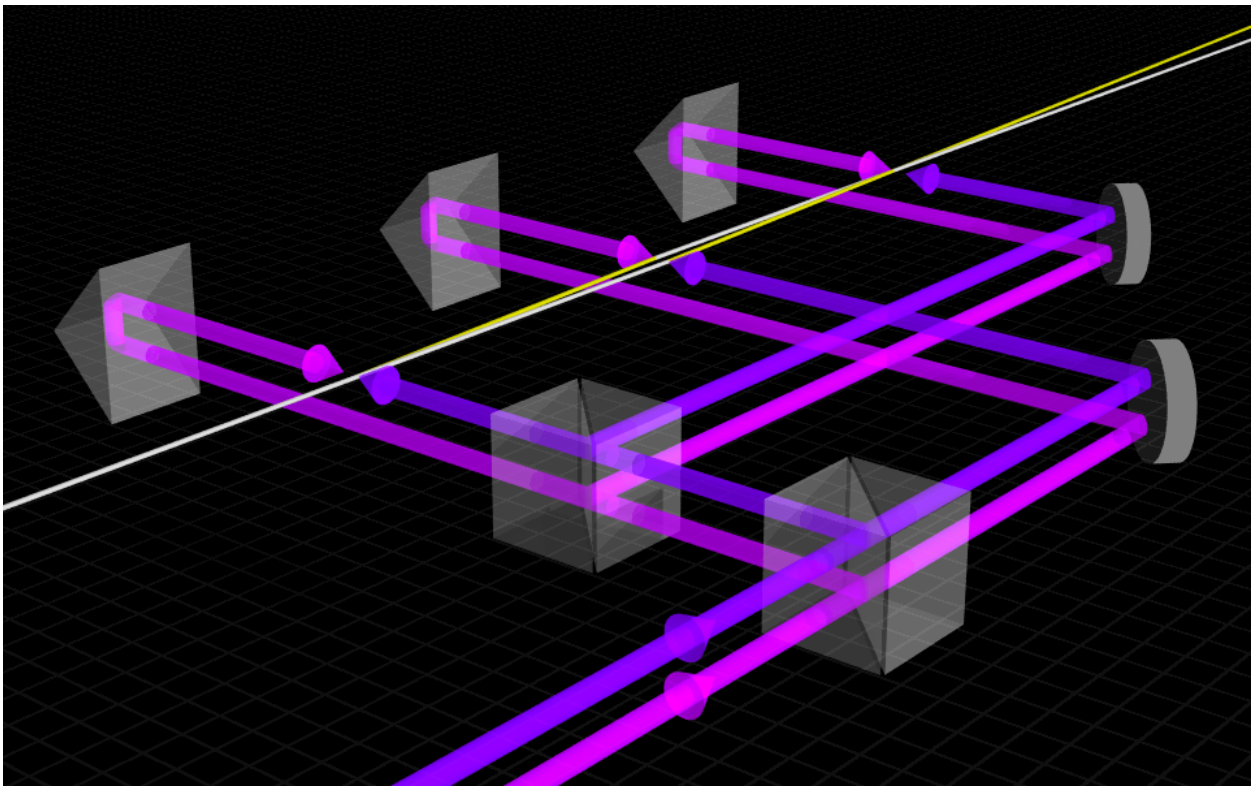

Figure 8.2 Illustration of the basic optics layout for the interferometer. The two lasers used to drive the Raman transition are sent through the same optics displaced from each other in the vertical direction. Beam splitters are used to create the right intensities for the interferometer's pulses. Right-angle prisms are used to direct the beams counter-propagating across the ion beam.

correct intensity of light for all three pulses. The beams are aligned vertically so that when sent through the vacuum chamber one beam passes through the ion beam while the other one passes bellow it. On the opposite side of the chamber a right-angle prism then reflects the beams sending each one counter-propagating against the other. In this configuration, phase noise is minimized since the beams all strike the same optics. However, motion of the prism or vacuum chamber windows will not be common to the two beams, and a rotating optic will also affect one beam differently than the other.

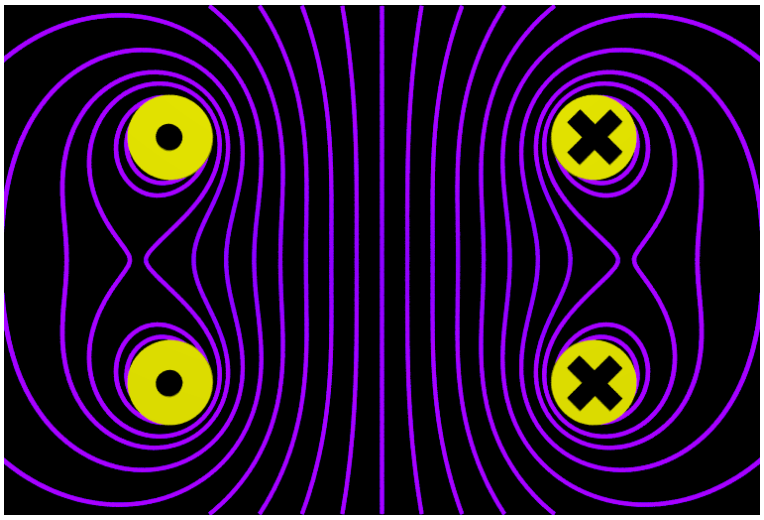

Figure 8.3 The configuration of currents used to produce a magnetic field to establish a quantization axis for the experiment and the field lines are shown. The "x's" indicate current flowing into the page while the dots indicate current flowing out of the page. The wires are spaced to make the B field very uniform across the propagation axis of the ions.

8.4.3 Quantization Axis and Magnetic Field

In order to ensure that stray magnetic fields do not cause an ion's quantization axis to "wander" through the experiment, a magnetic field is created to provide a definite quantization axis. While the quantization axis for an experiment can be easily created by circular coils, the length and bulk of our chamber prevented implementing coils of this nature. Instead, in our apparatus a uniform vertical magnetic field through the center of the chamber will be created by running a current through two sets of elongated coils. As seen in Figure 7.9 the coils loop the chamber lengthwise and are essentially a 2-dimensional analog of Helmholtz coils.

Figure 8.3 shows the current configuration and the magnetic field lines. Looking down the axis of the interferometer chamber the magnetic field is produced by currents that flow in opposite directions for each side of the chamber, much like a Ioffe-Pritchard trap.

The field strength on the axis of the ion beam can be calculated from standard electro-magnetic theory by Ampere's Law. The magnetic field created at a point r by an infinite current carrying

wire centered on the z-axis can be written in Cartesian coordinates as

$$B = \frac{\mu_0 I}{2\pi(x^2+y^2)}(-y\hat{i}, x\hat{j}), \tag{8.15}$$

where μ_0 is the permeability of free space, and I is the current. If we shift the location of the wire off axis by an amount b in the $\hat{i}$ direction and an amount a in the $\hat{j}$ direction we can then write

$$B = \frac{\mu_0 I}{2\pi((x-b)^2+(y-a)^2)}(-(y-a)\hat{i}, (x-b)\hat{j}). \tag{8.16}$$

By writing down the magnetic field for the four current carrying wires shown in Figure 8.3, the field on the propagation axis of the ions $\hat{z}$ can be shown to be

$$B = \frac{-4b\mu_0 I}{2\pi(a^2+b^2)}\hat{j}. \tag{8.17}$$

The first and second derivatives of the magnetic field on the z-axis can be made to go to zero by carefully selecting the values for a and b. This results in a very uniform field near the axis. The relationship between a and b that satisfies this condition is

$$\frac{b}{a} = \sqrt{3}. \tag{8.18}$$

It is important that the magnetic field generated by the coils be sufficient to overwhelm all other magnetic fields that may be present. The strongest magnetic field with the potential to affect the ions is the one created by the magnet attached to the ion pump that sits directly above chamber. A field of ~ 1 Gauss is generated by this magnet at the location of the ion beam as measured by a Hall probe, which is on the order of the earth's magnetic field.

The two coils each comprised of 7 loops of 10 gauge enamel coated wire with the dimensions .24 m x 2.7 m and separated by .207 m have been mounted around the interferometer chamber and can create a uniform vertical magnetic field of ~ 100 G down the central axis of the chamber when 50 A passes through the wires.

Chapter 9

Conclusions and Outlook

9.1 Current State of the Experiment

To date nearly the entire apparatus necessary for realizing an ion interferometer has been constructed. The only piece of equipment not yet operational is the laser system needed to drive the Raman transition between the ion's hyperfine ground states, which is the subject of a Master's thesis for another group member.

Figure 9.1 shows a picture of the equipment assembled for this project. The vacuum chambers are mounted to and supported by a structure of Faztek extruded aluminum pieces, and the laser systems are protected by HDPE enclosures. All of the optics, mounts, and other required parts necessary to finish the project have been purchased and are on hand.

As of this writing we have witnessed the formation of MOTs for the ^{87}Sr, ^{88}Sr, and ^{86}Sr isotopes. Figure 9.2 shows a picture of the ^{88}Sr MOT. Based on rudimentary tests consisting of blocking portions of the trapping beams, it is likely that the trap is in fact centered such that atoms are streaming through the hole in the elliptical mirror along the desired beam path.

We have not yet conclusively measured the output of the LVIS with fluorescence. To detect the

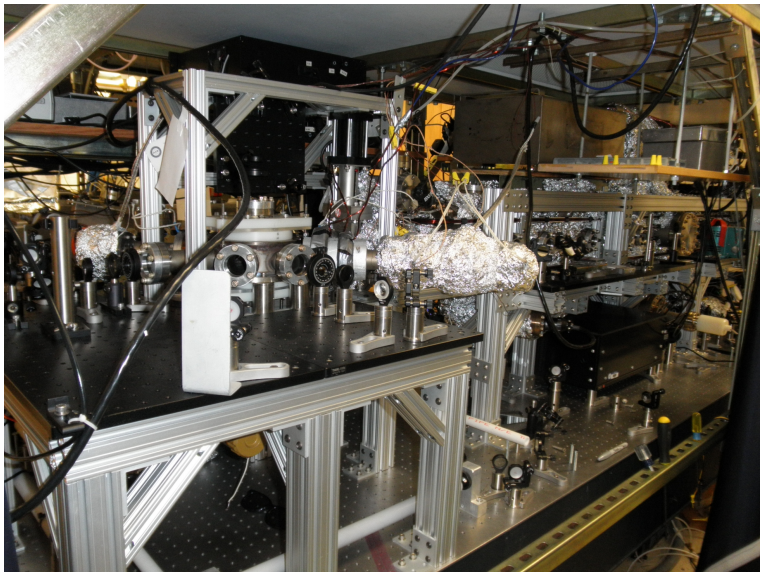

Figure 9.1 Picture of the experimental apparatus. The MOT chamber is at the left of the picture and the interferometer chamber extends down the length of the optics table. On the table but unseen are the 461 nm laser system, the 407.8 nm fluorescence probe, the vapor cell and saturated absorption lock setup, and the existing parts of the Raman laser system.

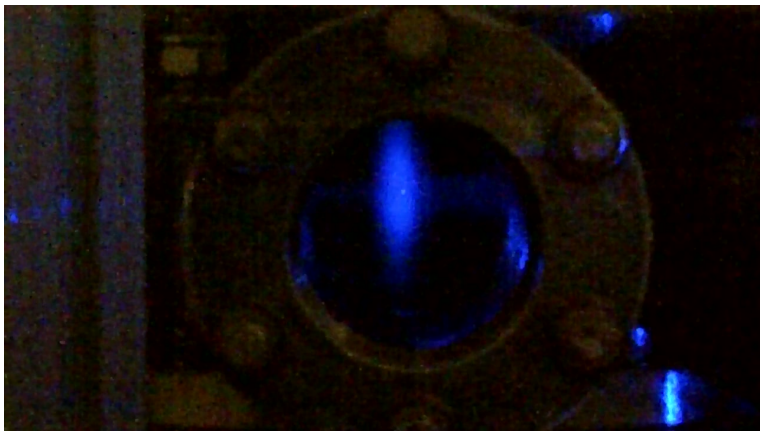

Figure 9.2 Picture of an ^{88}Sr MOT formed in the experiment. The bright dot in the middle of the fluorescence is the trapped atoms.

atom beam light from the 461 nm laser system was focused down and scanned in position across the path of the beam. However, since the expected LVIS flux is on the order of $10^5 - 10^6$ atoms per second there only exists $\sim 10^9 - 10^{10}$ atoms/m^3 in the volume illuminated by the laser, which density is several orders of magnitude lower than the density required to absorb just 1% of the light. This amounts to a fluorescence signal that is much weaker than the significant amount of scattered light from the MOT optics.

Other work that remains to be done, aside from the raman laser system, includes: ionizing the LVIS beam, tuning the fluorescence laser to the correct wavelength, frequency locking the fluorescence laser, and aligning the interferometer optics.

9.2 Contributions to Science

Over the course of building this apparatus scientific equipment has been built and experimental techniques have been developed that may be of interest and use in the greater scientific community.

9.2.1 MOT/LVIS Fields from Permanent Ring Magnets

One powerful technique is the application of ring magnets in the LVIS as described in section 3.2. In applications where constant or fine control over the magnetic field is not required a pair of permanent ring magnets can replace a pair of anti-Helmholtz coils as well as coils for Zeeman slowing and accelerating.

This is because in an anti-Helmholtz configuration the ring magnets create two maxima in the magnetic field extending away from the center of a MOT. The first maxima is smaller and the gradient leading to it can be set for trapping atoms. The second maxima is typically an order of magnitude larger providing a gradient that extends out over a significant distance. The field rising to or dying away from the second maxima can then be used for both slowing atoms approaching

the trap and for accelerating atoms leaving the trap. This feature makes the ring magnets well suited for LVIS applications.

The permanent ring magnets are much more compact in size than a wire coil that could produce an equivalent field gradient. This makes it easy to incorporate them into an apparatus with many parts or in confined spaces. Also, the permanent magnets do not require a current source of any kind, or thermal management.

9.2.2 Vapor Cells for Alkaline Earth Metals

In the course of building the experiment I also developed several different vapor cells suited for spectroscopy with alkaline earth metals as well as any other strongly reactive element having a similar vapor pressure.

Colleagues in the scientific community have already inquired about one of the vapor cells developed: the high-temperature cell built for the calcium interferometer experiment mentioned in the Foreword and described in [1, 7]. This cell features a dual-chamber design: the central portion of the cell contains a calcium sample separated from the rest of the cell by two calcium-fluoride windows with a seal similar to a Kasevish seal [108], and the rest of the cell is sealed with standard CF viewports. A calcium vapor at temperatures $> 800^\circ$ C can be maintained in the central portion of the cell.

The vapor cell described in Chapter 6 has a versatile design for producing vapors at lower temperatures. As is, the cell can easily be run in much the same way as a standard "six-shooter" vapor cell (where a broad optically-dense thermal beam of atoms crosses between two viewports), or as a standard vapor cell. Despite the fact that the cell is constructed from standard CF fittings with one simple modification, the atomic sample can easily be heated to $\sim 730^\circ$ C, which is well above the temperature rating of the fittings.

The bottom part of this cell can also be removed and used as a small-form vapor cell. In

this configuration the viewports have a direct line of sight to the atomic sample. This would normally cause the windows to become coated while evaporating oxide layers off of the sample. To prevent this I developed a technique to protect the windows during the oxide evaporation, which is described in Appendix D. In the method a foil sheet of a low melting-point metal, such as lead or indium, is used to cover the viewports while the oxide is evaporating. The oxide layer then coats the foil rather than the windows. Afterward, the foil can be melted and the cell run at temperatures within the range of the CF flanges.

9.2.3 Diode Laser Electronics Systems

The electronics suite developed for ECDL laser systems described in Chapter 5, have to date generated the most interest from colleagues in the field of AMO research.

In particular, there has been a lot of interest in the low-noise laser diode current driver I developed. The driver can be configured for either current polarity and has produced currents up to 5 A [3, 4]. The current can be modulated by an external source and the modulation bandwidth has a 3 dB point of ~ 50 MHz, which is much higher than the gain-bandwidth product of the op-amps used in circuit. The current driver is quieter than any other commercial or home-built driver known to us with the noise spectral density hitting a floor of 10^{-10} A/$\sqrt{\text{Hz}}$ at ~ 50 Hz.

I also developed a PID lock circuit built from the same low noise components as the current driver. The lock circuit allows the overall gain of the feedback for the diode current and the piezo mount voltage in an ECDL to be controlled independently. The ratio of the voltages for each side of the piezo mount is controlled with a trimpot, which does not affect the gain of the signals. Additionally, the circuit allows for scanning the frequency of an ECDL. The AR coated 922 nm master laser described in section 4.1.1 has been scanned ~ 30 GHz mode-hop free using this circuit. The PID controller was also used for building a laser system on the calcium interferometer with 1 kHz stability, where the total servo bandwidth was ~ 4 MHz [7].

We have also built several other circuits for driving and locking ECDL systems. For use with the PID lock-circuit we built a photo-amplifier that also performs homodyne detection. A temperature controller and piezo voltage amplifier have also been constructed.

To simultaneously control and monitor these electronics, I have built and programmed microcontroller units. These units allow ECDL electronics to be positioned close to the laser head and even sealed inside a laser's enclosure, keeping path lengths short. Meanwhile, the microcontroller unit can be placed just about anywhere convenient. Each microcontroller unit can change and monitor a current driver, temperature controller, and PID lock circuit at once.

In order to quickly assemble and modify these circuits I developed small-production scale surface mount soldering techniques, which are described in Appendix B. These techniques not only allow the use of high-speed low-noise chips that are only available in surface mount packages, but also allow a circuit to be built much more quickly than hand soldering each circuit component.

A half-dozen labs have contacted us asking for advice on building and using our current drivers. They have told us that the driver has become a lab standard for their diode laser systems. The temperature controllers, PID controllers, and the microchip controllers have also been inquired about. Additionally, the surface mount soldering techniques used to build the circuits have been implemented in at least two other labs who have contacted us for advice.

9.3 Conclusion

This work describes the development of a charged particle interferometer that will have unprecedented sensitivity. This apparatus will serve to open up a new sub-field in atomic physics based on ion interferometry. No other apparatus will be able to perform the same experiments in precision metrology as the ion interferometer described here. It will be capable of greater sensitivity and contrast than current state of the art experiments for measuring free space electric fields. This ion

interferometer is uniquely suited for conducting fundamental studies of electro-magnetic theory, precision measurement of free-space electric fields, and material properties.

The fundamental and practical studies planned for this interferometer may help in developing better hydrogen storage materials, superconducting materials, optical materials, and characterizing the patch charge potentials in various materials. Other experiments the apparatus might be suited for include "delayed-choice" experiments and studies of the Aharanov-Bohm effect. Future developments of the apparatus will lead to performing fundamental studies of the currently accepted electromagnetic theory. Specifically, small deviations from Coulomb's Law could be detected with a matter-wave ion interferometer improving the current measurement by two orders of magnitude [5].

Appendix A

Tutorial: Stabilizing Laser Diodes

Diode lasers have come to play a fundamental role in experimental physics. They can provide an inexpensive alternative to solid state, gas, and dye lasers. Additionally, laser diode wavelengths are tunable in a small range around their central frequency and are now commonly available at many wavelengths in the near infrared and visible spectrum. Their availability and efficiency have made them a standard part of many experiments from spectroscopy to atomic trapping and cooling.

While many turn-key tunable diode laser systems are available, there are several disadvantages when considering them for use in fundamental research. To begin with, they are fairly expensive and they are not easily serviced on site when problems arise. It is also difficult to modify them for a special purpose or specific experiment. Above all, commercial laser systems do not provide the extreme stability and narrow linewidths needed for many of the applications in fundamental research. As such, many atomic physics laboratories build their own laser diode systems tailored to their needs. In our lab we have developed many diode laser systems including one with a linewidth of a few kHz.

A.1 Single-mode Diode Lasers

Diode lasers, like all lasers, consist of mirrors forming an optical cavity that contains a gain medium capable of coherently amplifying light quanta. Just like a string instrument, the optical cavity is a resonator. Light can only persist in the cavity if it is oscillating in a resonant mode of the cavity. Each mode in an optical cavity is characterized by a spatial distribution and a mean oscillation frequency.

Unlike an instrument's strings, the optical cavity's modes are three-dimensional but it is possible to separate each mode of the optical cavity into one function of the distance along the axis of the cavity and another function of the radial distance from the axis of the cavity. The first function is known as the "longitudinal" mode, and the second as the "transverse" mode. The longitudinal mode is very similar in form to the modes of a vibrating string. The transverse mode determines the shape of the laser beam, or rather the shape of the spot the laser beam will make if projected onto a piece of paper. The frequency of a laser mode depends on both its longitudinal and transverse components [115].

The properties of the beam produced by the laser depend strongly on the interplay between gain medium and the different modes of the diode's cavity. The bandwidth of the gain medium, or the range of wavelengths or frequencies of light that the medium can amplify, typically extends over many cavity modes, making it possible for the laser to oscillate in multiple modes simultaneously. To achieve a stable, narrow linewidth beam of laser light, however, it is important for the laser to produce oscillation in only one cavity mode. Otherwise the light from the laser will be composed of many different frequencies, which is undesirable when performing scientific experiments where narrow atomic transitions are driven. Additionally, multiple transverse modes create poorly shaped beams that are hard to manipulate spatially and which cannot to be collimated well or focused to a small spot.

It is possible to achieve single mode operation of a laser by taking advantage of the fact that

"stronger" cavity modes tend to quench "weaker" modes. This is because the gain medium of a laser is nonlinear. It amplifies light through the process of stimulated emission, and the rate of stimulated emission in any mode is proportional to the optical power in that mode. If one mode has greater gain, or lower loss, than all of the others power in that mode will grow faster than in other modes. As the power in the mode increases, the rate of stimulated emission in the mode also increases causing the mode to pull energy out of the gain medium at an increasing rate, leaving less energy behind for other modes. Eventually the laser will reach the point where the dominant mode is extracting energy from the gain medium almost as fast as energy is put into the gain medium, and all of the other modes will "starve." This struggle for energy within the gain medium is known as "mode competition."

Since the gain medium's response curve in a diode laser is very flat near its maximum, tiny perturbations due to current or thermal drift and mechanical noise can alter the properties of the laser enough to change which modes are dominant, causing the laser to "mode hop." Additionally, laser diodes suffer from "spatial hole burning," a process where regions of the gain medium not "filled" by the dominant mode can be dominated by a mode with a different spatial distribution. As such, unstable multi-mode operation typically occurs in free running diode lasers, resulting in linewidths on the order of a GHz or even much larger.

To achieve single mode operation in a diode laser, it is necessary to create a situation where one mode is clearly dominant. This is typically done in two steps: first by making one particular transverse mode dominant, and then by introducing conditions that favor a particular longitudinal mode. Typically the first step is done by introducing losses at the edge of the cavity. This favors the TEM_{00} transverse mode who's spatial distribution is confined around the axis of the optical cavity. In a diode laser this is typically done in manufacturing by limiting the radial size of the gain medium, or by changing the index of refraction of the gain medium radially from the axis to introduce losses in other transverse modes. Not all diodes are made this way, so it is important

to purchase a laser diode which is labeled "single mode" if single transverse mode operation is required. It is also important to note that the label of "single mode" on a laser diode typically refers to the transverse mode only. Except for some very special diode lasers (known as DBR and DFB lasers), a bare "single mode" diode can simultaneously oscillate in multiple longitudinal modes.

A.2 Achieving Single-mode Operation

To achieve single longitudinal mode operation, a frequency selective element is introduced to favor one mode above all others. In practice, this is commonly done by placing a reflective diffraction grating in front of the laser diode in a Littrow configuration [116–118]. In this setup the first order diffraction of a particular frequency component of the incident laser beam is directed back into the laser diode. The diffracted light feeding back into the diode effectively reduces losses for the longitudinal mode at the frequency that couples back to the diode most efficiently, resulting in single mode operation. Furthermore, this frequency depends on the angle of the grating – so the grating angle can be used to tune the laser to the desired wavelength.

The Littrow configuration is known as a type of extended cavity diode laser (ECDL) because a second laser cavity is effectively created between the diffraction grating and the rear facet of the diode. In order for this configuration to be stable, the distance to the diffraction grating must now be controlled to match an existing longitudinal mode of the laser diode as depicted in Figure A.1. As the angle of the grating is tuned, the position of the grating and the laser diode current or temperature must be adjusted to keep the external and internal cavity modes matched. Furthermore, to keep the laser running in a single mode, the temperature of the diode, the current running through the diode, and the position and angle of the grating relative to the diode must be stable.

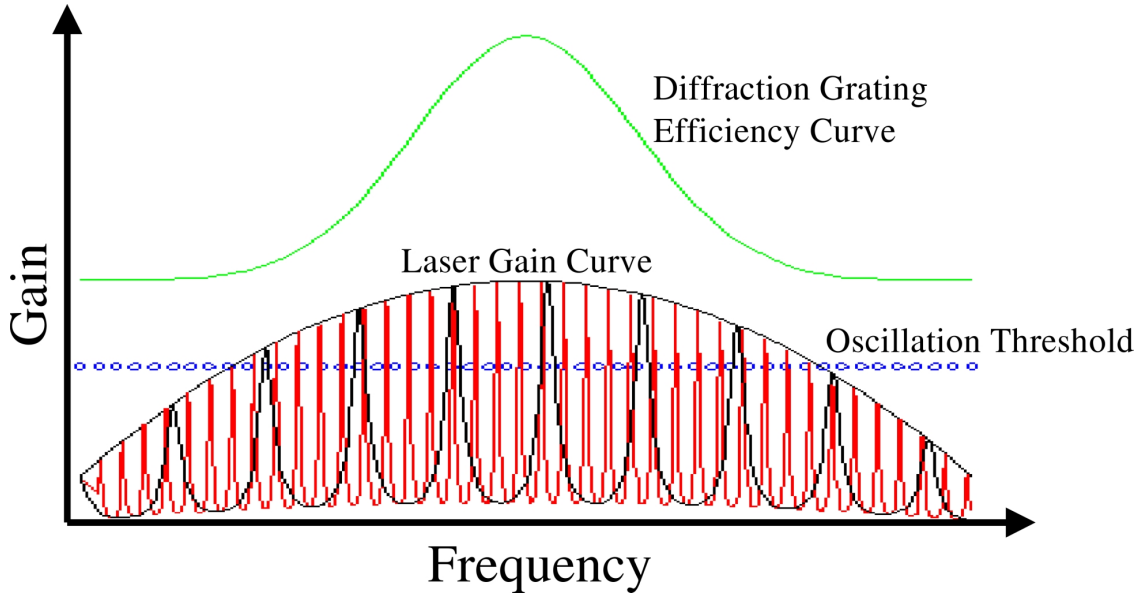

Figure A.1 The gain envelope of a laser is displayed in reference to the resonant frequencies of both the laser cavity and an external cavity formed by a diffraction grating. The fewer wide-band frequencies correspond to the free spectral range of the diode's cavity while the narrower closely-spaced frequencies correspond to the external cavity. Without the external cavity, several of laser diode's modes fall within a range of the gain curve that boosts them above the oscillation threshold leading to multi-mode output of the laser. With the external cavity only frequencies that match resonance modes for both cavities and have enough gain to overcome the laser oscillation threshold (dotted line) and that fall within a favorable area of the diffraction grating efficiency curve (top curve) can exist.

A.3 Schawlow-Townes Limit

Even when single mode operation is achieved, the laser will not generate a pure sine wave. Any fluctuations in the laser – thermal drift, motion of the grating, noise on the current driving the laser diode – will rapidly influence the laser frequency at any given time. The average of this drift over the time that an atom in an experiment interacts with the light results in an effectively broadened linewidth. These sources of technical noise are typically the limiting factor in a high-stability laser, but there is also a fundamental limit on the linewidth of a laser, even under ideal operating conditions.

One of the fundamental limits on a laser's operation is the finite linewidth of the optical cavity modes. Due to losses in the laser cavity, the frequency of each cavity mode is not sharply defined; light in any given single mode will contain a band of frequencies. The amplitudes of the various frequency components typically have a Lorentzian distribution about a central frequency. The characteristic width of this distribution is known as the linewidth of the mode and it is fundamentally determined by damping of the mode due to imperfect mirror reflectivity and other mechanisms that reduce intra-cavity light. This damping is characterized by a parameter known as the finesse. In short, a cavity with greater losses will have a lower finesse and the cavity modes will have a broader linewidth.

Another fundamental limitation in a laser is spontaneous emission of light within the gain medium. Even though the gain medium amplifies light through stimulated emission, spontaneous emission will also occur. A photon emitted spontaneously into a cavity mode will have a random phase, uncorrelated with the light already present. This results in a random phase shift given to the light in the mode. The broadening due to spontaneous emission depends on the ratio of the spontaneous emission rate to the rate of stimulated emissions in the mode. Whereas spontaneous emission is independent of the power within the lasing mode, the rate of stimulated emission increases with the power in the mode. As such, reducing losses in the cavity and increasing the

power circulating within a given cavity mode will decrease the fundamentally achievable linewidth. However, this also increases susceptibility to technical noise such as thermal drift.

From an understanding of the above-stated processes the fundamental limit on a laser's linewidth, known as the Schawlow-Townes linewidth, can be derived. This limit is given by

$$\Delta\nu_{laser} = \frac{\pi h\nu(\nu_{cav})^2}{P_{out}} \tag{A.1}$$

where $h\nu$ is the energy of a single photon, ν_{cav} is the bandwidth of the laser resonator, and P_{out} is the output power [119]. Note that, due to the nonlinear nature of lasers, this limit is smaller than the linewidth of the standalone optical cavity of the laser as long as there are enough photons in the cavity at any given time.

Due to the very short length of laser diode cavities and the high transmission of the output mirror on the cavity, laser diodes have large resonator bandwidths resulting in typical Schawlow-Townes linewidths of a few MHz [120]. This limiting linewidth can be improved by lowering the resonator bandwidth ν_{cav}, which is automatically achieved by using a Littrow configuration to force single mode operation. The diffraction grating lowers ν_{cav} in two ways. First, it lengthens the effective laser cavity. Second, it provides a more reflective surface for the output coupler of the laser than the diode face itself.

Even in the most stable lasers, technical noise typically broadens the linewidth of the laser far beyond the Schawlow-Townes limit. As such, much effort is put into reducing the technical noise of diode based laser systems.

A.4 Passive Stabilization

Mechanical/acoustical vibrations are very often a major factor in broadening the linewidth of the laser. Small movements of not only the laser diode head but the mirrors and optics that guide the laser beam can write phase noise onto the light and create instability even with an active frequency

lock. In an effort to minimize the effects of all such noise, laser systems are often constructed on "floating" optics tables: optical tables whose legs have air filled bladders to decouple the table surface from vibrations through the floor. Additionally, the entire system can be housed on a single breadboard which can be further isolated by resting it on a material such as Sorbothane. For laser systems with extremely narrow linewidths even more extreme measures have been taken employing surgical tubing, leaded foam, and other such materials in an attempt to damp out mechanical vibrations.

Thermal noise and drift will also significantly affect the laser. The diode is susceptible to thermal drift because the optical path length of the resonator cavity as well as the wavelength dependence of the gain curve depends on temperature [120]. Thermal fluctuations can therefore cause instability in the output wavelength of an ECDL. To minimize this effect, all of the laser diodes in this experiment are housed in a temperature-stabilized mount. While we do not actively temperature-stabilize any mount's base or surroundings, they are somewhat protected from room temperature fluctuations as they are mounted to individual breadboards and enclosed in high-density poly-ethylene (HDPE) boxes. The temperature controller for the mount is discussed in section 5.2.

Electrical noise and drift is the final major source of noise on the laser. First, the injection current of the laser diode affects the temperature of the diode through Joule heating. Additionally, the injection current affects the carrier density of the semiconductor, which in turn affects the index of refraction of the gain medium. In order to reduce issues arising from electrical drift, all of the diodes are driven with a low-noise current driver that is discussed in section 5.1.

A.5 Active Stabilization

Once technical noise has been minimized through passive stabilization, the linewidth of an ECDL can be actively narrowed by locking the laser to the mode of a high finesse Fabry-Perot optical reference cavity. The act of locking the laser to the cavity has two effects. First, by adjusting the laser to keep the light resonant with a mode of the reference cavity, long-term drift in the laser is reduced. Second, by using a high-bandwidth lock circuit, high-frequency technical noise can be detected and compensated for, bringing the linewidth of the laser closer to the Schawlow-Townes limit.

Since the reference cavity is not a laser, it does not contain an active gain medium, it does not require a frequency-selective element to assist mode competition, and it does not need a fairly transparent output coupler to let laser power out of the cavity. This allows the cavity to have a much higher finesse, greater stability, and much narrower resonance linewidths than the laser cavity.

When the laser is on resonance with a mode of the optical cavity, light is transmitted through the cavity. Locking the laser to the optical cavity can therefore be realized by monitoring the amount of light transmitted through the cavity. Unfortunately, there is significant delay between the onset of a drift between the laser or cavity frequencies and the change in the power of the light transmitted through the cavity. As the finesse of an optical cavity is increased so is the number of times a photon reflects off of the cavity mirrors before leaving the cavity. This increases the time it takes for the cavity to ring down and reduces the bandwidth of the lock to the cavity.

For this reason the light reflected from the front mirror of the cavity is typically monitored. The light off the front mirror of the cavity can be thought of as a sum of light reflecting directly back from the incident beam and circulating light in the cavity leaking out. When the laser is on resonance with the cavity, these two components are exactly out of phase and cancel each other. If the laser drifts off of resonance, the phase of the laser will evolve at a different rate than the phase of the light in the cavity. The two components of light will not be completely out of phase

anymore, resulting in an immediate change in the measured reflected power.

A typical technique used to lock a diode laser to a high-finesse optical cavity using light off of the input mirror is the Pound-Drever-Hall technique [121]. In this method the phase of the light incident on the optical cavity is modulated at a frequency ω_{rf}. If the laser has a frequency ω_{laser} before being phase modulated, after phase modulation the laser will consist of a main carrier frequency at ω_{laser} plus two sidebands at frequencies $\omega_{laser} + \omega_{rf}$ and $\omega_{laser} - \omega_{rf}$. The beating of the lower sideband with the carrier is out of phase with the beating of the upper sideband of the carrier, resulting in no net beating at ω_{rf}.

When the laser is on resonance, the reflected carrier is 180° out of phase with the light leaking from the cavity, resulting in a reflected carrier beam that is lower in power but unchanged in phase. However, if the laser drifts off of resonance, then the carrier will no longer be in phase with the cavity mode, and the reflected carrier will not exactly cancel with the light leaking from the cavity. This represents a shift in phase in the reflected carrier that changes the relative amplitudes of the two beat notes with the upper and lower sidebands resulting in two beat notes that no longer completely cancel. As such, a beat note with frequency ω_{rf} is created whose amplitude is directly related to the frequency drift between the reference cavity resonance and the laser frequency. Additionally, the phase of the ω_{rf} beat note indicates the direction the frequency has drifted. By detecting the light signal off the input mirror with a photodiode and mixing it with the original modulation frequency, a DC error signal is generated which can be used to lock the laser to the optical cavity mode resulting in lower technical noise. A more complete treatment of the technique including its mathematics can be found in [86].

Appendix B

Surface Mount Soldering and Electronics

B.1 Motivation for Using Surface Mount Chips

Thanks to the consumer markets of devices such as cellular phones, GPS receivers, and hand-held PDAs, integrated chip makers have developed many high-speed, high-bandwidth, low-noise, and low-power consumption chips to fuel new technology. The characteristics of some of these chips make them ideal for use in sensitive scientific instruments. While scientific laboratories stand to benefit from the capabilities of the new technology, many of these chips are only available in surface mount packages designed for use in an automated assembly process.

Including surface mount components in circuits appears troublesome because the standard laboratory method of soldering electronic components by hand is messy and extremely time consuming for surface mount chips on a printed circuit board (PCB). Surface mount parts tend to be small with closely spaced pins, and some components have pin pads in the center of the chip that are completely inaccessible with a soldering iron. Furthermore, these chips must be soldered directly to the board without the use of sockets, which increases the risk of damage. This also increases the time, effort, and frustration that go into debugging a prototype where parts may need to be

removed or changed out several times. Purchasing industry standard equipment to perform the soldering is expensive. Table top wave-flow soldering stations start at a few thousands of dollars and are designed for mass production [122], an unnecessary capability for most laboratories.

We set out to find an easy and inexpensive way to implement surface mount technology in our designs. What we found was that lab-built electronic devices based on surface mount chips are not only possible, but in many cases preferable. Obviously, surface mount electronics are necessary when required parts are only available in surface mount packaging. However, a surface mount package's low profile and short leads help reduce parasitics, such as stray inductance and capacitance, as well as coupling to ambient noise. Surface mount boards are also preferable to through-hole designs when making more than one copy of a circuit with many components. We present useful techniques for prototyping and debugging surface mount electronics, and two inexpensive methods for reflow soldering surface mount components on a PCB, one of which is already in use among some electronics hobbyists [123].

B.2 Advantages and Disadvantages

There are several barriers to overcome when implementing surface mount chips into laboratory circuits, the most imposing of which is prototyping. Generally, electronics are developed on a breadboard where passive components and integrated circuits (ICs) can be easily interchanged without the use of a soldering iron. This makes it easy to optimize and finalize the design for the electronic circuit before building a more permanent version on a copper-clad board. In contrast, surface mount chips do not slip in and out of sockets and each version of the circuit must be soldered to a custom PCB. Additionally, surface mount pads are prone to separating from the PCB when chips are repeatedly soldered to and removed from them.

Surface mount parts are typically much smaller than leaded components, making them harder

to handle and making it more difficult to read part numbers and codes. In particular, surface mount capacitors do not have values printed on them and many of them are similar in appearance. For these reasons, if a specific surface mount component is not required and the circuit contains only a few components, a surface mount design is not desirable.

Conversely, using surface mount chips can be much more efficient when assembling larger circuits. Many chips can be placed on the board and then soldered all at one time, as opposed to soldering each individual component one at a time. Their smaller size allows for more compact circuitry and a more efficient use of space. Also, unlike through-hole components, every chip in a surface mount circuit only contacts one side of the PCB. This not only makes it easier to remove a chip, but allows a chip to be removed without removing the PCB from its housing (see section B.5).

In short, a circuit design with few components or that is a one time build probably gains little from implementing a surface mount design unless a necessary chip is only available in a surface mount package. However, circuits with many components, and especially standard designs for commonly used circuits, are perfect candidates for employing surface mount technology.

B.3 Circuit Design and Layout

There are several methods that we use to streamline the design process discussed below. We also discuss circuit design and layout techniques that have simplified the debugging and optimization processes.

B.3.1 SPICE Modeling

Because it takes significant time and cost to manufacture a PCB, we recommend characterizing the circuit with a SPICE model first. SPICE is a standard numerical modeler for circuits and several

versions are available for free from chip manufacturers. These typically come with extensive libraries of the chips made by that company, as well as the ability to load SPICE models for chips from other makers. SPICE modeling can be useful for quickly checking different combinations of chips and other components to optimize the circuit. Also, when debugging a circuit, SPICE can be used to model the effect of probes on the circuit's performance.

B.3.2 Prototyping with Leaded Components

If through-hole components are available, the design can still be prototyped on a springboard using leaded components. Prototyping in this way allows for quick adjustment of the physical circuit design before manufacturing a PCB. However, this method may not work well for highly specialized or high-speed electronics where path lengths and feedback loops are critical.

We have taken the time to design our own copper clad prototyping board with a section on it for surface mount components with various pin configurations (shown in Figure B.1). With these boards we can take advantage of sockets and leaded passive components to prototype the majority of the circuit while still using key surface mount chips.

B.3.3 Express PCB

Several resources exist for the design of the circuit, its PCB layout, and its production. ExpressPCB offers two free CAD programs to aid in board design and production. The first, ExpressSCH, is a program for drawing the circuit schematics that also checks for inconsistencies such as broken wires. This program comes with a library of common electronic components and allows for the creation of new components.

The schematic file can then be linked to a second program for design of the PCB layout, ExpressPCB, to help ensure the physical circuit is wired correctly. ExpressPCB comes with a library of pad layouts for various standard packages as well as specific chips accompanied by their Digikey

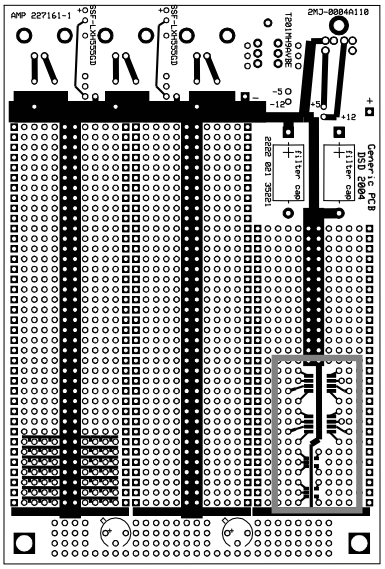

Figure B.1 Layout of one side of our prototyping board that allows for surface mount components to be integrated into the prototyping process. The surface mount section of the board is outlined by the gray box.

part number. It also allows for the creation of new components. The program offers such features as highlighting pins that should be connected together and providing warnings for missing components. It will also print the board in its actual dimensions. Unfortunately, the program does not output an industry standard Gerber file so boards can only be ordered through ExpressPCB.

B.3.4 Common Part Sizes That We Use

Many surface mount components come in a variety packages. For op-amps and other active components, small outline integrated circuit (SOIC) packages with 8 or 14 pins and 0.05 inch pin spacing (pitch) are common and are easier to deal with than packages with smaller pitch. For passive components we found 1206 and 0805 chip sizes were common and not too small to handle. We have reserved 0805 chip sizes for precision components and 1206 sizes for standard components in our lab to avoid confusion when stuffing boards.

B.3.5 Spacing Between Components

While surface mount components are small and can be packed tightly together, it is good practice to leave a generous amount of space between the pads when designing a board. This prevents confusion as to which two pads define a passive component when stuffing the board, and it also makes it easier to correct errors that occur during the soldering process (see section B.4.1). Care should also be taken to allow adequate room around through-hole components, which can overshadow or cover surface mount parts on a poorly designed board, making it harder to access them if needed.

B.3.6 Physical Design of the PCB

A well thought-out PCB design can simplify the debugging process. When designing a PCB it is generally recommended that traces run the width of the board on one side, and the length of the board on the other. This helps to keep the design simple and avoids the need to weave traces around each other with vias.

It is good practice to place a myriad of ground vias on the board, and especially around sensitive components. This will introduce many small, closely-placed ground loops, which negate large ground loops and minimize inductance between the components and ground. Also, clearing ground planes from under fast chips and around traces carrying sensitive signals will reduce problems created by stray capacitance. Checking over the board to eliminate islands (areas of metal that are not connected to either ground or a trace) is also a good idea.

There are also a few things that can be done when designing the PCB to make it easier to stuff. It helps to lay out the circuit in an intuitive manner so that different areas of the board are easily recognizable by function. Also, we label all components with either their part number or their value (for passive chips), and we make short notes or legends on a blank area of the board to aid in assembling the circuit. This can be done with a silk screening process, which adds extra cost to the production, or by simply writing such information into the copper ground plane. Also, as

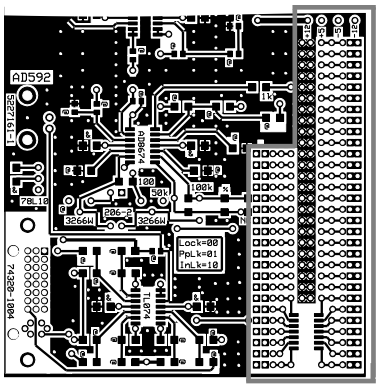

Figure B.2 The gray box encloses a prototyping grid that was added to a temperature controller circuit. This area can be used for debugging and/or adding components to the board including an SOIC chip. Power for this area can be accessed from the vias in the upper right corner of the figure, which provide ± 5 and ± 12 volts.

with standard circuit design, active components should all be placed so that pin 1 is in the same orientation for every IC.

We recommend placing as many surface mount components on the same side of the board as is possible. Only one side of the board can be soldered at once and any components on the opposite side will have to be soldered by hand. This also allows the circuit to be probed and debugged in situ as parts can be removed and added without extracting the circuit from its enclosure.

Finally, we recommend including a small prototyping area on the PCB where ground, power, and extra pin pads can support additional components that might be needed after the initial implementation. An example of one is shown in Figure B.2.

B.4 Building the Circuit

B.4.1 Layout of Surface Mount Components

When building surface mount circuitry, all of the chips on one side of the PCB can be set in place at the same time with solder paste before being placed in an oven to complete the soldering

process. Solder paste is a semi-solid mixture composed mainly of solder particles and flux. It is solid enough to hold its form, but soft enough to hold chips in place when they are pressed into it on a PCB. We found it easiest to use a 20 gauge tip on a syringe to apply the paste to individual pads. A component can then be carefully placed on the board using a pair of tweezers.

Taking care when placing components on the board can reduce errors when the solder is melted. The majority of such problems are the result of either too much solder paste or an uneven amount of solder paste being placed on the different pins of a component. Too much solder paste will form solder bridges to adjacent pads or the ground plane during the cooking process. Placing more solder one side of the chip than the other can cause chips to rotate, slide, or tombstone (lift one side of the chip off of the board).

For components with a pitch less than .05 inches, it may be easiest to place a bead of solder across all the pins. This will result in solder bridges, but solder bridges can be corrected after the cooking process (see section B.5).

Typically, it takes less than an hour to stuff a PCB containing on the order of a 100 surface mount components once a person has a little practice.

B.4.2 Through-Hole and Additional Surface Mount Components

Some circuits will contain through-hole components and may even have surface mount chips placed on both sides of the board. We layout the side of the board with the most surface mount components on it first and solder those pieces in the toaster oven. Afterward, we will add through-hole components and the remaining surface mount parts by hand.

B.4.3 Reflow Soldering

We experimented with two methods of melting the solder paste: using a common toaster oven, and using a hot plate and heat gun. We perform all of our surface mount soldering using the toaster

oven method, but also took data with the hot plate method for comparison.

Soldering in a Toaster Oven

As far as we are aware, the idea of using a common toaster oven for surface mount soldering was first put forth in a newsletter of the Seattle Robotics Society [123]. The group insisted that this method is only made possible by the use of an expensive, water-soluble solder paste from Kester that requires refrigeration and has a short shelf-life. We found that while the water-soluble paste did work, its storage and shelf life were inconvenient and its flux left a slightly sticky residue on the PCB. In searching for an alternative, we found that ChipQuik's 63/37 w/flux solder paste worked just as well. The ChipQuik paste features a "no-clean" flux that leaves no residue on the board. Additionally, it is cheap and requires no refrigeration. While the official shelf-life of the paste is 6 months without refrigeration, we were able to use an unrefrigerated syringe of the paste over a year after its original use without any problems.

According to [123], after placing a prepared board in the cold oven, the oven's temperature needs to be increased discretely in carefully timed intervals. This is done to allow the oven to come to equilibrium and, presumably, to prevent overshoot at each step. We investigated whether this was necessary or whether setting the oven to the maximum temperature setpoint from the start would work just as well.

The temperature across a blank, previously unused PCB was measured by strapping five K-type thermocouples across it. The thermocouples were connected to an Agilent 34970A Data Acquisition Switch Unit, which collected the data during the process. We then heated the board using an Oster model 6232-015 1500 Watt toaster oven.

The maximum and minimum temperature data from five points across the PCB during two trials is shown in Figure B.3. Even though the temperature was changed in steps at the intervals given in [123] for one data set, there is no evidence of it in the figure, indicating that the temperature

setpoint is changed before the oven approaches equilibrium at any step. No gradual steps are needed in cooking the board, all that is required is to get the solder to its melting point, after which the board can be cooled. The entire cooking process takes about 6.5 minutes per board with our oven.

The most temperature sensitive chip we use is an AD5541 16-bit digital to analog converter. Its maximum lead temperature of 215° C for 60 seconds [124] seems to be exceeded by the 246° C oven temperature seen when we took the data shown in Figure B.3 with no chips on the PCB. However, we have used the toaster oven to successfully create all of the electronics described in Chapter 5. We have repeated this process multiple times and have never damaged a sensitive integrated chip through the process. Given that no chips have been damaged, we can infer that lead and chip temperatures are kept below their maximum limit. This may be due to the thermal mass of the chips and the evaporation of the liquid components of the solder paste.

When soldering a PCB with this method it is important to make sure that the rack inside the oven is level. This prevents gravity from moving any of the chips off of their pads once the solder melts. The PCB with its components and solder paste is placed in the oven at room temperature and the oven is turned on and set to 450° F (232° C). During the process of cooking the board, the solder paste flattens out as the oven begins to heat. Near the 5.5 minute mark for our oven, solder will begin to melt, bead around individual pins, and become shiny. Once solder has beaded on all the pins, the oven can be turned off and the door opened. While [123] advises tapping the board immediately after cooking, we found that bumping, tipping, or jarring the board and oven immediately after the cooking process can cause chips to slide, rotate, and solder to adjacent pads, traces, or ground planes. Allowing a few seconds before actually moving the board will allow the solder to cool enough so that chips are not misaligned with their pads.

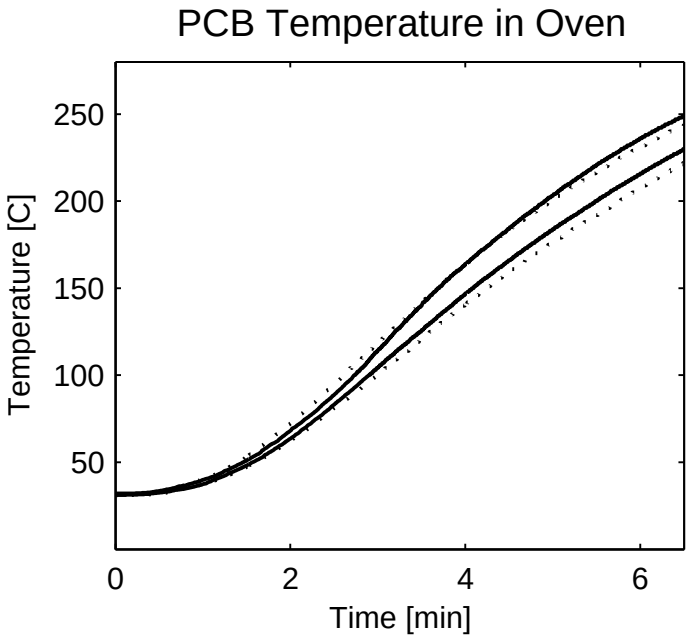

Figure B.3 Five thermocouples measured the temperature across a blank PCB while it was run through the soldering process using two different methods. The solid lines represent the maximum and minimum of the five temperature points observed on the board while the setpoint was changed in timed steps: 250° F (121° C) for 4 minutes, 375° F (191° C) for 2 minutes, and 450° F (232° C) for 30+ seconds to melt and bead the solder. The dotted lines represent the maximum and minimum profiles observed while heating the board with the oven temperature set to 450° F (232° C) from the start.

Soldering with a Hot Plate and Heat Gun

We tested a second method of doing surface mount soldering in the lab by bringing the PCB up to a uniform temperature and then locally heating the PCB to melt the solder paste. A commercial system for this approach is available and includes a heated board rack and a hot air pencil for melting the solder at individual pins [125]. We tested to see if this method could be successfully implemented using a heat gun and hot plate already available in our lab. First, we prepared the PCB as described for the toaster oven and then heated it to $\sim 150^\circ$ C on a hot plate. This is just below the 183° C melting point of the solder. By then passing a standard heat gun across the board the solder was brought to its melting point.

We took data using a blank PCB with five thermocouples as in the oven trials. Figure B.4 shows data taken by melting solder paste with a heat gun after the board had been brought to $\sim 150^\circ$ C with a hot plate. The solid trace shows local temperature spikes that occur while melting solder to a blank area on the board, while the dotted trace shows the temperature spike from actually soldering an 8 pin SOIC package with this method.

We found that when using non-specialized standard lab equipment this method is time consuming and potentially dangerous to chips. The process of getting the hot plate to the right temperature with its crude dial took about 20 minutes, while actually soldering a single chip took 10 to 20 seconds. Overall, it would take the better part of an hour to solder one board. Also, the broad nozzle of the heat gun prevented accurate targeting of chip leads, which led to overheating of the entire chip. We have never used this method for actually building surface mount circuitry.

B.5 Moving or Replacing Components

When the board has cooled, solder bridges, loose pins, tombstones, and misplaced pieces can be checked for. The use of a magnifying glass and a bright light are invaluable here. Usually, errors

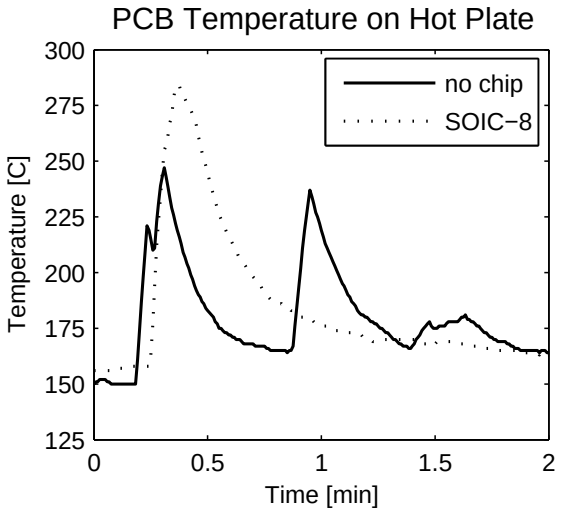

Figure B.4 After the PCB is brought to a 150° C, solder is melted by a heat gun to both the grounding plane and a trace with and without components present. The solid line shows temperatures spikes that occur from melting solder to a blank area of the board, while the dotted line shows the temperature spike due to soldering an 8 pin SOIC op-amp. Although the temperature spike for soldering the chip is short and local, the board and chip reach temperatures potentially dangerous to the chip.

such as a solder bridges can be corrected by simply touching a soldering iron to the spot and dragging the solder onto the correct traces. However, some errors require removing solder from the board or resetting a chip. Also, some chips may need to be changed out if damaged through a design error or by later abuse in the circuit's life. Components with only two leads, such as capacitors and resistors, are often easily removed by gently tugging on them with a pair of fine tweezers while heating first one end, then the other with a soldering iron. Additional solder can be added to one of the leads to increase its heat capacity so that it stays molten longer.

Removing components with more than two pins can be done through one of two methods. The first method is described in [126]. A braided copper desoldering wire or a desoldering pump is used to remove excess solder, as is commonly done on circuit boards of all designs. Then, a thin strand of enameled or nichrome wire can be threaded around the pins. Gently tugging on the wire while heating a pin will cause the pin to come off of its pad. This method is time consuming and

can easily result in a damaged pad or pin.

A better way to desolder chips is to use ChipQuik Removal Alloy [125]. This is an easy and clean way to remove most surface mount chips. The alloy mixes with the solder and creates a low melting point eutectic that will stay liquid for some time after having been heated by a soldering iron. By applying a small amount of the removal alloy to all the pins of a component, the component can simply be lifted off of the board. The alloy can then be removed from the board with a desoldering pump or wire braid. Care should be taken when using the desoldering pump since its spring action combined with heat from the iron can easily lift or rip pads and traces off of a PCB.

The quickest way that we have found to desolder chips is with the use of a solder pencil. The pencil operates by blowing a small column of hot air ($\sim 720^\circ$ F). This method requires a little bit more care than using a soldering iron to avoid damaging heat sensitive chips. However, it has the benefit of working well for chips with pads on the under side of the chip, which the other two methods cannot do. Since the pencil transfers heat to the board over a larger area than an iron it removes chips quickly without the use of ChipQuik.

B.6 Conclusion

Using a toaster oven to solder surface mount PCB electronics is fast and safe. We have produced dozens of circuits with this method and it has consistently produced beautiful circuits with no cold solder joints and no dead chips.

We have found that implementing surface mount technology in lab-built devices is not only viable, but often preferable. Prototyping for these types of circuits can take longer than for through-hole circuits. However, once an electronic device is developed, the time and cost required for assembling and building a surface mount based device are lower.

Surface mount technology opens up a whole new world of electronic solutions for scientific instruments. Using this technology in the laboratory to update existing electronic devices and develop new ones can allow for better process control, more precise measurement, higher speed, and lower noise in experiments.

Appendix C

Optical Cavity Calculations

The Octave/Matlab code used to calculate the parameters for the doubling cavity as well as output from the code is presented here. This code is specific for our cavity using four flat mirrors and a lens. While the code shown is also specific to the non-linear crystal included in our cavity as well as the lens used (125 mm), these parameters could easily be adjusted for a different cavity configuration. One important caveat is that Octave and Matlab codes are not completely compatible. The notable exception with the code shown below is the use of the function fsolve, whose options are different in each program. The code used here specifies options for Octave.

The program titled CavityLvsBeamW.m plots the size of the cavity beam waist in the arm opposite the lens as a function of the half length of the cavity. CavityLvsBeamW.m calls the function FindCavityW.m, which calculates the beam waist using the ray transfer matrix for beam propagation in the cavity and the complex parameter q defined by

$$\frac{1}{q} = \frac{1}{R} - \frac{i\lambda}{\pi\omega^2}, \tag{C.1}$$

where R is the radius of curvature of a Gaussian beam, λ is the wavelength of the light, and ω is the beam radius at any given point. FindCavityW.m in turn calls Mdist.m, Mrefract.m, and Qparam.m. Qparam.m is a function corresponding to Eq. C.1. Mdist.m and Mrefract.m are functions used to

calculate the values of the ray transfer matrix based on the given optics in the cavity.

The following section presents the graphical output of the code and the subsequent sections present the code itself.

C.1 Cavity Stability and Beam Waist

The motivation for doing these types of calculations is two-fold. First, it is important to find the configuration of optics that provides the desired beam waist for a given non-linear crystal or gain medium. Second, the cavity increasingly becomes sensitive to mechanical vibrations and technical noise if the length of the cavity necessary to produce a given beam waist is near the edge of stability for that cavity. The doubling cavity in section 4.1.3 is described by the code given here. Figure C.1 shows the beam waist as a function of the half cavity length. For cavity lengths that produce smaller beam waists, the beam waist becomes more and more sensitive to small changes in the cavity length and therefore to technical noise that affects the cavity length. Drift in the cavity beam waist affects the coupling of incident light to the cavity and therefore the intra-cavity power, and in our case the power of the frequency doubled light.

C.2 CavityLvsBeamW.m

```
clear;
close all;
%************************************************************
% Calculate parameters for doubling cavity
%************************************************************
% Declare variables
global lambda = 922*10^(-9);
global Lc = .03; % Crystal length
global nc = 1.8367; % index of refraction of crystal
global Ll = .0033; % lens length
```

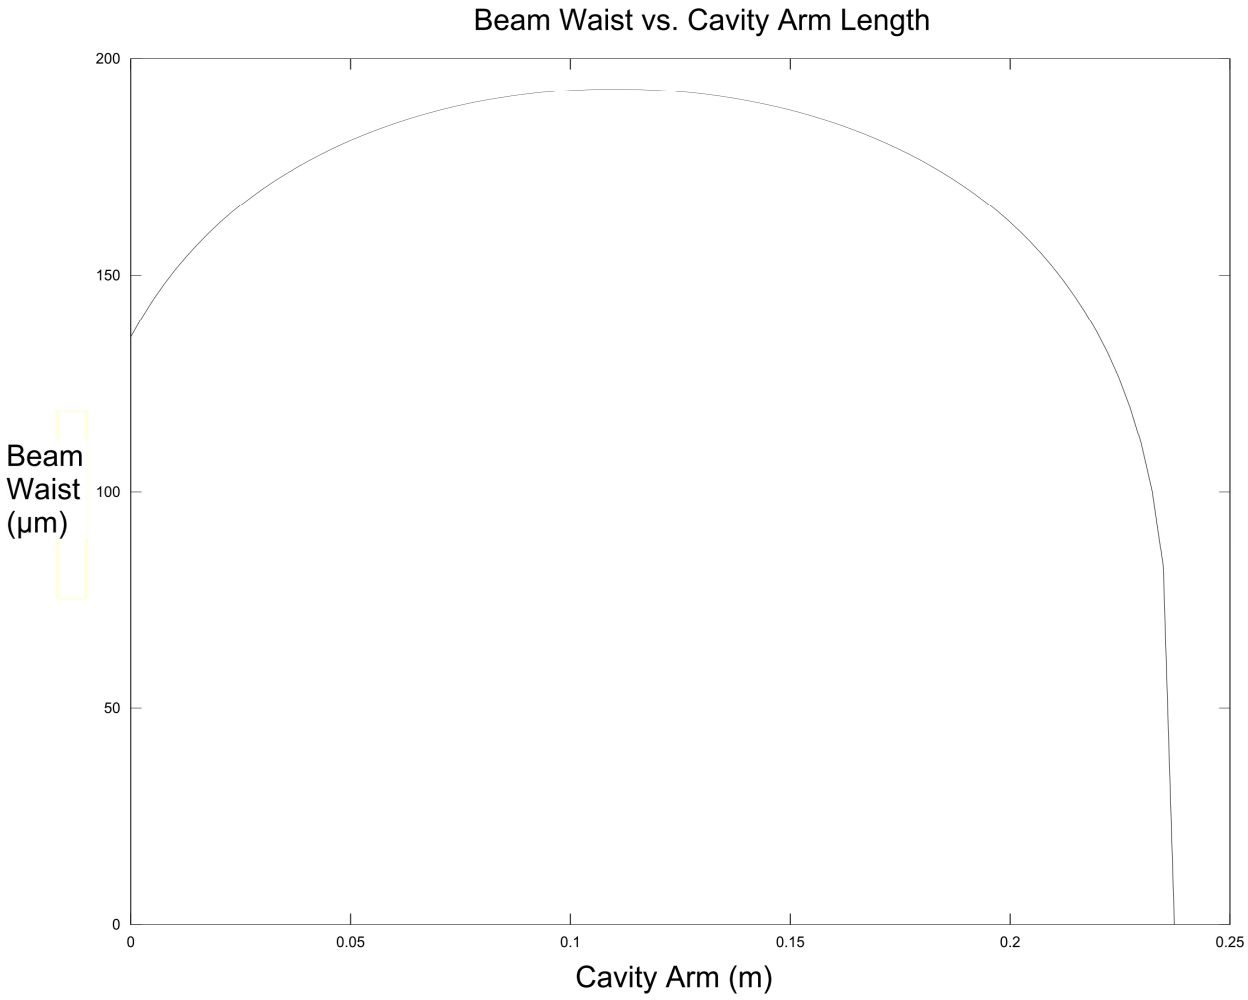

Figure C.1 The beam waist of the doubling cavity as a function of cavity length. The most stable beam waist (i.e. most insensitive to cavity length fluctuations) occurs for a cavity length of ~ 11.25 cm.

```
global nl = 1.5087; % lens index of refraction
global Rlens = .1282; % radius of curvature of side of the lens
global Tbk7 = exp(-Ll*(-.01/.9975)); %Transmission through bk-7
%                                    922nm has .9975 through 10 m)
% Enter the half cavity length you want to find the beam waist for
global z = .12-.015-.00155+.001;
% Calculate effective matrix of traversing the lens. Note that
%for the lens to be convex we must propogate so that we hit the
%positive curvature face of the lens first otherwise the lens is
%concave.
Mlens = Mrefract(nl,1,-Rlens)*Mdist(nl,Ll)*Mrefract(1,nl,Rlens);
% Calculate the first half, second half, and round trip matrices
M1 = Mdist(nl,Ll/2)*Mrefract(1,nl,Rlens)*Mdist(1,z)*Mrefract(nc,1,inf)
*Mdist(nc,Lc/2);
M2 = Mdist(nc,Lc/2)*Mrefract(1,nc,inf)*Mdist(1,z)*Mrefract(nl,1,-Rlens)
*Mdist(nl,Ll/2);
Mabcd = M2*M1;
% find the beam waist given the cavity length specified
[w0c, fval, info] = fsolve("FindCavityW", .00012);
% Print relevant values to the command window
w0c
fval
info
% Plot w0 in the crystal for various cavity lengths
figure (1)
Lmin = 0;
Lmax = .25;
wlast = .0001;
N = 100;
for a = 1:N
z = Lmin + (Lmax-Lmin)*(a-1)/(N-1);
Lplot(a) = z;
wlast = wlast;
[wc, fval, info] = fsolve("FindCavityW", wlast);
if(info == 1) wplot(a) = wc;
wlast = wc;
else
wplot(a) = 0;
end
end
plot(Lplot,wplot*1e6);
title("Beam waist vs Cavity arm length");
```

```
xlabel("Cavity Arm m");
ylabel("Beam Waist microns");
```

C.3 FindCavityW.m

```
function wmin = FindCavityW(w)
global lambda;
global Lc;
global nc;
global Ll;
global nl;
global Rlens;
global Tbk7;
global z;
% Calulate the matrix of propogating from the beam waist in the crystal
%to the beam waist in the mirror
M1 = Mdist(nl,Ll/2)*Mrefract(1,nl,Rlens)*Mdist(1,z)*Mrefract(nc,1,inf)
*Mdist(nc,Lc/2);
M2 = Mdist(nc,Lc/2)*Mrefract(1,nc,inf)*Mdist(1,z)*Mrefract(nl,1,-Rlens)
*Mdist(nl,Ll/2);
% calculate the ABCD matrix for the round trip in the crystal
Mabcd = M2*M1;
Qm = Mabcd*[Qparam(inf,lambda, nc, w);1];
Qm = Qm/Qm(2); Qf = Qm(1);
Qadj = -imag(Qf); wm = sqrt(lambda*Qadj/nc/pi);
wmin = w-wm ;
endfunction
```

C.4 Qparam.m

```
function q = Qparam(R, lambda, n, w)
q = w^2*(w^2/R + i*lambda/(n*pi))^(-1);
endfunction
```

C.5 Mdist.m

```
function M = Mdist (n,d)
M = [1 n*d; 0 1];
endfunction
```

C.6 Mrefract.m

```
function M = Mrefract(n1, n2, Rc)
M = [1 0; (n1-n2)/Rc/n2 n1/n2];
endfunction
```

Appendix D

Development of the Sr Vapor Cell

D.1 Raw Strontium Samples

Typically, a sample of strontium preserved in an argon filled glass or metal ampoule is available in small fractions of a gram. For use in the experiment we purchased a 25 g ingot of dendritic strontium stored in mineral oil from Metallium for $52 that was to provide enough strontium for both the vapor cell and the MOT/interferometer. Only recently has the same company made available a .5 g argon sealed sample for $35. While this sample is covered with a dark gray "skin", it is dry and free from the problems and impurities introduced by the mineral oil.

The strontium that comes baptized in mineral oil has a dark, almost black sheen to it. Scraping or breaking off pieces of the strontium in the oil reveals the soft silvered color of the strontium which remain shiny for days to weeks if the sample is sealed and left to sit still. The dendritic nature of the strontium is apparent from the visual appearance of the ingot and pieces of strontium break off from the ingot along parallel lines, much the way a piece of string cheese breaks apart, with veins of the dark coating reaching deep into the body of the ingot. The purity of the strontium metal itself is specified at 98%.

D.2 First Strontium Vapor Cell

The first strontium vapor cell was designed and constructed by Stuart Harper. It consisted simply of a 1.5 in. 304 stainless-steel (SS) tee with two standard 7056 glass viewports mounted at an angle to reduce etaloning, a high temperature valve on the middle port, and two stems of stainless steel threaded rod welded opposite the central port for mounting the cell on ceramic standoffs. The theory of operation was for a sample of strontium to be placed between the viewports before the valve was installed and the chamber evacuated. The cell was heated from the viewports to prevent strontium from condensing on those surfaces.

To load the strontium it was initially thought that it might be sufficient to simply dry off the strontium sample without too much preparation. With the vapor cell prepared and full of argon, the strontium was removed from the oil, wiped dry with a lint-free disposable wipe, and cut to expose fresh strontium under flowing argon. It was then placed in the cell and the cell was quickly evacuated. Even though the cell had previously been baked out, because the strontium was only simply wiped free of the oil the cell was brought up again to bake-out temperatures while still on the pump. However, as the cell heated the visible shiny surface of the Sr would begin to visibly whiten with oxidation.

This was first thought to be due to the presence of residual atmosphere in the chamber which was slow to leave through the low conductance of the small valve and tube affixed to the cell. We took careful precautions to ensure this was not the case but after each iteration we still found a slow oxidation of the strontium as the cell heated. Finally, we used a helium leak-detector to inspect the cell. While no particular part of the cell revealed a breach of the vacuum chamber when helium was blown over it locally, a consistently flowing trace amount of helium was detected entering the cell when the entire cell was placed in a helium environment. As this was done at room temperature it was easy to assume that the leak only worsened at higher temperatures.

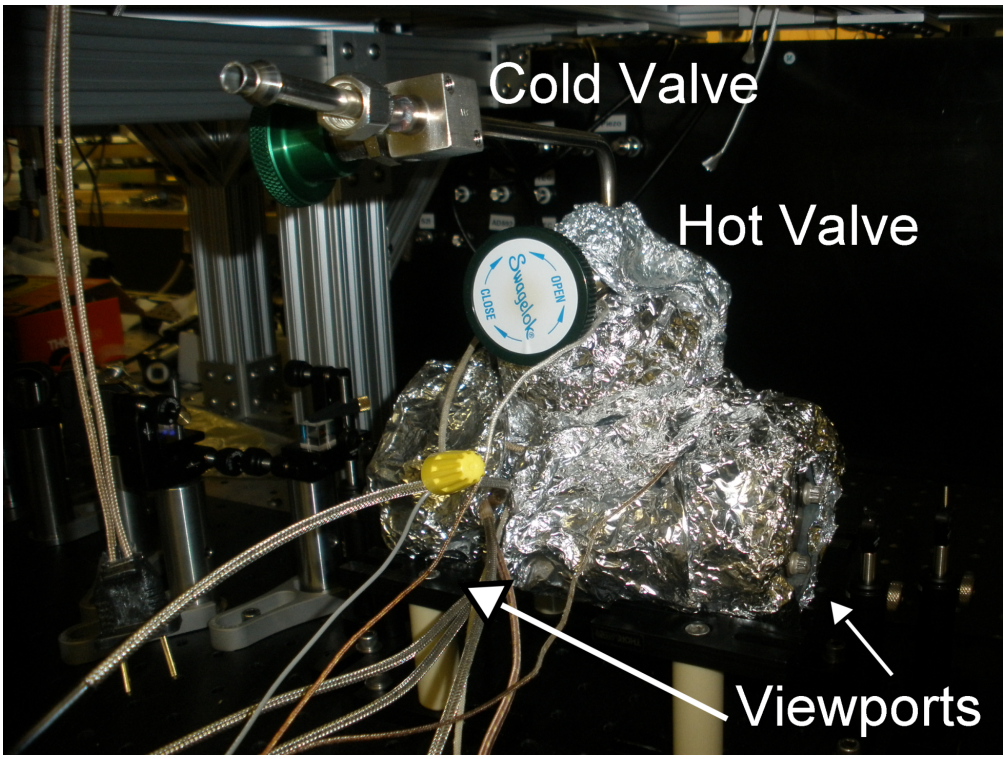

Figure D.1 Picture of the second vapor cell tried in the experiment. The vapor cell is a CF tee with standard viewports and a high temperature valve. A second cold valve is used to prevent atmosphere from corroding the hot valve.

D.3 Second Strontium Vapor Cell

Although we were unable to determine the exact cause of the leak in the first cell, it was likely due to an imperfect weld given that the CF flanges were cycled through several copper gaskets with the same effect and there was no evidence of damage to any of the knife edges on the seals. In order to move forward quickly and reduce the number of required weld joints, we constructed the next version of the cell out of a simple, off the shelf CF tee using standard 7056 CF viewports and a high temperature valve from Swagelok welded to a CF blank for loading and evacuating the cell. This vapor cell is shown in Figure D.1.

We did not witness any extra oxidation of the strontium as we heated this cell. However, even when running the cell at the maximum rated temperature for its flanges, which should have

produced a strontium vapor much denser than needed for our purposes, no fluorescence could be seen in the cell as the laser was scanned over the resonance transition.

D.4 Development of a Lab Sink Glove Box

Even though the appearance of exposed parts of the strontium was still fairly shiny after the transition to the vacuum chamber, our method for preparing the strontium seemed fairly crude. We believed it was possible that an extremely thin layer of oxide may have been forming on the surface of the exposed strontium given that the strontium was not prepared in a completely pure argon or vacuum environment.

In order to provide a better environment we were lent a homemade "glove bag" by John Ellsworth. After evacuating and refilling the bag with 99.999% argon gas several times strontium was wiped of the oil, cut, and placed in the vapor cell. However, we could not create an environment that did not witness rapid oxidation of the Sr.

To create a better environment the lab sink was converted to a make-shift glove box as seen in Figure D.2. An acrylic plate was machined to fit over the top of the sink. The plate was also machined to be fitted with a dyer hose to connect with the lab exhaust pipe, a .25 in. tube for use as a gas inlet and holes with matching cuffs to pressure fit rubber gloves into the top of the box.

The box was prepared by sealing the drain with a rubber plug and the edge of the acrylic to the sink counter-top with duct tape after the materials needed for placing the strontium in the cell were placed in the sink. After flushing the sink several times with argon a better environment was created as witnessed by the lack of reaction in the Sr.

Despite the fact that the vapor pressure of mineral oil (specified to be $< .1$ kPa at 25° C [127]) causes it to boil away rapidly in low vacuum, we also sought to reduce possible contamination of and reaction with the Sr from the oil. In an attempt to do this we began "washing" strontium

Figure D.2 Picture of the glove box fashioned from the lab sink.

samples in pure dodecane before placing them in the vapor cell.

Regardless of our efforts, no strontium vapor was witnessed in the vapor cell. At our wits end, we began contacting many of our colleagues in the field known to work with strontium. However, in each case their lab simply purchased the Sr in argon ampoules. An in-depth literature search eventually turned up work with strontium where a strontium rod was stored and prepared under argon for use in vacuum. The researchers explicitly stated that no matter the precautions taken in transferring strontium to a vacuum chamber under good argon, the strontium was reactive enough to develop an oxide layer that had to be evaporated off in vacuum at 600° C or greater [100].

This was an issue for the current cell as the maximum rated temperature for copper gasket sealed CF flanges is 450° C . To overcome this problem others have used designs that function as a jet or thermal stream of atoms [128]. This type of cell has the downside of needing to be serviced every time the source runs out. True vapor cell designs for these species have been designed where metal, glass, and ceramic parts are joined together using a high temperature glass fritz [100]. These designs can be frustrating to assemble and are fragile. We, ourselves, have previously presented a

high temperature vapor cell for spectroscopy on the calcium intercombination line which was bulky and required the use of fragile calcium fluoride windows and custom made copper gaskets [1].

D.5 Third Sr Vapor Cell

In order to evaporate off the oxide layer we modified the existing stainless steel tee with the cup described in Chapter 6. A new problem was now encountered: whether or not the viewports were kept at room or bake-out temperatures when the cup was brought up above 600° C the evaporating oxide would increase the cell pressure several decades, and it had a clear line of sight from the cup to the windows. As such the viewports quickly coated with a fine dark powder.

D.5.1 The Lead Solution

Inspiration for protecting the windows was drawn from a strontium jet design found in [128]. In this apparatus, strontium sealed in a stainless steel tube with an indium cap is placed in vacuum. After the bulk of the chamber is evacuated and baked out, heat is applied to the ampoule to melt the indium cap. The small amount of argon inside is pumped away and the ampoule can then be heated to produce the strontium vapor desired. We imagined that we might employ a similar design by placing indium foil covers in front of the windows while the the oxide layer was evaporated, after which the viewports could be heated to melt the foil and uncover the windows.

Indium foil in large enough sheets was available for $\sim$ \$100 per viewport with a significant lead time for packaging and shipping. However, lead was available on campus immediately. Lead looked like a promising substitute given that its vapor pressure is significantly lower than strontium's, it has a low melting point, and nothing in the phase diagram of possible Pb-Sr compounds indicated undesirable effects from the presence of lead at temperatures the cell would be exposed to [129].

Purified lead in short .25 in. rods was purchased from the chemical stock room on campus and rolled into foil .1 mm thick. The foil was then cleaned with acetone and alcohol and inspected for grit that may have been pressed into the foil during its production. The foil was placed in the chamber roughly inset .25 in. from the viewport's glass surface and a small opening near the edge of the foil was made to make sure the foil would not be dislodged by during evacuation by pressure trapped between it and the glass. The cup was then heated to evaporate the oxide layer until the pressure in the cell returned to pre-heating levels. Each viewport was then heated one at a time above the melting point of the lead with the cell tilted in the appropriate way for gravity to pull the foil away from the glass as it melted.

D.6 Motivation for the End Cell Design

D.6.1 Temperature Gradients

To produce the desired strontium vapor at this point, the cell was heated from four points: the two viewports, the valve, and the cup. The viewports were set at 380° C to be the hottest point in the cell in order to drive strontium off the glass before it could chemisorb into the surface. The cup and valve port were then heated to maintain the necessary vapor pressure produced by temperatures closer to $\sim$ 300° C. However, the strontium quickly condensed on the viewports.

Even though the cell was insulated well with layers of ceramic tape and aluminum foil, we found that heat loss by convection through the surface of the viewport was enough to cause a gradient of $\sim$ 200° C radially from the edge of the viewport to the center of the glass when the band heaters were brought to 380 °C. Gradients from the tee's flanges to the cup were found to be typically less than 100° C and depended on how well insulated the cell was. Clearly, the center of the glass was the coldest point in the cell even though the band heaters maintained the hottest temperature of the cell at the viewport's edge.

To reduce the temperature gradient across the face of the viewports aluminum blanks were machined to match the outer diameter of the viewports and provide a small .25 in. hole in the center for accessing the cell optically with a laser. The blank sat up against the flange bolts for each viewport and was held in place by a mica band heater for the flange, which is wide enough to enclose both the blank and the viewport. A .25 in. tube was then used to mark the position of the holes in the blanks while the entire end was insulated to ensure that most the material facing the glass would be near the outer flange temperature. After installing the aluminum blanks we found temperature gradients to the center of the viewports to be less than 50° C when the band heaters were at 380° C.

While this prevented strontium from visibly condensing on the glass, each viewport darkened beyond use quickly when a decent strontium vapor was formed in the cell. Despite our efforts to drive strontium off of the glass before it could react, the chemical bond with the glass was strong enough to capture strontium from the constant vapor. In light of this we replaced the glass viewports with sapphire ones. While strontium can also react with sapphire it does not chemisorb into the sapphire at temperatures below 600° C where it can also be driven off if it manages to condense there [100]. However, sapphire has a much lower transmittance at our wavelength than the glass and its birefringent properties slightly scramble the polarizations used in our saturated absorption setup.

D.6.2 Saturated Absorption Signal

With the sapphire viewports on the cell spectroscopy could finally be performed using the cell. Analysis of the saturated absorption signal provided in this configuration was not positive. The width of the Lamb dip was significantly broadened from the natural linewidth to $\sim$ 190 MHz and the depth was significantly less than expected. This was undesirable since the drift of the laser frequency while locked can be directly related to the width of the Lamb dip. The optimal detuning

of the cooling laser is $\Gamma/2$ and as seen in Figure 2.1 the cooling force can turn off with a change in the detuning on the order of $\Gamma/2$. Since 190 MHz is $\sim 6\Gamma/2\pi$, this Lamp dip was deemed too broad to provide a tight enough lock for the laser. We initially suspected this signal was a result of the presence of lead or perhaps the minute impurities introduced with it.

Before breaking down the cell to clean out the lead, we decided to run it as a true vapor cell. In this configuration the lamb dip depth was still small but the width had narrowed to $\sim 3\Gamma$ FWHM. This was still a very different signal than was calculated. While there was no reason to believe that the lead itself was the cause, it still contained trace amounts of impurities including zinc, cadmium, antimony, and arsenic, a few of which have vapor pressures significantly larger than strontium.

Additionally, we noticed that as temperature gradients in the cell were changed and tested, strontium could condense on and be driven off of the sapphire viewports, but a separate dark gray coating would degrade the transmission and could not be driven off by temperature. This contamination was frustrating, and prevented running the apparatus as a true vapor cell.

In an attempt to improve the laser lock and the absorption signal the cell was thoroughly cleaned. Lead and other contaminants were removed with nitric acid, baked out, and then the cell was reloaded with strontium. Surprisingly, this did not improve the broadening or size of the Lamb dip or the presence of dark gray contamination that could be deposited on the sapphire. The pressure broadening appeared to be related to the strontium itself. We speculate that the dendritic structure of the strontium has absorbed some amount of mineral oil, like a sponge. No amount of simple preparation will rid a sample of such microscopic veins of oil that might run through the metal. Oil trapped in this manner would then be released along with sublimated strontium rather than in a single bake-out, so that it is continually present in the cell.

In light of this, a second tee was added to the one with the cup for the final design presented in Chapter 6. The tees are mounted together at their central ports and the viewports were placed on the unmodified tee so that they do not have a direct line of sight to the strontium source. In this

configuration the windows can be kept cold regardless of the cup temperature to run as a thermal atom beam, or the entire cell can be heated to produce the required vapor density.

Appendix E

Vacuum Preparation

In order to achieve vacuum levels in the very high ($\sim 10^{-6} - 10^{-9}$ Torr) and ultra-high ($< 10^{-9}$ Torr) range in a reasonable amount of time it is often necessary to clean and bake-out a chamber prior to use. At these vacuum pressures the main source of gas load in the chamber is usually water desorption from the walls of the chamber. In vacuum chambers with pressures of $< 10^{-10}$ Torr the main gas load is typically hydrogen gas diffusing from the metal chamber walls.

Lower pressures become increasingly difficult to achieve in elastomer sealed systems as the permeation of atmosphere through the seal becomes the main gas load. For this reason, in systems with high vacuum requirements where elastomer seals must be used often a second low vacuum environment is maintained around the seal to limit the permeation of gas through the gasket. However, because of the number of flanges and size of the interferometer chamber, providing a low pressure environment on the outside of the flanges was not feasible.

The vacuum systems in this experiment required very high to ultra high vacuum environments in spite of the use of Viton sealed flanges. For this reason, good preparation of the vacuum environment was essential.

E.1 Cleaning with Solvents

In preparing the components for vacuum they were first washed with several solvents to rid them of dirt and oils as well as to chemically desorb as much water as possible. Most of the vacuum parts used in the experiment were cleaned thoroughly in reagent grade acetone followed by either reagent grade methanol and/or propanol.

While acetone is a stronger solvent than alcohol, the last rinse of a vacuum fitting should never been done with even very pure grades. This is because acetone will leave its own residue on the surface of vacuum parts.

Rinsing with alcohol after cleaning with acetone will remove the acetone's residue. Both methanol and propanol are also hydrophylic and lipophilic. The lipophilic property allows them to bond with oils and residue left by the acetone. The hydrophilic property causes them to bond with water molecules on the chamber surfaces and carry them away.

Almost all of the fittings and pieces that comprise our vacuum chambers were purchased from a company that specializes in vacuum technology. Typically, these parts have already been pre-cleaned and come sealed in plastic. Some pieces, such as the 8 ft. long 6 in. diameter pipe that makes up the bulk of the interferometer chamber, were purchased from a local steel manufacturer with no previous treatment. These parts were assembled and welded together in house. Prior to cleaning these parts with acetone and alcohols, they were washed with Alconox, a strong basic soap.

The turbo pumps were also cleaned in preparation for creating the best initial vacuum environment. According to the process specified by the manufacturer the pumps were cleaned first with acetone, then methanol, and finally with propanol.

Before sealing each flange on the chamber, the viton gasket used was wiped down with propanol. Only alcohols are suitable for elastomers because they are quickly attacked acids and many other solvents.

In the case of stainless steel parts that have been severely contaminated, or that have deposits of another metal, nitric acid can be used. Only stainless steel and a few other specific alloys are impervious to attack from nitric acid, and the acid should be used with care. Strong concentrations of the acid react violently with metals and organics alike and the fumes are often poisonous. If necessary, the use of nitric acid should only occur under a fume hood.

E.2 Baking Out Vacuum Parts

All of the fittings and pieces that were used on the vacuum systems were baked-out at least twice. First, in a high temperature bake-out without elastomer seals and windows, and then again after the chambers were completely assembled in a lower temperature bake-out.

E.2.1 Initial Bake-out Procedures

The majority of the parts and fittings on the experiment were initially baked out in a vacuum environment without elastomer seals. A good rule of thumb for baking out vacuum parts is to bakeout at $\sim 250°$ C for a week, where each increase in temperature of 50° C can decrease the required bakeout time by a factor of ~ 10. Before assembling the vacuum apparatus, each individual small form stainless steel, alumina, and copper part and fitting was baked out at a temperature of $\sim 350-450°$ C for ~ 24 hours.

While baking out vacuum chambers is typically done under vacuum, it can also be done in air for stainless steel parts. Typically, the bakeout in air is only done as a time saving alternative for ultra-high vacuum applications where temperatures of $\sim 400°$ C in air produce a comparable result to bakeout temperatures of $\sim 950°$ C in vacuum [130]. Baking out in air also has the benefit of creating a nitride layer on the surface of stainless steel that is less absorptive of water than the raw material.

Figure E.1 Picture of the setup for baking out the interferometer chamber in air. After the chamber was wired with heat taped and wrapped in foil, crumpled aluminum foil was used to fill the remaining space in the cavity, and long sheets of foil were laid across the top. The chamber was baked at $\sim 350°$ C for 4 days.

The only part to not initially baked out under vacuum was the interferometer chamber. In order to to achieve the same bakeout temperatures as the smaller fittings, this part was necessarily baked-out first outside of the lab with no elastomer seals or windows. An enclosure of solid cinder bricks was made around the chamber which was wrapped in high-density heat tapes capable of temperatures up to $\sim 400°$ C. The enclosure was then filled with crumpled aluminum foil and sheets of aluminum foil were laid across top. The chamber is shown being prepared for the bakeout in Figure E.1. Dry nitrogen was flowed through the chamber during the bakeout to limit the humidity of the environment. The chamber was baked out in this configuration at an average temperature of $\sim 350°$ C for 4 days.

All of the viton gaskets used were first baked out in vacuum at temperatures of $200 - 250°$ C for $4 - 5$ hours. The benefits of baking out Viton gaskets are well detailed in [131]. The gaskets were baked out en-masse, uncompressed in a 6 in. CF nipple where un-reacted monomer and other contaminants could be cold trapped and easily cleaned from the vacuum fittings after the bakeout.

E.2.2 In Situ Bake-out

Once the individual parts of all the chambers had been baked out at high temperature, they were assembled and baked out again in situ along with the ion pumps and gauges. At this point the MOT chamber was completely assembled with all of the windows mounted. The viton sealed flanges were brought up to $100-130^\circ$ C, the ion pump was brought up to $\sim 280^\circ$ C (below the Curie temperature for the magnet), and the rest of the chamber was heated to temperatures $< 350^\circ$ C. The magnet on the ion pump could not be removed to reach higher temperatures because of how the chamber was mounted. While the rule of thumb is not to bake Viton gaskets above 100° C, these gaskets can be baked out to 150° C if they are not compressed by more than 20% [132].

The interferometer chamber was also assembled completely for a second bakeout with viton sealed windows, the ion pump, RGA, titanium sublimation pump, and ion gauge installed. Band heaters were used to bring all of the 2.75 in. flanges up to $100-130^\circ$ C. Heat tapes were placed on the main chamber body and were able to achieve temperatures of $\sim 300^\circ$ C despite significant heat loss through the chamber's support structure. The magnet was removed from the ion pump and plate heaters were used to bring it to $> 400^\circ$ C, and the current was run through the sublimation pump according to the manufacturer's bakeout instructions.

As a final step the pneumatic gate valve on the MOT chamber was opened, connecting the chambers through the differential pumping section described in section 7.3. The magnet was then replaced on the large ion pump, and both chambers were baked out together.

During the bakeout, the voltage of the 2.75 in. band heaters was controlled by a bank of standard light dimmer switches to adjust the temperature. The plate heaters and heat tapes were plugged into high current variacs to control their temperature.

Figures E.2 and E.3 show pictures taken during the bakeout process.

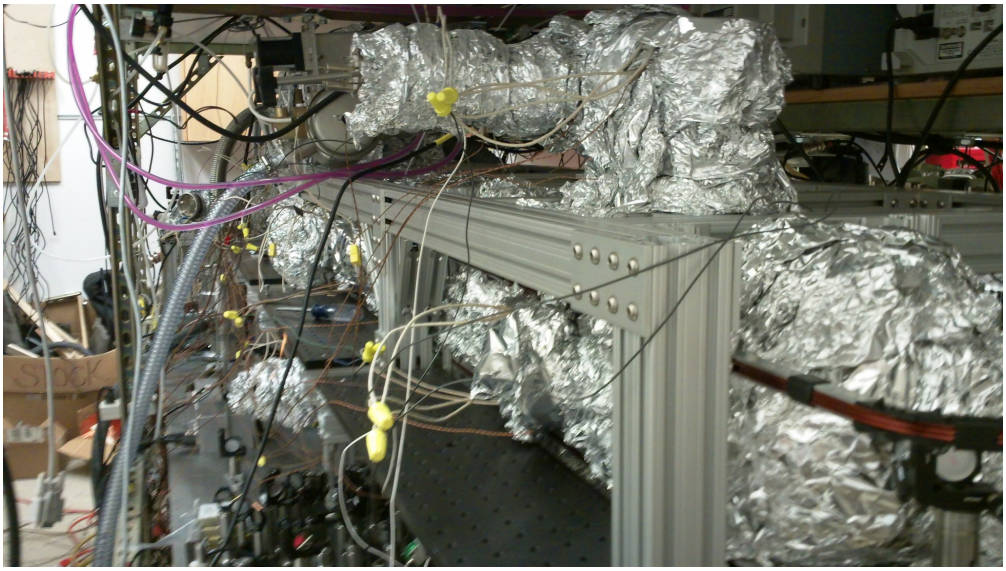

Figure E.2 Picture of the interferometer chamber wrapped and wired for heating.

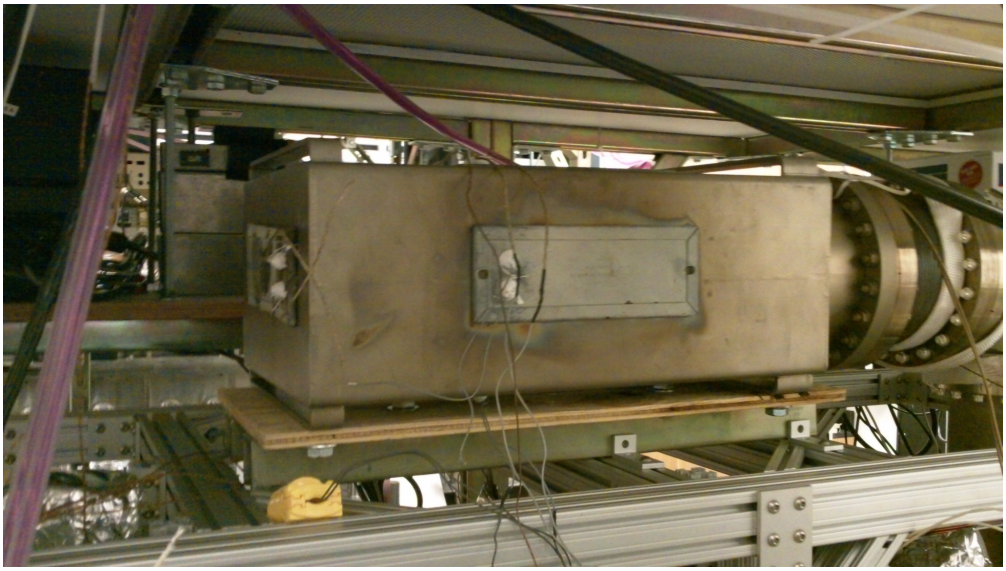

Figure E.3 Picture showing the location of the plate heaters on the ion pump. The magnet was remove from the pump in order reach temperatures above its Curie temperature.

Appendix F

PIC18F4550 Micro-controller Programming

The PIC18F4550 controller used to manage the current, temperature, and lock of the laser systems in this experiment has three types of memory: 2 kB of random access memory (RAM), 32 kB of Flash memory, and 256 B of electrically-erasable-programmable-read-only-memory (EEPROM). The chip's central processing unit (CPU) reads a short program known as a bootloader upon start up. The bootloader is stored in a protected section known as the boot block at the top of program memory. Its main purpose is to setup the chip's oscillator frequency and it contains code that allows programming the remainder of the Flash memory via USB.

After the bootloader program configures the microcontroller the chip runs whichever program is stored in the remainder of the Flash memory. The software can write to and read from the EEPROM memory, which is used to store important values and settings. The information stored in the EEPROM is preserved even when the chip is turned off or reset.

The software program was developed in MPLAB, a free programming environment from the manufacturer Microchip Technologies. The chip and development software use Microchip's ver-

sion of the C programming language, and it is necessary to have the MPLAB linker file for this particular chip (rm18f4550.lkr) in the directory when compiling the software. The main program and the functions it calls are given in the remainder of this appendix. Each section of the appendix represents a section of code with subsections for each .c and .h file for that piece of code.

F.1 Main

F.1.1 Main.c

```
/********************************************************************
*
*      Durfee Lab Microcontroller Board Software
*      Used to control Current, Temp, and PID board from one microchip
*
********************************************************************

* FileName:        main.c
* Dependencies:    See INCLUDES section below
* Processor:       PIC18F4550
* Compiler:        C18 V3
* Lab:          Durfee Group

* Software License Agreement
* Whoever uses this software assumes responsibility for all possible
* defects and consequences that may or may not arise from its intended
* or unintentional use including but not limited to:
* experimental failure including the destruction of expensive optical
* components and/or the release of smoke from electronics,
* the wrath of Durfee, loss of job, loss of self-respect, onset of
* clinical depression, increased desire for violent or unproductive,
* activities, global warming, the advent of the heat death of the
* universe, and civil liability for the breach of the terms and
* conditions of this license.

* THIS SOFTWARE IS PROVIDED IN AN AS IS CONDITION. NO WARRANTIES,
* WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED
```

```
* TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
* PARTICULAR PURPOSE APPLY TO THIS SOFTWARE. THE LAB SHALL NOT,
* IN ANY CIRCUMSTANCES, BE LIABLE FOR SPECIAL, INCIDENTAL OR
* CONSEQUENTIAL DAMAGES, FOR ANY REASON WHATSOEVER.

* Author               Date        Comment
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
* Chris Erickson       05/21/08    Original
* Dallin Durfee        06/27/08    Got analog read to work
* Chris Erickson       10/27/09    Added PID controller
* Chris Erickson       10/28/09    Organized funcs, added TOC
* Chris Erickson       11/04/09    Autolatch funtion works,
*                                  fixed LCD temp screen bug
* Dallin S. Durfee     02/07/11    Added code for chu lock,
*                                  fixed warnings, messages,
*                                  cut out dead code, etc.
* Dallin S. Durfee     04/07/11    Specified whether int resist
*                                  shorts during scan, lock/in,
*                                  lock/out
* Chris Erickson       07/05/11    modified 04/07/11 code to
*                                  short int resistor
*****************************************************************/

/*****************************************************************
TABLE OF CONTENTS
*****************************************************************
* INCLUDES
*
* VARIABLES
*
* DECLARATIONS
*
* INITIATE INTERRUPTS
*
* LOAD PARAMETER FUNCTIONS
       Load_Params_I, Load_Params_T, Load_Params_P,
       Load_Params_main
*
* CHANGE PARAMETER FUNCTIONS
       Change_Parameter_boolean, Change_Parameter,
       Change_Params_I, Change_Params_T, Change_Params_P,
       Change_Params_main
```

```
*
* CHECK PARAMETER FUNCTIONS
        Check_Params_I, Check_Params_T, Check_Params_main,
        Check_all_params
*
*  OTHER FUNCTIONS
        Make_Default
*
* MAIN()
*
* MODE FUNCTIONS
        Mode_Cleanup, Mode_SAD, Mode_PID, Mode_Temp,
        Mode_Current
*
* INITIALIZATION FUNCTIONS
        InitializeProcessor
*
* ADC FUNCTIONS
        Read_ADC, Run_ADC
*
* CHANGE VALUE RELATED FUNCTIONS
        Multi_change, Update_val, Clear_char_buffers, RtoK
*
* SCREEN UPDATE FUNCTIONS
        Clear_screen, PID_screen_update, Temp_screen_update,
        Current_screen_update
*
* TOGGLE FUNCTIONS
        Toggle, Toggle_lock
*************************************************************/

/** I N C L U D E S ***********************************************/

#include "common.h"
#include "LCD.h"
#include "DAC.h"
#include "EERW.h"
#include <adc.h>
#include "main.h"
#include <math.h>

/** V A R I A B L E S *********************************************/
```

```
//See "common.h"

/** Globals ************************************************/

int currentfunct = 0;
/// the mode that we are in (current, temp, pid, sad)
#define IRES_F     (double)ires_val*MIN_RES_STEP
#define MIN_CURRENT_STEP
((double)RS_MAX_VOLTS*MILLI)/((double)UINT_MAX)/((double)IRES_F)

#define ICALIB_F   ((double)icalib_val*MIN_CALIB_STEP)
#define TCALIB_F   ((double)tcalib_val*MIN_CALIB_STEP)
#define REF_VOLT_F (double)ref_volt*(double)MIN_REFVOLT_STEP

int encoder_turns = 0; // Stores number of dial turns on encoder
// since the last time the software got around to processing
//them

int i, check_e = 0, units = 1, m_space = 13, tog = 0, avg = 0;

double adc_val = 0;
unsigned char check_r = 0, check_s = 0;
char charbuf[40], buf[40]; //Buffers for writing to LCD
unsigned int check_d = 0, multiplier = 1000;

// Determine which functions to run
char functionI, functionT, functionP;

// Digital Board Reference Voltage
unsigned int ref_volt = 50000;

// Parameters for current driver
unsigned int ires_val = 5000, imax_val = UINT_MAX,
idac_val = 0, icalib_val = 10000, ipossup_val = 1,
ichulockStepSize=10, ichulockNumAvg=10;
char ichulock = 0, ichulockdir=1;
unsigned long ichulockavg=0, ichulocklast=0;
unsigned int ichulockN=0;

// Parameters for temperature controller
unsigned int tcalib_val = 10000, tdac_val = 40355,
```

```
tthermistor_val = 10000, tscale_val = 0, tratiomin_val=26894,
tratiomax_val=42180, tdisp_val=0;

//Parameters for PID
unsigned int pdac_val, p_int = 1, p_mode = 0;
char lock_flag = 0, auto_flag = 0;

/** D E C L A R A T I O N S **********************************/

#pragma code

/******************************************************************
* Function:        void main(void)
*
* PreCondition:    None
*
* Input:           None
*
* Output:          None
*
* Side Effects:    None
*
* Overview:        Main program entry point.
*
* Note:            None
*******************************************************************/

/** I N I T I A T E  P O R T S / P I N S ***********************

** TRIS registers determine whether a pin is an input or output,
** while LAT registers allow one to write a value to an output
** pin. A TRIS bit will set a pin to output if cleared and input
** if high. A LAT bit will clear a pin set to 0 and raise a pin
** if high.
**
** TRISC &= 0xF8 = 1111 1000 means that C0, C1, and C2 (pin 17)
** will be outputs and the rest of C will be inputs.
**
** LATC &= 0xF8 sets the output value of C0, C1 and C2 to be 0.
**
** TRISD &= 0x78 sets D0, D1, D2, D7 to be outputs and the rest
** to be inputs.
```

```
**
** LATD &= 0x78 sets the output value of D0,1,2,7 to be 0.
** D4,5,6 are the pins that control the DAC.
** D0,1 are LED's.
**
** TRISA &= 0xCF sets A4,5 to be outputs, the rest are
** inputs.
**
** LATA &= 0xCF sets A outputs to 0.
**
** TRISE &= 0xF8 sets E0,1,2 to outputs.
**
** LATE &= 0xF8 makes E output values 0.
**
** TRISB &= 0x10 sets B0,1,2,3,5,6,7 to be outputs.
**
** LATB &= 0x10 sets value of B0,1,2,3,5,6,7 to 0.
**
** ALL Pins NOT SPECIFICALLY MENTIONED are now INPUTS.
*************************************************************/

#define PORTINIT() LATC &= 0xF8; TRISC &= 0xF8; LATD &= 0x78;
TRISD &= 0x78; LATA &= 0xCF; TRISA &= 0xCF; LATE &= 0xF8;
TRISE &= 0xF8; LATB &= 0x10; TRISB &= 0x10;

/*************************************************************/
/** I N I T I A T E   I N T E R R U P T S ********************/

#define INTERRUPT_FLAG_CLEAR();
INTCONbits.TMR0IF=0; // Clear timer overflow flag bit
                        // (gets cleared at end of interrupt)

/** V E C T O R  R E M A P P I N G ***************************/

void timer_isr (void); // Declaration of timer interrupt function
extern void _startup (void); // See c018i.c in C18 compiler dir
#pragma code _RESET_INTERRUPT_VECTOR = 0x000800
//Redirects the reset interrupt address to be after
//the bootloader.
void _reset (void) //Directs processor to the code we want to run
                   //when the interrupt occurs.
{
```

```
_asm goto _startup _endasm
}

#pragma code _HIGH_INTERRUPT_VECTOR = 0x000808
void _high_ISR (void) //Interrupt code that runs when encoder
                      //dial is turned
{
_asm
goto timer_isr
_endasm
}

#pragma code _LOW_INTERRUPT_VECTOR = 0x000818
//Low priority interrupt not used.
void _low_ISR (void)
{
;
}

#pragma code
//Go back to normal compiling.
//Pragmas are assembly intsructions that
//the compiler takes into account first.

#pragma interrupt timer_isr
void timer_isr (void)
//Function that gets run when timer interrupt occurs.
{
//This particular encoder puts out 2 pulses delayed by 90
//deg. to go through a full gray code cycle with each click
//(detent). If the timer interrupt occurs we check to see if
//there is a pulse, wait until that pulse ends, and then
//check to see if the second pulse is still occuring or not.
//Depending on the last condition we increment or decrement
//"encoder_turns", then reset the timer flag bit so the
//interrupt can occur again. "encoder_turns" gets reset
//in the main code after it's change is processed.

if(ENCODE2 == 1u)
{
while(ENCODE2 == 1u){;}
if(ENCODE1 == 1u)
```

```
encoder_turns ++;
else
encoder_turns --;
}     INTERRUPT_FLAG_CLEAR();
}

/*****************************************************************
LOAD PARAMETER FUNCTIONS
*****************************************************************/

void Load_Params_I(void)
{
////////////////// Current Driver Params ///////////////////////

ires_val = EE_readUSInt(EEADDR_IR);
//Load stored integer representation of resistor value
idac_val = EE_readUSInt(EEADDR_IINIT);
//Load saved current setpoint
imax_val = EE_readUSInt(EEADDR_IMAX);
//Load value of saved current maximum
if(idac_val > imax_val)
{
idac_val = imax_val;
}
icalib_val = EE_readUSInt(EEADDR_ICALIB);
ipossup_val = EE_readUSInt(EEADDR_IPOSSUP);
ichulockStepSize = EE_readUSInt(EEADDR_ICHUSTEP);
ichulockNumAvg = EE_readUSInt(EEADDR_ICHUNAVG);
}

void Load_Params_T()
{
///////////////// Temperature Controller Params ////////////////

tscale_val = EE_readUSInt(EEADDR_TSCALE);
tdisp_val = EE_readUSInt(EEADDR_TDISP);
tthermistor_val = EE_readUSInt(EEADDR_TR);
tratiomin_val = EE_readUSInt(EEADDR_TRATIOMIN);
tratiomax_val = EE_readUSInt(EEADDR_TRATIOMAX);
tdac_val = EE_readUSInt(EEADDR_TINIT);
tcalib_val = EE_readUSInt(EEADDR_TCALIB);
}
```

```
void Load_Params_P()
{
//////////////////// PID Controller Params ////////////////////
pdac_val = EE_readUSInt(EEADDR_PINIT);
}

void Load_Params_main()
{
//////////////////// Determine which functions to run
functionI = EE_readUSInt(EEADDR_FUNCTIONI);
functionT = EE_readUSInt(EEADDR_FUNCTIONT);
functionP = EE_readUSInt(EEADDR_FUNCTIONPID);
ref_volt = EE_readUSInt(EEADDR_REFVOLT);
}

/*****************************************************************
CHANGE PARAMETER FUNCTIONS
*****************************************************************/

unsigned int Change_Parameter_boolean(const rom char* sTheString,
int EEProm_Addr, unsigned int oldval, const rom char* strYes,
const rom char* strNo)
{
LCDClear();
LCDWriteConstString("*",1,1);
LCDWriteConstString(sTheString,1,2);
tog = oldval;
if(oldval)
LCDWriteConstString(strYes,2,1);
else
LCDWriteConstString(strNo,2,1);
while(FUNCTIONSW != 0u)
{
Toggle(strYes, strNo);
//Read encoder and toggle between "yes" and "no" for the answer
}
while(FUNCTIONSW == 0u); // Wait for button to be released
EE_writeUSInt(EEProm_Addr, (unsigned int)tog);
return(tog);
}
```

```
unsigned int Change_Parameter(const rom char* sTheString, int
EEProm_Addr,unsigned int oldval, unsigned int minint, unsigned
int maxint, float scalefactor)
{
unsigned int check, val;
LCDClear();
LCDWriteConstString("*",1,1);
LCDWriteConstString(sTheString,1,2);
LCDWriteConstString("",2,m_space);

val = oldval;
check = val+1;

while(FUNCTIONSW != 0u)
{
if(encoder_turns != 0)
//Check to see if the encoder has been turned
{
val = Update_val(val,(long)minint,(long)maxint);
}
if(MULTISW == 0u)
{
Multi_change();
//times when the multi button is hit on the front panel. The loop
//runs faster than the typical time it takes to press the button.
LCDWriteConstString("",2,m_space);
}

while(MULTISW == 0u);
if(check != val)
{
check = val;
FastftoStr((double)val*(double)scalefactor,charbuf,4,
4);
LCDWriteVarString(charbuf,2,1);
sprintf(buf,(const rom far char *)"%5u", val);
LCDWriteVarString(buf,2,12);
LCDWriteConstString("",2,m_space);
  }
}
while(FUNCTIONSW == 0u);
```

```
EE_writeUSInt(EEProm_Addr, val);
return(val);
}

void Change_Params_I(void)
{
ires_val = Change_Parameter(" Rs Value (Ohm)", EEADDR_IR,
ires_val, (UINT_MIN+1), UINT_MAX, MIN_RES_STEP);
imax_val = Change_Parameter("     Max I (mA)", EEADDR_IMAX,
imax_val, UINT_MIN, UINT_MAX, MIN_CURRENT_STEP);
Change_Parameter("    I Init (mA)", EEADDR_IINIT,
EE_readUSInt(EEADDR_IINIT), UINT_MIN, imax_val,
MIN_CURRENT_STEP);
icalib_val = Change_Parameter("          I Cal",
EEADDR_ICALIB, icalib_val, UINT_MIN, UINT_MAX,
MIN_CALIB_STEP);
ipossup_val = Change_Parameter_boolean(" Type of Supply",
EEADDR_IPOSSUP, ipossup_val, "    Positive",
"    Negative");
ichulockStepSize = Change_Parameter("Chu Lck Step Sz",
EEADDR_ICHUSTEP, ichulockStepSize, UINT_MIN, UINT_MAX,
MIN_CURRENT_STEP);
ichulockNumAvg = Change_Parameter("   Chu Num Avgs",
EEADDR_ICHUNAVG, ichulockNumAvg, UINT_MIN, 9999, 1);

EE_writeUSInt(EEADDR_NOTSET_I,EE_SETFLAG);
}

void Change_Params_T(void)
{
tscale_val = Change_Parameter_boolean("     Temp Scale",
EEADDR_TSCALE, tscale_val, "    Kelvin", "    Celsius");
tdisp_val = Change_Parameter_boolean("  Display Units",
EEADDR_TDISP, tdisp_val, "    Temperature",
"    Resistance");
tthermistor_val = Change_Parameter("Thermistr(kOhm)",
EEADDR_TR, tthermistor_val, (UINT_MIN+1), UINT_MAX,
MIN_THERMISTOR_STEP);
tratiomin_val = Change_Parameter("  Min DAC Ratio",
EEADDR_TRATIOMIN, tratiomin_val, (UINT_MIN+1), UINT_MAX,
1.0/(double)UINT_MAX);
tratiomax_val = Change_Parameter("  Max DAC Ratio",
```

```
EEADDR_TRATIOMAX, tratiomax_val, (UINT_MIN+1), UINT_MAX,
1.0/(double)UINT_MAX);
tcalib_val = Change_Parameter("       Meas Cal",
EEADDR_TCALIB, tcalib_val, UINT_MIN, UINT_MAX,
MIN_CALIB_STEP);
EE_writeUSInt(EEADDR_NOTSET_T,EE_SETFLAG);
}

void Change_Params_P(void)
{
//unsigned int old_int;
//old_int = p_int;
//p_int = Change_Parameter_boolean("Intg. Fbk.
// Loop",EEADDR_PINT ,p_int,"Open","Closed");
//if(p_int != old_int)
// Toggle_lock(PID_INT);

Run_Autolatch();
EE_writeUSInt(EEADDR_NOTSET_T,EE_SETFLAG);
}

void Change_Params_main()
{
functionI = Change_Parameter_boolean(" Current Driver",
EEADDR_FUNCTIONI, functionI, "Enabled", "Disabled");
functionT = Change_Parameter_boolean("Temp Controller",
EEADDR_FUNCTIONT, functionT, "Enabled", "Disabled");
functionP = Change_Parameter_boolean(" PID Controller",
EEADDR_FUNCTIONPID, functionP, "Enabled", "Disabled");
ref_volt = Change_Parameter("Ref Voltage (V)",
EEADDR_REFVOLT, ref_volt, (UINT_MIN+1), UINT_MAX,
MIN_REFVOLT_STEP);
EE_writeUSInt(EEADDR_NOTSET_MAIN,EE_SETFLAG);
}

/*************************************************************
CHECK PARAMETER FUNCTIONS
*************************************************************/

void Check_Params_I()
{
unsigned int val;
```

```
val = EE_readUSInt(EEADDR_NOTSET_I);
if(val != EE_SETFLAG)
{
Display_Msg("I Funct not set ");
Change_Params_I();
}
}

void Check_Params_T()
{
unsigned int val;
val = EE_readUSInt(EEADDR_NOTSET_T);
if(val != EE_SETFLAG)
{
Display_Msg("T Funct not set ");
Change_Params_T();
}
}

void Check_Params_main()
{
unsigned int val;
val = EE_readUSInt(EEADDR_NOTSET_MAIN);
if(val != EE_SETFLAG)
{
Display_Msg("No funct enabled");
Change_Params_main();
}
}

void Check_all_params()
{
if(functionI)
{
Check_Params_I();
Load_Params_I();
}
if(functionT)
{
Check_Params_T();
Load_Params_T();
}
```

```
if(functionP)
{
Load_Params_P();
}
}

/*************************************************************
OTHER FUNCTIONS
*************************************************************/

char Make_Default(unsigned int addr)
{
LCDClear();
LCDWriteConstString("*   Set Default?",1,1);
tog = 0;
LCDWriteConstString("       No",2,1);

while(FUNCTIONSW != 0u)
{
Toggle("       Yes","       No");
//Read encoder and toggle between "yes" and "no" for the answer
}
while(FUNCTIONSW == 0u);
// Wait for button to be released
if(tog)
{
switch(addr)
{
case EEADDR_IINIT : EE_writeUSInt(EEADDR_IINIT,
idac_val); break;
case EEADDR_TINIT : EE_writeUSInt(EEADDR_TINIT,
tdac_val); break;
case EEADDR_PINIT   : EE_writeUSInt(EEADDR_PINIT,
pdac_val); break;
}
return(1);
}
return(0);
}

/*************************************************************/
/** MAIN
```

```
*************************************************************/

void Write_Default_Vals(void)
{
if(functionI)
DAC_write(idac_val,DAC_IDEV,ipossup_val);
//Write current to the DAC on the current controller
if(functionT)
DAC_write(tdac_val,DAC_TDEV,1);
//Write value to the DAC
if(functionP)
DAC_write(pdac_val,DAC_PDEV,1);
p_mode = 0; PID_MODESW = 0; p_int = 0; PID_INTSW=1;
}

void main(void)
{
unsigned char check_max;
//Used to check for a saved current setpoint
int tempint;
InitializeProcessor();
DAC_Init();
LCDinit();
// Clears LCD memory
Clear_char_buffers();
LCDWriteBinary(LCD_CURS_ON);
Check_Params_main();
Load_Params_main();
Check_all_params();
Write_Default_Vals();
LCDWriteConstString("-BYU Durfee Lab-",1,1);
LCDWriteConstString("uCont 2011_07_05",2,1);
// write them again every 100ms for a second in case DACs
//are slow powering up
for(tempint=0; tempint<10; tempint++)
{
DELAY_100ms();
Write_Default_Vals();
}

while(1)
{
```

```
if(MULTISW == 0u)
{
Display_Msg("Changing Functs ");
Change_Params_main();
Check_all_params();
}
else
{
while(MODESW == 0u);
switch(currentfunct)
{
case MODE_I: if(functionI) Mode_Current(); break;
case MODE_T: if(functionT) Mode_Temp(); break;
case MODE_PID: if(functionP) Mode_PID(); break;
default: Mode_SAD();
}
currentfunct++;
if(currentfunct > 3)
currentfunct = 0;
}
}
}

/*****************************************************************
MODE FUNCTIONS
*****************************************************************/

void Mode_Cleanup(void)
{
avg = 0;
adc_val = 0;
encoder_turns = 0;
}

void Mode_SAD(void)
// Situational Awareness Mode - displays all current values for
//all devices controlled
{
double current;
int sadloc;

Delay10TCYx(5);
```

```
clear_screen();
sadloc = 12;
while(MODESW != 0u)
{
LCDWriteConstString("",1,sadloc);
//Move Cursor
sadloc = sadloc + 1;
if(sadloc > 16)
{
sadloc = 12;
}

LCDWriteConstString("SAD",1,13);
LCDWriteConstString("",1,sadloc);
//Move Cursor

DoChuLock();

if(functionI)
{
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_CURRENT & ADC_INT_OFF & ADC_VREFPLUS_VDD
& ADC_VREFMINUS_VSS, 11);
  while(Run_ADC(ANALOG_CURRENT,ICALIB_F*
REF_VOLT_F*((double)I_RES_RATIO)*
((double)MILLI)/(((double)AD_MAX)*
IRES_F),0,"mA ",1)<1)
{
if(MODESW == 0u)
break;
}

if(ichulock)
{
if(ichulockdir > 0)
LCDWriteConstString(">",1,11);
else
LCDWriteConstString("<",1,11);
}
else
LCDWriteConstString("-",1,11);
```

```
}
else
{
LCDWriteConstString("I disabled",1,1);
}

LCDWriteConstString("",1,sadloc);
//Move Cursor

if(functionT)
{
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_TEMP & ADC_INT_OFF & ADC_VREFPLUS_VDD
& ADC_VREFMINUS_VSS, 11);
if(tscale_val == 0u)
while(Run_ADC(ANALOG_TEMP,TCALIB_F*(REF_VOLT_F/
(double)AD_MAX)*AD590VTOK,-273.15,"C ",2)<1)
//Update AD converter readout to main LCD screen
{
if(MODESW == 0u)
break;
}
else
while(Run_ADC(ANALOG_TEMP,TCALIB_F*(REF_VOLT_F/
(double)AD_MAX)*AD590VTOK,0,"K ",2)<1)
//Update AD converter readout to main LCD screen
  {
if(MODESW == 0u)
break;
}
}
else
{
LCDWriteConstString("T disabled",2,1);
}

LCDWriteConstString("",1,sadloc);
//Move Cursor
if(functionP)
{
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_PID & ADC_INT_OFF & ADC_VREFPLUS_VDD &
```

```
ADC_VREFMINUS_VSS, 11);
if(!p_mode)
LCDWriteConstString("~SCN~",2,12);
else
{
if(p_int == 0u)
LCDWriteConstString("LOCK ",2,12);
else
LCDWriteConstString("LOCK+",2,12);
}
}
else
{
LCDWriteConstString("NOPID",2,12);
}

LCDWriteConstString("",1,sadloc);
//Move Cursor
}
}

void Mode_PID(void)
{
clear_screen();
Delay10TCYx(5);
Mode_Cleanup();
check_d = pdac_val;
PID_screen_update();

while(MODESW != 0u)
{
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_PID & ADC_INT_OFF & ADC_VREFPLUS_VDD &
ADC_VREFMINUS_VSS, 11);
if(auto_flag)
Autolatch();

if(encoder_turns != 0)
//Check to see if the encoder has been turned
pdac_val = Update_val(pdac_val,(long)UINT_MIN,
(long)UINT_MAX);
//Update pdac_val with the appropriate change from the encoder
```

```
if(check_d != pdac_val)
//Check for change in "pdac_val"
{
check_d = pdac_val;
//Reset "check_d" to look for future change in
//"pdac_val"
DAC_write(pdac_val,DAC_PDEV,1);
//Write value to the DAC
PID_screen_update();
//Update the main LCD screen
}

Check_err_sig();

if(MULTISW == 0u)
{   Multi_change();
// runs faster than the typical time it takes to press the
//button.
while(MULTISW == 0u)
if(MODESW == 0u)
break;
}

DoChuLock();
LCDWriteConstString("",1,m_space);

if(EXTSW == 0u)
{
while(EXTSW == 0u){}
Toggle_lock(PID_MODE);
PID_screen_update();
}

if(FUNCTIONSW == 0u)
{
while(FUNCTIONSW == 0u){}
//wait until button is released before doing anything

if(!Make_Default(EEADDR_PINIT))
Change_Params_P();
```

```
PID_screen_update();
//Update the main LCD screen
while(FUNCTIONSW == 0u);
}
}
}

void Mode_Temp(void)
{
clear_screen();
Delay10TCYx(5);
Mode_Cleanup();
check_d = tdac_val;
Temp_screen_update();

while(MODESW != 0u)
{
if(encoder_turns != 0)
//Check to see if the encoder has been turned
tdac_val = Update_val(tdac_val,(long)UINT_MIN,
(long)UINT_MAX);
//Update idac_val with the appropriate change from the encoder

if(check_d != tdac_val)
//Check for change in "idac_val"
{
check_d = tdac_val;
//Reset "check_d" to look for future change in "idac_val"
DAC_write(tdac_val,DAC_TDEV,1);
//Write value to the DAC
Temp_screen_update();
//Update the main LCD screen
}

OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_TEMP & ADC_INT_OFF & ADC_VREFPLUS_VDD &
ADC_VREFMINUS_VSS, 11);

if(tscale_val == 0u)
Run_ADC(ANALOG_TEMP,TCALIB_F*(REF_VOLT_F/(double)
AD_MAX)*AD590VTOK,-273.15,"C ",2);
//Update AD converter readout to main LCD screen
```

```
else
Run_ADC(ANALOG_TEMP,TCALIB_F*(REF_VOLT_F/(double)
AD_MAX)*AD590VTOK,0,"K ",2);
//Update AD converter readout to main LCD screen

if(MULTISW == 0u)
{
Multi_change();
// runs faster than the typical time it takes to press the
//button.
while(MULTISW == 0u)
if(MODESW == 0u)
break;
  }
DoChuLock();
LCDWriteConstString("",1,m_space);

if(FUNCTIONSW == 0u)
{
//Hitting the function button enters the first menu item, which
//asks if the user wants to change the temperature scale.
while(FUNCTIONSW == 0u){}
//wait until button is released before doing anything
if(!Make_Default(EEADDR_TINIT))
Change_Params_T();
Temp_screen_update(); //Update the main LCD screen
while(FUNCTIONSW == 0u);
}
}
}

void Mode_Current(void)
{
clear_screen();
Delay10TCYx(5);
Mode_Cleanup();
check_d = idac_val;
Current_screen_update();

while(MODESW != 0u)
{
if(encoder_turns != 0)
```

```
//Check to see if the encoder has been turned
{
idac_val = Update_val(idac_val,(long)UINT_MIN,
(long)imax_val);
//Update idac_val with the appropriate change from the encoder
}
if((check_d != idac_val)&&(!ichulock))
//Check for change in "idac_val"
{
check_d = idac_val;
//Reset "check_d" to look for future change in
//"idac_val"
DAC_write(idac_val,DAC_IDEV,ipossup_val);
//Write value to the DAC
Current_screen_update();
//Update the main LCD screen
}
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
ANALOG_CURRENT & ADC_INT_OFF & ADC_VREFPLUS_VDD &
ADC_VREFMINUS_VSS, 11);
Run_ADC(ANALOG_CURRENT,ICALIB_F*REF_VOLT_F*((double)I_RES
_RATIO)*((double)MILLI)/(((double)AD_MAX)*IRES_F), 0,
"mA ",2);
//Run_ADC(ANALOG_CURRENT,10000.0*ICALIB_F*(REF_VOLT_F/
// double)AD_MAX)*((double)I_RES_RATIO)*((double)
// MILLI/IRES_F),0);
//Update AD converter readout to main LCD screen
//Run_ADC(ANALOG_CURRENT,TCALIB_F*(REF_VOLT_F/(double)
// AD_MAX)*AD590VTOK,0);
//Run_ADC(ANALOG_CURRENT,(double)(REF_VOLT_F/(double)
// AD_MAX)*(double)I_RES_RATIO*((double)
// MILLI/IRES_F)*(icalib_val*(double)MIN_CALIB_STEP),
// 0);
//Update AD converter readout to main LCD screen

if(MULTISW == 0u)
{
  Multi_change();
// runs faster than the typical time it takes to press the
//button.
while(MULTISW == 0u)
if(MODESW == 0u)
```

```
break;
  }
LCDWriteConstString("",1,m_space);
if(EXTSW == 0u)
  // engage / disengage chu lock
{
//this is where the code got that one time we had problems with
//bad connections on the board
while(EXTSW == 0u){}
Toggle_lock(I_CHU);
ichulockavg = 0;
ichulocklast = 0;
ichulockN = 0;
Current_screen_update();
}

DoChuLock();
if(FUNCTIONSW == 0u)
{
//Hitting the function button enters the first menu item, which
//asks if the user wants to change the resistor value.
while(FUNCTIONSW == 0u){}
//wait until button is released before doing anything
if(!Make_Default(EEADDR_IINIT))
Change_Params_I();
Current_screen_update();
//Update the main LCD screen
while(FUNCTIONSW == 0u);
}
}
}

//////////////////////////////////////////
//////////////  Chu Lock /////////////////
//////////////////////////////////////////

void DoChuLock(void)
{
long ichulockchange = 0;
if(ichulock)
{
OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
```

```
ANALOG_CHU & ADC_INT_OFF & ADC_VREFPLUS_VDD &
ADC_VREFMINUS_VSS, 11);
ichulockavg += Read_ADC(ANALOG_CHU);
ichulockN++;
if(ichulockN >= ichulockNumAvg)
{
if(ichulockavg > ichulocklast)
{
ichulockdir=ichulockdir*(-1);
}
ichulockchange = (long)ichulockdir*
(long)ichulockStepSize;
if((long)idac_val+(long)ichulockchange >
(long)imax_val)

idac_val = imax_val;
else if((long)idac_val+(long)ichulockchange
< (long)UINT_MIN)

idac_val = UINT_MIN;
else
idac_val += ichulockchange;
DAC_write(idac_val,DAC_IDEV,ipossup_val);
//Write value to the DAC
if(currentfunct == MODE_I)
Current_screen_update();
//Update the main LCD screen
ichulockN=0;
ichulocklast = ichulockavg;
ichulockavg = 0;
}
}
}

/*****************************************************************/
/** INITIALIZATION FUNCTIONS
*****************************************************************/

void InitializeProcessor(void)
{
INTCON = 0x20;
//Disable Global and enable TMR0 interrupts
```

```
INTCON2 = 0x84;
        //TMR0 high priority
RCONbits.IPEN=1;
//Enables priority levels on interrupts
TMR0H = 0;
//Clear timer high bit
TMR0L = 0;
//Clear timer low bit
T0CON = 0xC8;
//set up timer0 - bypass prescaler. 256 instructions or 10.67 us.
INTCONbits.GIEH=1;
//Enables high priority interrupts
    PORTINIT();
// Initializes all ports
// Set up ADC ports.
    // OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
// ADC_CH3 & ADC_INT_OFF & ADC_VREFPLUS_VDD &
// ADC_VREFMINUS_VSS, 11);
     // OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
// ADC_CH2 & ADC_INT_OFF & ADC_VREFPLUS_VDD &
// ADC_VREFMINUS_VSS, 11);
    // OpenADC(ADC_FOSC_32 & ADC_RIGHT_JUST & ADC_20_TAD,
// ADC_CH3 & ADC_INT_OFF & ADC_VREFPLUS_VDD &
// ADC_VREFMINUS_VSS, 11);
/* The 11 indicates that ports A0-A3 will be used as analog
inputs. The chart below shows what value to use
depending on how many analog ports you want to use.  For
example, if you want to use ports A0-A5, you would use
the value 9, if you want to use all of the ports as
analog input, you would use 0, and if you wanted to use
no ports as analog input, you would enter 15 as the last
argument to this function.
 0=0000 A A A A A A A A A A A A A
 1=0001 A A A A A A A A A A A A A
 2=0010 A A A A A A A A A A A A A
 3=0011 D A A A A A A A A A A A A
 4=0100 D D A A A A A A A A A A A
 5=0101 D D D A A A A A A A A A A
 6=0110 D D D D A A A A A A A A A
 7=0111 D D D D D A A A A A A A A
 8=1000 D D D D D D A A A A A A A
 9=1001 D D D D D D D A A A A A A
```

```
10=1010 D D D D D D D D A A A A A
11=1011 D D D D D D D D D A A A A
12=1100 D D D D D D D D D D A A A
13=1101 D D D D D D D D D D D A A
14=1110 D D D D D D D D D D D D A
15=1111 D D D D D D D D D D D D D
*/
}

/*****************************************************************/
/** ADC FUNCTIONS
*****************************************************************/

int Read_ADC(unsigned int channel)
{
SetChanADC(channel);
ConvertADC();
while( BusyADC() );
return(ReadADC());
}

int Run_ADC(unsigned int channel, float scalefactor, float
offset, const rom char* units, int row)
{
avg++;
adc_val += (double)Read_ADC(channel);// /(double)ANALOG_AVG;
//adc_val = Read_ADC(channel);

if(avg > ANALOG_AVG)
{
adc_val = (adc_val/(double)ANALOG_AVG)*
(double)scalefactor+offset;
// adc_val = adc_val*scalefactor;
FastftoStr((double)adc_val,buf,4,3);
// FastftoStr((double)adc_val,buf,5,1);//638
LCDWriteConstString("           ",row,1);
LCDWriteVarString(buf,row,1);
LCDWriteConstString(units,row,9);
Clear_char_buffers();
LCDWriteConstString("",1,m_space);
//Reset the cursor the right position
avg = 0;
```

```
adc_val = 0.0;
return(1);
}
return(0);
}

/*************************************************************
CHANGE VALUE RELATED FUNCTIONS
*************************************************************/

void Multi_change(void)
{
multiplier /= 10;
//Change the digit of "idac_val" that will be inc/dec when the
//encoder is turned
m_space ++;
//Change the cursor position to indicate that digit on the LCD
if(multiplier < 1u)
//Cycle through the digits by wrapping back to the "ones" place
//after the "thousands" place. This prevents changing of the
//"ten-thousands" place
multiplier = 1000;

if(m_space > 16)
//Cycle through corresponding cursor position, and wrap-around
  m_space = 13;
}

unsigned int Update_val(unsigned int val, long min, long max)
{
if((max - (long)encoder_turns*multiplier) < (long)val)
//If the dial turn would write a value above
val = max;
// max, then set the value to be the max.
else if((min - (long)encoder_turns*multiplier) > (long)val)
//If the dial turn would write below min
val = min;
// then set it to be the min.
else
//Otherwise, update appropriately
val += encoder_turns*multiplier;
```

```
encoder_turns = 0;
//Reset "encoder_turns" for accurate counting
return(val);
}

void Clear_char_buffers(void)
{
for(i=0; i < 40; i++)
//Clear buffers for writing and reading EEPROM memory
{
charbuf[i] = 0;
buf[i] = 0;
}
}

double RtoK(double R, double Ro)
// converts thermistor resistance to Kelvin
{
double temp;
temp = log(R/Ro);
temp = (3.3540170E-3)+(2.5617244E-4)*temp+(2.1400943E-6)*
temp*temp-(7.2405219E-8)*temp*temp*temp;

//return(1/temp);
return(1/temp);
}

/*************************************************************
SCREEN UPDATE FUNCTIONS
*************************************************************/

void clear_screen(void)
{
LCDWriteConstString("                ",1,1);
LCDWriteConstString("                ",2,1);
}

void PID_screen_update(void)
{
double voltage;
voltage = ((double)pdac_val/((double)UINT_MAX))*
((double)DAC_MAX);
```

```
FastftoStr(voltage,charbuf,1,2);
LCDWriteConstString("Set ",1,2);
LCDWriteVarString(charbuf,1,6);
LCDWriteConstString("V ",1,10);
sprintf(charbuf,(const rom far char *)"%5u",pdac_val);
LCDWriteVarString(charbuf,1,12);
LCDWriteConstString("Mode:",2,1);
LCDWriteConstString("Ri:",2,11);
if(p_int == 0u)
LCDWriteConstString("IN ",2,14);
else
LCDWriteConstString("OUT",2,14);
if(p_mode == 0u)
LCDWriteConstString("SCAN ",2,6);
else
LCDWriteConstString("LOCK ",2,6);
Clear_char_buffers();
}

void Current_screen_update(void)
{
double current;
current = ((double)idac_val/((double)UINT_MAX))*((double)
RS_MAX_VOLTS*MILLI)/((double)IRES_F);
FastftoStr(current,charbuf,4,3);
LCDWriteVarString(charbuf,1,1);
LCDWriteConstString("mA ",1,9);
sprintf(charbuf,(const rom far char *)"%5u",idac_val);
LCDWriteVarString(charbuf,1,12);
LCDWriteConstString("meas",2,12);
if(ichulock)
{
if(ichulockdir > 0)
LCDWriteConstString(">",2,16);
else
LCDWriteConstString("<",2,16);
}
else
LCDWriteConstString("-",2,16);
Clear_char_buffers();
}
```

```
void Temp_screen_update(void)
{
double val;
double A;
A = ((double)(tratiomax_val-tratiomin_val)*((double)
(UINT_MAX-tdac_val)/((double)UINT_MAX))+(double)
tratiomin_val)/((double)UINT_MAX);
val = (A/(1-A))*(double)(R_TEMP);
if(tdisp_val==0u)
{
LCDWriteConstString("kOh",1,9);
}
else
{
val = RtoK(val,(double)tthermistor_val*(double)
MIN_THERMISTOR_STEP);
if(tscale_val==0u)
{
val = val - ZEROC;
LCDWriteConstString("C  ",1,9);
}
else
LCDWriteConstString("K  ",1,9);
}

FastftoStr(val,charbuf,3,3);
LCDWriteConstString("       ",1,1);
LCDWriteVarString(charbuf,1,2);
sprintf(charbuf,(const rom far char *)"%5u",tdac_val);
LCDWriteVarString(charbuf,1,12);
LCDWriteConstString("meas ",2,12);
Clear_char_buffers();
}

/*****************************************************************
TOGGLE FUNCTIONS
*****************************************************************/

void Toggle(const rom char* strYes, const rom char* strNo)
{
if(check_e != encoder_turns)
{
```

```
tog = !tog;
encoder_turns = 0;
LCDWriteConstString("                ",2,1);
if(tog)
LCDWriteConstString(strYes,2,1);
else
LCDWriteConstString(strNo,2,1);
}
}

void Toggle_lock (const unsigned int sw)
{
switch(sw)
{
case PID_MODE : switch(p_mode)
// 0=scan, 1=lock w/ resistor in, 2=lock w/out resistor
case 0 : p_mode = 1; PID_MODESW = 1; p_int = 0;
PID_INTSW=1; break;
case 1 : p_mode = 2; PID_MODESW = 1; p_int = 1;
PID_INTSW=0; break;
default : p_mode = 0; PID_MODESW = 0; p_int = 0;
PID_INTSW=1; break;
// I had the logic reversed for PID_INTSW - I though that
// a zero meant that the resistor was out (switch open), but,
//in fact, sending a zero closes the switch.  I didn't want
//to change all of the places p_int shows up, so I just
//reversed the value sent to PID_INTSW
// here and where it is initialized - DSD July 5, 2011
}
break;
/* case PID_INT : if(PID_INTSW)
PID_INTSW = 0;
else
PID_INTSW = 1;
break;
*/
case I_CHU : if(ichulock)
ichulock = 0;
else
ichulock = 1;
break;
}
```

```
}

void Check_err_sig(void)
{
if(Read_ADC(ANALOG_PID) != 0)
{
if(lock_flag){}
else
{
LCDWriteConstString("!",1,1);
lock_flag = 1;
}
}
else
{
if(!lock_flag){}
else
{
LCDWriteConstString(" ",1,1);
lock_flag = 0;
}
}
}

void Run_Autolatch(void)
{
LCDClear();
LCDWriteConstString("* Run Autolatch?",1,1);
tog = 0;
auto_flag = 0;
LCDWriteConstString("       No",2,1);
while(FUNCTIONSW != 0u)
{
Toggle("       Yes","       No");
//Read encoder and toggle between "yes" and "no" for the
//answer
}
while(FUNCTIONSW == 0u);
// Wait for button to be released
if(tog)
{
auto_flag = 1;
```

```
LCDClear();
LCDWriteConstString("Engaged Scan Now",1,1);
LCDWriteConstString("  Press Funct   ",2,1);
while(FUNCTIONSW != 0u){}
}
}

void Autolatch(void)
{

LCDWriteConstString(" ",1,1);
if(p_mode)
Toggle_lock(PID_MODE);
PID_screen_update();
while(auto_flag)
{
if(encoder_turns != 0)
//Check to see if the encoder has been turned
pdac_val = Update_val(pdac_val,(long)UINT_MIN,
(long)UINT_MAX);
//Update pdac_val with the appropriate change from the
//encoder
if(check_d != pdac_val)
//Check for change in "pdac_val"
{
check_d = pdac_val;
//Reset "check_d" to look for future change in "pdac_val"
DAC_write(pdac_val,DAC_PDEV,1);
//Write value to the DAC
PID_screen_update();
//Update the main LCD screen
}

if(Read_ADC(ANALOG_PID) != 0)
{
auto_flag = 0;
Toggle_lock(PID_MODE);
}

if(MULTISW == 0u)
{
Multi_change();
```

```
// runs faster than the typical time it takes to press the
//button.
while(MULTISW == 0u)
if(MODESW == 0u)
break;
}
LCDWriteConstString("",1,m_space);
if(FUNCTIONSW == 0u)
{
while(FUNCTIONSW = 0){};
auto_flag = 0;
LCDWriteConstString("   Autolatch    ",1,1);
LCDWriteConstString("   CANCELED!!   ",2,1);
DELAY_100ms();
DELAY_100ms();
DELAY_100ms();
DELAY_100ms();
DELAY_100ms();

LCDWriteConstString(" ",1,1);
PID_screen_update();
}
}
}
```

F.1.2 Main.h

```
void InitializeProcessor(void);
// This function takes care of all of the initialization for the
// 18F4550 microcontroller
void Mode_SAD(void);
void Mode_Current(void);
void Mode_Temp(void);
void Mode_PID(void);
void DoChuLock(void);
void Toggle_lock(const unsigned int sw);
void Check_err_sig(void); void Run_Autolatch(void);
void Autolatch(void);
int Run_ADC(unsigned int channel, float scalefactor, float
```

```
offset, const rom char* units, int row);
int Read_ADC(unsigned int channel);
// Reads and returns a value from the requested analog channel
double Convert_mA(unsigned int setpoint);
// Returns the digital setpoint converted to actual milliamps
void Multi_change(void);
unsigned int Update(unsigned int val, long max, long min);
void Clear_char_buffers(void);
void Current_screen_update(void);
void Temp_screen_update(void);
void PID_screen_update(void);
void Toggle(const rom char* strYes, const rom char* strNo);
unsigned int Flip_Uint(unsigned int val);
void Set_saved_value(unsigned int val_select, unsigned int max,
unsigned int min);
void clear_screen(void);
unsigned int Update_val(unsigned int val, long min, long max);
void Change_Params_I(void);
void Change_Params_T(void);
void Change_Params_P(void);
void Change_Params_main();
void Check_Params_I();
void Check_Params_T();
void Check_Params_main();
void Check_all_params();
void Load_Params_I(void);
void Load_Params_T(void);
void Load_Params_P(void);
void Load_Params_main(void);
```

F.2 LCD

These programs are used to write to and update the LCD display.

F.2.1 LCD.c

```
#include "LCD.h"
#include "common.h"
#include <math.h>

void LCDinit(void)
{
// this is the sequence recommended by the LCD data sheet
RS = 0;
RW = 0;
DELAY_LCDINIT();
DELAY_LCDINIT();
DELAY_LCDINIT();
E=0;
DELAY_LCDINIT();
DELAY_LCDINIT();
DELAY_LCDINIT();
LCDWriteBinary(LCD_WAKEUP);
DELAY_LCDINIT();
LCDWriteBinary(LCD_WAKEUP);
DELAY_LCDINIT();
LCDWriteBinary(LCD_WAKEUP);
// these first three numbers appear to get the device started
// and set up
LCDWriteBinary(LCD_PARAM);
// Sets parameters for the actual display (8-bit, 2 row, 5x8 dot
// characters)
LCDWriteBinary(LCD_OFF);
// Turn off display and cursor
LCDWriteBinary(LCD_CLEAR);
// Clear the display
LCDWriteBinary(LCD_CURSOR_DIR);
// Set cursor dir
LCDWriteSlow(LCD_ON);
// Turn on display
}

void LCDPutDataOnLines(long iValue)
{
RS = ((0xFFFF&iValue)>>9);
RW = ((0xFFFF&iValue)>>8);
DB7 = ((0xFFFF&iValue)>>7);
DB6 = ((0xFFFF&iValue)>>6);
```

```
DB5 = ((0xFFFF&iValue)>>5);
DB4 = ((0xFFFF&iValue)>>4);
DB3 = ((0xFFFF&iValue)>>3);
DB2 = ((0xFFFF&iValue)>>2);
DB1 = ((0xFFFF&iValue)>>1);
DB0 = (0xFFFF&iValue);
}

void LCDWriteSlow(long iValue)
{
LCDPutDataOnLines(iValue);
DELAY_LCDON();
E = 1;
DELAY_LCDON();
E = 0;
DELAY_LCDINSTRUCT();
}

void LCDWriteBinary(long iValue)
{
LCDPutDataOnLines(iValue);
DELAY_LCDE();
E = 1;
DELAY_LCDE();
E = 0;
DELAY_LCDINSTRUCT();
}

void LCDWriteConstString(const rom char* sTheString, int row,
int col)
{
int charpos;
int ioffset = 0;
if(row == 1)
charpos = col-1 + LCD_ROW1START;
else
charpos = col-1 + LCD_ROW2START;
LCDWriteBinary(charpos+LCD_SETPOS_OFFSET);
LCDWriteBinary(LCD_CURSOR_FWD);
while(*(sTheString+ioffset) > 0)
{
//LCDWriteBinary(*sTheString);
```

```
LCDWriteBinary(LCD_WRITECHAR_OFFSET+*(sTheString+
ioffset));
ioffset++;
}
}

void LCDWriteVarString(char sTheString[], int row, int col)
{
int charpos;
int ioffset = 0;
if(row == 1)
charpos = col-1 + LCD_ROW1START;
else
charpos = col-1 + LCD_ROW2START;
LCDWriteBinary(charpos+LCD_SETPOS_OFFSET);
LCDWriteBinary(LCD_CURSOR_FWD);
while(*(sTheString+ioffset) > 0)
{
//LCDWriteBinary(*sTheString);
LCDWriteBinary(LCD_WRITECHAR_OFFSET+*(sTheString+
ioffset));
ioffset++;
}
// LCDWriteBinary(LCD_WRITECHAR_OFFSET+sTheString[0]);
}

#define SPACESFORINT 3
#define SPACESFORFRAC 4
#define ROUNDHALF 0.00005 // = 0.5 / 10^SPACESFORFRAC
#define FRACMULTIPLIER 10000 // = 10^SPACESFORFRAC

int findleadingspaces(unsigned int thevalue, int spaces)
{
unsigned int spacemult = 1;
int i,extraspaces;
for(i=0; i < spaces-1; i++)
{
spacemult *= 10;
}
extraspaces = 0;

for(i = 0; i < spaces-1; i++)
```

```
{
//if( ((float)thevalue/(float)spacemult) < 1.0)
if( thevalue/spacemult < 1.0)
{
extraspaces++;
spacemult /= 10;
}
else
{
return(extraspaces);
}
}
return(extraspaces);
}

char * FastftoStr(double f, char * str, int spacesbefore, int
spacesafter)
{
unsigned int intpart;
unsigned int fracpart;
int negative = 0;
//float roundhalf = 0.5;
float fracmult = 1;
unsigned int spacemult=1;
int i,extraspaces, stringlocation;
char formatstring[10];
if (f<0)
{
f=-f;
negative = 1;
}

for(i=0; i < spacesafter; i++)
{
fracmult *= 10;
}
f = f+0.5/fracmult;
// required for rounding
intpart = (unsigned int)floor(f);
fracpart = (unsigned int)floor((f - intpart)*fracmult);
extraspaces = findleadingspaces(intpart, spacesbefore-
negative);
```

```
for(stringlocation=0; stringlocation < extraspaces;
stringlocation++)

str[stringlocation] = ' ';

if(negative)
{
str[stringlocation] = '-';
stringlocation++;
}
itoa(intpart, &str[stringlocation]);
str[spacesbefore] = '.';
stringlocation = spacesbefore + 1;
extraspaces = findleadingspaces(fracpart, spacesafter);
  for(i = 0; i < extraspaces; i++)
str[stringlocation+i] = '0';
itoa(fracpart, &str[stringlocation+i]);
//str[spacesbefore+1+spacesafter] = 0;
return(str);
}

char * EEftoStr(double f, char * str)
{
// 3 digits, a decimal, and then 4 more
unsigned int intpart;
unsigned int fracpart;
unsigned int negative = 1;
if (f<0)
{
f=-f;
negative = -1;
}
f = f+ROUNDHALF;
// required for rounding
intpart = (unsigned int)floor(f);
fracpart = (unsigned int)floor((f - intpart)*FRACMULTIPLIER);

if((negative<0u)&(intpart==0u))
sprintf(str, (const rom far char *)"  -0.%04u",
fracpart);
// "  -0.%0SPACESFORFRACd"
else
```

```
sprintf(str, (const rom far char *)"%3u.%04u",
(intpart*negative), fracpart);
//
// "%SPACESFORINTd.%0SPACESFORFRACd"
return(str);
}

void LCDClear()
{
LCDWriteConstString("                ",1,1);
LCDWriteConstString("                ",2,1);
}

void Display_Msg(const rom char* sTheString)
{
LCDClear();
LCDWriteConstString(sTheString,1,1);
LCDWriteConstString("  Press Funct   ",2,1);
while(FUNCTIONSW != 0u);
while(FUNCTIONSW == 0u);
}
```

F.2.2 LCD.h

```
void LCDinit(void);
// Initiallizes LCD (see datasheet)
void LCDPutDataOnLines(long iValue);
void LCDWriteBinary(long iValue);
// Write one binary word to the LCD
void LCDWriteConstString(const rom char* sTheString, int row,
int col);
void LCDWriteVarString(char sTheString[], int row, int col);
void LCDWriteSlow(long iValue); void LCDClear(void);
void Display_Msg(const rom char* sTheString);
char * FastftoStr(double f, char * str, int spacesbefore, int
spacesafter);
#define LCD_CLEAR 0x001
// word to send to clear the display
#define LCD_WRITECHAR_OFFSET  0x200
```

```
// add this to ASCII char code and send to LCD to write char at
// cursor position
#define LCD_SETPOS_OFFSET 0x80
// add this to position to get word to set cursor position
#define LCD_CURSOR_FWD 0x006
// set cursor to move forwards
#define LCD_CURSOR_BAK 0x004
// set cursor to move backwards
#define LCD_WAKEUP 0x030
// call 3 times to "wake up" unit
#define LCD_PARAM 0x038
// word to set parameters for the actual display  (8-bit, 2 row,
// 5x8 dot characters)
#define LCD_OFF 0x008
// word to turn off cursor and display
#define LCD_CURSOR_DIR 0x006
// word to set cursor dir
#define LCD_ON 0x00C
// word to turn on display
#define LCD_CURS_ON 0x00E
// word to turn on cursor
#define LCD_CURS_OFF 0x00C
// word to turn cursor back off
#define LCD_CHARSINROW 16
#define LCD_ROW1START 0x00
// first char position on row 1
#define LCD_ROW2START 0x40
// first char position on row 2
#define DB7 LATAbits.LATA4
// db7 to db0 are the data bits for the LCD display
#define DB6 LATAbits.LATA5
#define DB5 LATEbits.LATE0
#define DB4 LATEbits.LATE1
#define DB3 LATEbits.LATE2
#define DB2 LATCbits.LATC0
#define DB1 LATCbits.LATC1
#define DB0 LATCbits.LATC2
#define E    LATDbits.LATD0
// e is clock to send data to LCD (after setting db7-db0 and rw
// and rs, set this to 1 for at least 230 ns, and then set back
// to 0
#define RW   LATDbits.LATD1
```

```
// Read/write bit for LCD
#define RS  LATDbits.LATD2
// rs bit for LCD
```

F.3 DAC

These files are called for programming the AD5541 DAC on each circuit.

F.3.1 DAC.c

```
#include "common.h"
#include "DAC.h"

void DAC_Init()
{
DAC_CS_I = 1;
}

void DAC_clock()
// clock cycle on DAC
{
Delay10TCYx(10);
  DAC_SCLK = 0;
Delay10TCYx(10);
  DAC_SCLK = 1;
Delay10TCYx(10);
}

void DAC_write(unsigned int value, char device, char invert)
{
int x = 0;
// long value = val;
if(invert)
value = UINT_MAX - value;
DAC_SCLK = 1;
Delay10TCYx(5);
```

```
switch(device)
{
case DAC_IDEV : DAC_CS_I = 0; break;
case DAC_TDEV : DAC_CS_T = 0; break;
case DAC_PDEV : DAC_CS_P = 0; break;
}
for(x=15; x>-1; x--)
// reads serially values into DAC
{
DAC_DIN = ((value>>x) & 0x01);
  DAC_clock();
}
switch(device)
{
case DAC_IDEV : DAC_CS_I = 1; break;
case DAC_TDEV : DAC_CS_T = 1; break;
case DAC_PDEV   : DAC_CS_P = 1; break;
}
DAC_clock();
DAC_SCLK = 0;
}
```

F.3.2 DAC.h

```
void DAC_Init(void);
void DAC_clock(void); // clock cycle on DAC
void DAC_write(unsigned int value, char device, char invert);
// write value to DAC
#define DAC_CS_I              LATBbits.LATB1
#define DAC_CS_T  LATBbits.LATB0
#define DAC_CS_P  LATBbits.LATB2
#define DAC_SCLK              LATBbits.LATB3
#define DAC_DIN               LATBbits.LATB6
#define DAC_IDEV 0
#define DAC_TDEV 1
#define DAC_PDEV 2
//#define DAC_MID 32767
//NOTE: DAC_CS_T used to be B2, but is now B0
```

F.4 EERW

The code in this section is called for reading and writing to the PIC's EEPROM.

F.4.1 EERW.c

```
#include "common.h"
#include "EERW.h"
#include "LCD.h"

/*****************************************************************
The functions WriteEEPROM and ReadEEPROM are set up write and read
character strings from the EEPROM memory. They are not completely
debugged and checked out, but should work with minimal tampering.
They are not called in the demo.
*****************************************************************/

void WriteEEPROM(unsigned int addr, unsigned char length,
unsigned char *source)
{
  unsigned char *ptr;
  unsigned char loop;
  ptr = source;
  INTCONbits.GIEL = 0;
  EEADR = (unsigned char)(addr & 0x00FF);
  loop = length - 1;
  EECON1bits.EEPGD = 0;
  EECON1bits.CFGS  = 0;
  EECON1bits.WREN  = 1;
  do
  {
    INTCONbits.GIEL = 0;
    EEDATA = *ptr;
EECON2 = 0x55;
EECON2 = 0xAA;
    EECON1bits.WR = 1;
    INTCONbits.GIEL = 1;
    while(EECON1bits.WR == 1u);
    ptr++;
```

```
    EEADR++;
  }
  while (loop--);
  EECON1bits.WREN  = 0;
}

void ReadEEPROM(unsigned int addr, unsigned char length,
unsigned char *dest)
{
  unsigned char *ptr;
  unsigned char loop;
  ptr = dest;
     INTCONbits.GIEL = 0;
  EEADR = (unsigned char)(addr & 0x00FF);
  loop = length - 1;
  EECON1bits.EEPGD = 0;
  EECON1bits.CFGS  = 0;
  do
  {
    EECON1bits.RD = 1;
    *ptr = EEDATA;
    ptr++;
    EEADR++;
  }
  while (loop--);
  INTCONbits.GIEL = 0;
}

/*****************************************************************
The functions EE_write and EE_read write and read a single
character to the EEPROM memory. They take a pointer to the
character and an address.
*****************************************************************/

void EE_write(unsigned char address, unsigned char *_data)
{
    EEDATA = *_data;
//Store value at "data" into the EEPROM register that will be
//written
    EEADR = address;
//Start write sequence as described in datasheet, page 91
    EECON1bits.EEPGD = 0;
```

```
//Select access to data EEPROM memory
    EECON1bits.CFGS = 0;
//Access data EEPROM memory
    EECON1bits.WREN = 1;
//Enable writes to data EEPROM
    INTCONbits.GIE = 0;
  //Disable interrupts
    EECON2 = 0x55;
//I have no idea what we are doing with the EECON2
EECON2 = 0x0AA;
// register. There is no description of it available
EECON1bits.WR = 1;
//Start writing
    while(EECON1bits.WR)
{
//Make sure nothing interrupts the chip while saving to the mem
     _asm nop _endasm;
}
    if(EECON1bits.WRERR)
{
//Send a message if the write faile for any reason
    LCDWriteConstString(" WRITE FAILED!! ",2,1);
}
    EECON1bits.WREN = 0;
//Disable writes to data EEPROM
    INTCONbits.GIE = 1;
//Enable interrupts
}

unsigned char * EE_read(unsigned char address, unsigned char
*_data)
{
    EEADR = address;
//Start read sequence as described in datasheet
     EECON1bits.CFGS = 0;
//Access data EEPROM memory
    EECON1bits.EEPGD = 0;
//Select access to the data EEPROM memory
    EECON1bits.RD = 1;
//Start read process
    *_data = EEDATA;
//Store value at the address "_data"
```

```
return(_data);
//Return a pointer to the character read
}

/*****************************************************************
EE_writeUSInt and EE_readUSInt write and read 16 bit unsigned
integers to the EEPROM memory using the functions EE_write and
EE_read.
*****************************************************************/

void EE_writeUSInt(unsigned char address, unsigned int data)
{
int i;
int size;
unsigned char ch;
size = sizeof(int);
//No magic #'s, make sure the size is right
for(i=0; i < size; i++)
//Take an integer one byte at a time and store it
// in the data EEPROM memory
{
ch = (data >> i*8)&0x00FF;
EE_write(address+i, &ch);
}
}

unsigned int EE_readUSInt(unsigned char address)
{
int i;
int size;
unsigned char ch;
unsigned int result;
size = sizeof(int);
result = 0;
for(i=0; i < size; i++)
//Starting at the specified address where a
// known integer was stored, read back the
// integer one byte at a time and reassemble it
{
EE_read(address+i, &ch);
result +=  ((int)ch) << 8*i;
}
```

```
return(result);
//Return the integer
}
```

F.4.2 EERW.h

```
void EE_write(unsigned char address, unsigned char *_data);
unsigned char * EE_read(unsigned char address, unsigned char
*_data);
void EE_write_hex(unsigned char address, unsigned char data);
unsigned char * EE_read_true(unsigned char address, unsigned
char data);
void WriteEEPROM(unsigned int addr, unsigned char length,
unsigned char *source);
unsigned int EE_readUSInt(unsigned char address);
void EE_writeUSInt(unsigned char address, unsigned int data);
```

Appendix G

Curriculum Vitae

The author was born in Salt Lake City, Utah on March 2, 1980. He attended Brigham Young University from 1998 to 2005, and graduated with a Bachelor of Science degree in Physics and a Bachelor of Arts degree in Music in 2005. He served a mission for the Church of Jesus Christ of Latter-day Saints in Buenos Aires, Argentina from 1999 to 2001. He continued at Brigham Young University in the Fall of 2005 and began graduate studies in Physics. He received a NASA Space Grant Fellowship in 2006, a SMART Fellowship in 2006 and 2007, and a BYU Graduate Studies Fellowship in 2009 and 2010. He pursued his research in Physics under the direction of Professor Dallin Durfee and received the Master of Science degree from Brigham Young University in 2007.

List of Publications and Articles Submitted for Publication:

D. Troxel, C. J. Erickson, D. S. Durfee, "Updates to an ultra-low noise current driver," Rev. Sci. Instrum., 82, 96101 (2011).

N. D. Zameroski, G. D. Hager, W. Rudolph, C. J. Erickson, and D. A. Hostutler, "Pressure broadening and collisional shift of the Rb D2 absorption line by CH4, C2H6, C3H8, n-C4H10 and He," Journal of Quantitative Spectroscopy \& Radiative Transfer, 112, 59 (2011).

C. J. Erickson, J.L. Archibald, and D.S. Durfee, "An ultra-stable 657 nm diode laser," The Journal of the Utah Academy of Sciences, Arts, and Letters, 89, 257 (2009).

C.J. Erickson, and D.S. Durfee, "Reflow soldering of surface mount electronic components in a laboratory," arXiv:0901.0136v1 (2008).

C. J. Erickson, M. Van Zijll, Greg Doermann, and D.S. Durfee, "An ultrahigh stability, low-noise laser current driver with digital control," Rev. Sci. Instrum., 79, 073107 (2008).

C. J. Erickson, D. Christensen, M. Van Zijll, M. Washburn, J.L. Archibald, D.S. Durfee, "Progress toward an ultra-high precision thermal beam interferometer," Proceedings of the Annual Fellowship Symposium of the Rocky Mountain Space Grant Consortium (2007).

C. J. Erickson, B. Neyenhuis, and D. S. Durfee, "A high temperature calcium vapor cell for spectroscopy on the 4s^2 ^{1}S$_0$ to 4s4p ^{3}P$_1$ intercombination line," Rev. Sci. Instrum., 76, 123110 (2005).

Bibliography

[1] C. J. Erickson, B. Neyenhuis, and D. S. Durfee, "A high-temperature calcium vapor cell for spectroscopy on the $4s^2\ {}^1S_0$ to $4s4p\ {}^3P_1$ intercombination line," Rev. Sci. Inst. **76,** 123110 (2005).

[2] C. J. Erickson, Senior Thesis, Brigham Young University, 2005.

[3] C. J. Erickson, M. V. Zjill, G. Doermann, and D. S. Durfee, "An ultrahigh stability, low-noise laser current driver with digital control," Rev. Sci. Inst. **79,** 1 (2008).

[4] D. L. Troxel, C. J. Erickson, and D. S. Durfee, "Updates to an ultra-low noise laser current driver," Rev. Sci. Inst. **82,** 96101 (2011).

[5] B. Neyenhuis, D. Christensen, and D. S. Durfee, "Testing non-classical theories of electromagnetism with ion interferometry," Phys. Rev. Lett. **99,** 200401 (2007).

[6] D. Christensen, B. Neyenhuis, and D. S. Durfee, "Numerical calculations of classical and non-classical electrostatic potentials," ArXiv:physics/0609128v1 (2006).

[7] C. J. Erickson, Master's thesis, Brigham Young University, 2007.

[8] A. A. Michelson and E. W. Morley, "On the relative motion of the Earth and the luminiferous ether," Am. J. of Sci. **34,** 22 (1887).

[9] C. F. McMillan, D. R. Goosman, N. L. Parker, L. L. Steinmetz, H. H. Chau, T. Huen, R. K. Whipkey, and S. J. Perry, "Velocimetry of fast surfaces using Fabry–Perot interferometry," Rev. Sci. Inst. **59,** 1–21 (1988).

[10] L. M. Barker and R. E. Hollenbach, "Laser interferometer for measuring high velocities of any reflecting surface," J. App. Phys. **43,** 4669–4675 (1972).

[11] Northrop Grumman, "MK 39 ring laser gyro: inertial navigation system," online, http://www.sperrymarine.northropgrumman.com/products/Inertial_Navigation/mk39, 2011.

[12] T. Suzuki, O. Sasaki, and T. Maruyama, "Phase locked laser diode interferometry for surface profile measurement," Appl. Opt. **28,** 4407–4410 (1989).

[13] C. Tiena, C. Lee, and C. Jaing, "The measurement of thin film stress using phase shifting interferometry," J. Mod. Opt. **47,** 839–849 (2000).

[14] B. Abbott *et al.*, "Detector description and performance for the first coincidence observations between LIGO and GEO," Nuc. Instrum. and Meth. in Phys. Res. A **517,** 154–179 (2004).

[15] T. L. Gustavson, A. Landragin, and M. A. Kasevich, "Rotation sensing with a dual atom-interferometer Sagnac gyroscope," Class. Quantum Grav. **17,** 2385 (2000).

[16] A. Peters, K. Y. Chung, and S. Chu, "High-precision gravity measurements using atom interferometry," Metrologia **38,** 25 (2001).

[17] S. Gupta, K. Dieckmann, Z. Hadzibabic, and D. E. Pritchard, "Contrast interferometry using Bose-Einstein condensates to measure h/m and α," Phys. Rev. Lett. **89,** 140401 (2002).

[18] V. A. Dzuba and V. V. Flambaum, "Relativistic effects in Sr, Dy, YbII, and YbIII and search for variation of the fine-structure constant," Phys. Rev. A **68,** 022506 (2003).

[19] J. R. Torgerson, "Method for precision test of fine structure constant variation with optical frequency references," arXiv:physics/0012054v3 (2001).

[20] J. D. Prestage, R. J. Tjoelker, and L. Maleki, "Atomic clocks and variations of the fine structure constant," Phys. Rev. Lett. **74,** 3511 (1995).

[21] H. Marion *et al.*, "Search for variations of fundamental constants using atomic fountain clocks," Phys. Rev. Lett. **90,** 150801 (2003).

[22] E. Peik, B. Lipphardt, H. Schnatz, T. Schneider, and S. G. Karshenboim, "Limit on the present temporal variation of the fine structure constant," Phys. Rev. Lett. **93,** 1708011 (2004).

[23] M. Fischer *et al.*, "New limits on the drift of fundamental constants from laboratory measurements," Phys. Rev. Lett. **92,** 2308021 (2004).

[24] E. Peik, B. Lipphardt, H. Schnatz, T. Schneider, C. Tamm, and S. G. Karshenboim, "Frequency comparisons and absolute frequency measurements of $^{171}Yb^{+}$ single-ion optical frequency standards," ArXiv:physics/0504101v1 (2005).

[25] R. M. Mansouri and R. U. Sexl, "A test of special relativity," Gen. Relativ. Gravit. **8,** 497 (1977).

[26] H. P. Robertson, "Postulate versus observation in the special theory of relativity," Rev. Mod. Phys. **21,** 378 (1949).

[27] R. J. Kennedy and E. M. Thorndike, "Experimental establishment of the relativity of time," Phys. Rev. **42,** 400 (1932).

[28] C. Braxmaier, H. Müller, O. Pradl, J. Mlynek, and A. Peters, "Tests of relativity using a cryogenic optical resonator," Phys. Rev. Lett. **88,** 010401 (2002).

[29] R. L. Walsworth, D. Bear, M. Humphrey, E. M. Mattison, D. F. Phillips, R. E. Stoner, and R. F. C. Vessot, "New clock comparison searches for Lorentz and CPT violation," ArXiv:physics/0007063 (2000).

[30] M. A. Humphrey, D. F. Phillips, E. M. Mattison, R. F. C. Vessot, R. E. Stoner, and R. L. Walsworth, "Testing CPT and Lorentz symmetry with hydrogen masers," Phys. Rev. A **68,** 063807 (2003).

[31] G. F. Missiroli, G. Pozzi, and U. Valdrè, "Electron interferometry and interference electron microscopy," J. Phys. E **14,** 649 (1981).

[32] H. Lichte, "Gottfried Möllenstedt and his electron biprism: four decades of challenging and exciting electron physics," J. Electron Microsc. **47,** 387 (1998).

[33] An ion interferometer has been demonstrated in the lab of Franz Hasselbach, but I have been unable to find any publications describing their work. Non-reviewed public presentation of the device can be found online at http://www.physik.unituebingen.de/hasselbach/ii/ii.htm and http://video.google.com/videoplay?docid=-7562244914138140 10.

[34] L. de Broglie, Ph.D. thesis, Université de Paris, 1924.

[35] E. Esterman and O. Stern, "Diffraction of a sodium atom beam," Zeits. Phys. **61,** 95 (1930).

[36] N. F. Ramsey, "A molecular beam resonance method with separated oscillating fields," Phys. Rev. **78,** 6 (1950).

[37] N. F. Ramsey, "History of atomic and molecular standards of frequency and time," IEEE Transactions on Instrumentation and Measurement **21,** 90–99 (1972).

[38] N. Ramsey, "Precise measurements of time," Am. Sci. **76,** 42 (1988).

[39] W. M. Itano and N. F. Ramsey, "Accurate measurement of time," Scientific American **269,** 56–65 (1993).

[40] O. Carnal and J. Mlynek, "Young's double slit experiment with atoms: a simple atom interferometer," Phys. Rev. Lett. **66,** 2689 (1991).

[41] P. L. Gould, P. J. Martin, G. A. Ruff, R. E. Stoner, J. L. Picque, and D. E. Pritchard, "Momentum transfer to atoms by a standing light wave: transition from diffraction to diffusion," Phys. Rev. A **43,** 585 (1991).

[42] F. Riehle, T. Kisters, A. Witte, J. Helmcke, and C. J. Bordé, "Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer," Phys. Rev. Lett. **67,** 177 (1991).

[43] *Atom interferometry*, P. R. Berman, ed., (Academic Press, 1997).

[44] D. S. Weiss, B. C. Young, and S. Chu, "Precision measurement of the photon recoil of an atom using atomic interferometry," Phys. Rev. Lett. **70,** 2706 (1993).

[45] A. Peters, K. Y. Chung, and S. Chu, "Measurement of gravitational acceleration by dropping atoms," Nature **400,** 849 (1999).

[46] M. J. Snadden, J. M. McGuirk, P. Bouyer, K. G. Haritos, and M. A. Kasevich, "Measurement of earth's gravity gradient with an atom interferometer-based gravity gradiometer," Phys. Rev. Lett. **81,** 971 (1998).

[47] J. M. McGuirk, G. T. Foster, J. B. Fixler, M. J. Snadden, and M. A. Kasevich, "Sensitive absolute-gravity gradiometry using atom interferometry," Phys. Rev. A **65,** 033608 (2002).

[48] M. K. Oberthaler, S. Bernet, E. M. Rasel, J. Schiedmayer, and A. Zeilinger, "Inertial sensing with classical atom beams," Phys. Rev. A **54,** 3165 (1996).

[49] T. L. Gustavson, P. Bouyer, and M. A. Kasevich, "Precision rotation measurements with an atom interferometer gyroscope," Phys. Rev. Lett. **78,** 2046 (1997).

[50] A. Lenef, T. D. Hammond, E. T. Smith, M. S. Chapman, R. A. Rubenstein, and D. E. Pritchard, "Rotation sensing with an atom interferometer," Phys. Rev. Lett. **78,** 760 (1997).

[51] T. L. Gustavson, A. Landragin, and M. A. Kasevich, "Rotation sensing with a dual atom interferometer Sagnac gyroscope," Class. Quantum Grav. **17,** 2385 (2000).

[52] C. R. Ekstrom, J. Schmiedmayer, M. S. Chapman, T. D. Hammond, and D. E. Pritchard, "Measurement of the electric polarizability of sodium with an atom interferometer," Phys. Rev. A **51,** 3883 (1995).

[53] D. M. Giltner, R. W. McGowan, and S. A. Lee, "Atom interferometer based on Bragg scattering from standing light waves," Phys. Rev. Lett. **75,** 2638 (1995).

[54] E. M. Rasel, M. K. Oberthaler, H. Batelaan, J. Schmiedmayer, and A. Zeilinger, "Atom wave interferometery with diffraction gratings of light," Phys. Rev. Lett. **75,** 2633 (1995).

[55] 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," Phys. Rev. Lett. **77,** 3331–3334 (1996).

[56] M. S. Jun, C. Y. Park, and D. Cho, "Low velocity intense source of rubidium atoms," Journal of the Korean Physical Society **33,** 260–264 (1998).

[57] C. Y. Park, M. S. Jun, and D. Cho, "Magneto-optical trap loaded from a low-velocity intense source," J. Opt. Soc. Am. B **16,** 994–997 (1999).

[58] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven, "Two-dimensional magneto-optical trap as a source of slow atoms," Phys. Rev. A **58,** 3891–3895 (1998).

[59] S. B. Nagel, C. E. Simien, S. Laha, P. Gupta, V. S. Ashoka, and T. C. Killian, "Magnetic trapping of metastable $^3P_2$ atomic strontium," Phys. Rev. A **67,** 011401 (2003).

[60] X. Xu, T. H. Loftus, J. L. Hall, A. Gallagher, and J. Ye, "Cooling and trapping of atomic strontium," J. Opt. Soc. Am. B **20,** 968 (2003).

[61] F. Sorrentino, G. Ferrari, N. Poli, R. Drullinger, and G. M. Tino, "Laser cooling and trapping of atomic strontium for ultracold atom physics, high-precision spectroscopy and quantum sensors," arXiv:physics/0609133v1 (2006).

[62] W. Mende, K. Bartschat, and M. Kock, "Near-threshold photoionization from the Sr $^1P_1$ state," J. Phys. B **28,** 2385–2393 (1995).

[63] S. L. Gilbert and C. E. Wieman, "Laser cooling and trapping for the masses," Optics and Photonics News pp. 8–14 (1993).

[64] L. A. Orozco, "Laser cooling and trapping of neutral atoms," AIP Conf. Proc. **464,** 67–90 (1999).

[65] A. M. Steane, M. Chowdhury, and C. J. Foot, "Radiation force in the magneto-optical trap," J. Opt. Soc. Am. B **9,** 2142–2158 (1992).

[66] H. J. Metcalf and P. van der Straten, in *Laser cooling and trapping*, R. S. Barry, J. L. Birman, J. W. Lynn, M. P. Silverman, H. E. Stanley, and M. Voloshin, eds., (Springer, 1999).

[67] L. Allen and J. H. Eberly, *Optical resonance and two-level atoms* (Wiley, New York, 1975).

[68] M. L. Citron, H. R. Gray, C. W. Gabel, and C. R. Stroud, "Experimental study of power broadening in a two-level atom," J. Phys. Rev. A **16,** 1507 (1977).

[69] W. D. Phillips, in *Proceedings of the International School of Physics Enrico Fermi: Course CXVIII*, E. Arimondo, W. D. Phillips, and F. Strumia, eds., (North-Holland, 1992), pp. 289–344.

[70] F. Riehle, *Frequency standards: basics and applications* (John Wiley and Sons, 2006).

[71] F. Ruschewitz, J. L. Peng, H. Hinderth, N. Schaffrath, K. Sengstock, and W. Ertmer, "Sub-kilohertz optical spectroscopy with a time domain atom interferometer," Phys. Rev. Lett. **80,** 3173 (1998).

[72] G. Ferrari, P. Cancio, R. Drullinger, G. Giusfredi, N. Poli, M. Prevedelli, C. Toninelli, and G. M. Tino, "Precision frequency measurement of visible intercombination lines of strontium," Phys. Rev. Lett. **91,** 243002 (2003).

[73] M. Yasuda and H. Katori, "Lifetime measurement of the $^3P_2$ metastable state of strontium atoms," Phys. Rev. Lett. **92,** 153004 (2004).

[74] T. Ido, T. H. Loftus, M. M. Boyd, A. D. Ludlow, K. W. Holman, and J. Ye, "Precision spectroscopy and density-dependent frequency shifts in ultracold Sr," Phys. Rev. Lett. **94,** 153001 (2005).

[75] R. L. Targat, X. Baillard, M. Fouché, A. Brusch, O. Tcherbakoff, G. D. Rovera, and P. Lemondé, "Accurate optical lattice clock with ^{87}Sr atoms," Phys. Rev. Lett. **97,** 130801 (2006).

[76] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," Phys. Rev. Lett. **59,** 2631 (1987).

[77] C. Monroe, W. Swann, H. Robinson, and C. E. Wieman, "Very cold trapped atoms in a vapor cell," Phys. Rev. Lett. **65,** 1571 (1990).

[78] K. Lindquist, M. Stephens, and C. Wieman, "Experimental and theoretical study of the vapor-cell Zeeman optical trap," Phys. Rev. A **46,** 4092 (1992).

[79] C. Monroe, H. Robinson, and C. Wieman, "Observation of the cesium clock transition using laser-cooled atoms in a vapor cell," Opt. Lett. **16,** 50–52 (1991).

[80] N. Ramsey, *Molecular beams* (Oxford University Press, 1985).

[81] K. Vant, J. Chiaverini, W. Lybarger, and D. J. Berkeland, "Photoionization of strontium for trapped-ion quantum information processing," arXiv:quant-ph/0607055v1 (2006).

[82] C. J. Dai, S. F. Hu, and J. Lu., "Multistep excitation of autoionizing states of neutral strontium," Journal of Quantitative Spectroscopy and Radiative Transfer **56,** 255 (1996).

[83] J. D. Bierlein and H. Vanherzeele, "Pottassium titanyl phosphate: properties and new applications," J. Opt. Soc. Am. B **6,** 622 (1989).

[84] Y. Li, Y. Zhao, Y. Peng, T. Yang, J. Cao, Z. Fang, and E. Zang, *Conference on Precision Electromagnetic Measurements Digest* (CPEM, 2008), pp. 190–191.

[85] R. L. Targat, J. Zondy, and P. Lemondé, "75%-efficiency blue generation from an intracavity PPKTP frequency doubler," arXiv:physics/0408031v1 (2004).

[86] E. D. Black, "An introduction to Pound-Drever-Hall laser frequency stabilization," Am. J. of Phys. **69,** 79–87 (2001).

[87] C. J. Dai, S. F. Hu, and J. Lu, "Multistep excitation of autoionizing states of neutral strontium," Journal of Quantitative Spectroscopy and Radiative Transfer **56,** 255 (1996).

[88] P. Bouyer, T. L. Gustavson, K. G. Haritos, and M. A. Kasevich, "Microwave signal generation with optical injection locking," Opt. Lett. **21,** 1502–1504 (1996).

[89] R. W. Fox, C. W. Oates, and L. W. Hollberg, in *Cavity-enhanced spectroscopies, volume 40 (experimental methods in the physical sciences)*, R. van Zee and J. P. Looney, eds., (Academic Press, 2003), Vol. 40.

[90] D. W. Preston, "Doppler-free saturated absorption: laser spectroscopy," Am. J. of Phys. **64,** 1432–1436 (1996).

[91] K. G. Libbrecht and J. L. Hall, "A low-noise high-speed diode laser current controller," Rev. Sci. Inst. **64,** 2133–2135 (1993).

[92] http://www.physics.byu.edu/faculty/durfee/electronics.php.

[93] "Precision, very low noise, low input bias current operation amplifiers: AD8671/AD8672/AD8674," Technical report, Analog Devices Inc. (2005) .

[94] "LT1028/LT1128 ultralow noise precision high speed op amps," Technical report, Linear Technology Corporation (1992) .

[95] "LM317 adjustable regulator," Technical report, Fairchild Semiconductor (2001) .

[96] S. Chiow, Q. Long, C. Vo, H. Müller, and S. Chu, "Extended-cavity diode lasers with tracked resonances," Appl. Opt. **46,** 7997–8001 (2007).

[97] A. Bennett, Senior Thesis, Brigham Young University, 2010.

[98] M. van Zijll, Senior Thesis, Brigham Young University, 2007.

[99] http://www.veeco.com/promos/vapor-pressure-curves-of-the-elements.aspx, Veeco Instruments Inc., 2011.

[100] K. Vogel, Ph.D. thesis, University of Colorado, 1999.

[101] D. Risold, B. Hallstedt, L. J. Gauckler, H. L. Lukas, and S. G. Fries, "Thermodynamic optimization of the Ca-Cu and Sr-Cu systems," Calphad **20,** 151–160 (1996).

[102] http://www.elementsales.com, Metallium Inc., 2011.

[103] S. P. Hansen, "Some vacuum basics," online, http://www.belljar.net/basics.htm (2008).

[104] M. L. Minges, in *Electronic materials handbook: volume 1 packaging*, ASM International Handbook Committee, ed., (ASM International, 1989).

[105] "Alumina," Technical report, Superior Techical Ceramics Corp. (2010) .

[106] C. Brauda, "Text book values of various dielectric materials," online, http://members.gcronline.com/cbrauda/0007.htm, 1999.

[107] "Dow Corning 732 multi-purpose sealant," Technical report, Dow Corning Corporation (2008) .

[108] A. Noble and M. Kasevich, "UHV optical window seal to conflat knife edge," Rev. Sci. Inst. **65,** 3042 (1994).

[109] P. Danielson, "Gas loads from elastomer and metal seals", online, http://www.vacuumlab.com/Articles/Gas%20Loads%20Elastomer-Metal%20Seals.pdf (2000).

[110] E. Korsunsky, "Atomic beamsplitter based on Raman transitions in four standing light waves," Quantum Semiclass. Opt. **10,** 477–483 (1998).

[111] F. Cecelja, M. Bordovsky, and W. Balachandrant, "Lithium niobate sensor for measurement of dc electric fields," IEEE Transactions on Instrumentation and Measurement **50,** 465–469 (2001).

[112] E. Ishiguro, T. Arai, M. Mizushima, and M. Kotani, "On the polarizability of the hydrogen molecule," Proc. Phys. Soc. A **65,** 178–187 (1952).

[113] D. Spelsberg and W. Meyer, "Static dipole polarizabilities of N_2, O_2, F_2, and H_2O," J. Chem. Phys. **101,** 1282–1288 (1994).

[114] P. T. van Duijnen and M. Swart, "Molecular and atomic polarizabilities: Thole's model revisited," Phys. Chem. A **102,** 2399–2407 (1998).

[115] P. W. Milonni and J. H. Eberly, *Lasers* (John Wiley and Sons, 1988).

[116] T. Hof, D. Fick, and H. J. Jansch, "Application of diode lasers as a spectroscopic tool at 670 nm," Opt. Commun. **124,** 283 (1996).

[117] A. S. Arnold, J. S. Wilson, and M. G. Boshier, "A simple extended-cavity diode laser," Rev. Sci. Inst. **69,** 1236 (1998).

[118] 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," Opt. Commun. **117,** 541 (1995).

[119] A. Schawlow and C. Townes, "Infrared and optical masers," Phys. Rev. **112,** 1940 (1958).

[120] C. E. Wieman and L. Hollberg, "Using diode lasers for atomic physics," Rev. Sci. Inst. **62,** 1 (1991).

[121] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator," Appl. Phys. B **31,** 97 (1983).

[122] Manncorp, "Manncorp Products Catalog," online, http://www.manncorp.com (2009).

[123] K. Maxon, “Have you seen my new soldering iron?,” Encoder: The Newsletter of the Seattle Robotics Society 6 (2000).

[124] “5V, serial input voltage output, 16-bt DACs AD5541/AD5542,” Technical report, Analog Devices (1999) .

[125] M. Cohen, "An alternative to the solder iron," online, http://www.chipquik.com/alt_sold_new.htm (2011).

[126] L. Enriquez, "A brief introduction to prototyping with Surface Mount Technology (SMT)," online, http://www.geocities.com/vk3em/smtguide/websmt.html (2000).

[127] Sciencelab, "Material safety data sheet mineral oil", online, http://www.sciencelab.com/msds.php?msdsId=9927364 (2010).

[128] E. M. Bridge, J. Millen, C. S. Adams, and M. P. A. Jones, “A vapor cell based on dispensers for laser spectroscopy,” Rev. Sci. Inst. 80 (2009).

[129] G. Bruzzone, E. Franceschi, and F. Merlo, “On the Sr-Pb system,” Journal of Less Common Metals **81,** 155–160 (1981).

[130] M. Bernardini *et al.*, “Air bake-out to reduce hydrogen outgassing from stainless steel,” J. Vac. Sci. Technol. A **16,** 188 (1998).

[131] P. Danielson, "Gas loads and o-rings", online, http://www.vacuumlab.com/Articles/Gas%20Loads%20and%20O-Rings.pdf (2000).

[132] P. Hait, “The application of polyimide to ultrahigh vacuum seals,” Vacuum **17,** 547–550 (1967).

Index