



All Theses and Dissertations

2007-11-28

Construction of a Calcium Matter-Wave Interferometer

Christopher Joseph Erickson
Brigham Young University - Provo

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>

 Part of the [Astrophysics and Astronomy Commons](#), and the [Physics Commons](#)

BYU ScholarsArchive Citation

Erickson, Christopher Joseph, "Construction of a Calcium Matter-Wave Interferometer" (2007). *All Theses and Dissertations*. 1227.
<https://scholarsarchive.byu.edu/etd/1227>

This Thesis is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in All Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact scholarsarchive@byu.edu, ellen_amatangelo@byu.edu.

CONSTRUCTION OF A CALCIUM MATTER-WAVE
INTERFEROMETER

by

Christopher Joseph Erickson

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Physics and Astronomy

Brigham Young University

December 2007

Copyright © 2007 Christopher Joseph Erickson

All Rights Reserved

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Christopher Joseph Erickson

This thesis has been read by each member of the following graduate committee
and by majority vote has been found to be satisfactory.

Date

Dallin S. Durfee, Chair

Date

Scott D. Bergeson

Date

Michael J. Ware

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Christopher Joseph Erickson in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

Date

Dallin S. Durfee
Chair, Graduate Committee

Accepted for the Department

Ross L. Spencer, Chair
Department of Physics and Astronomy

Accepted for the College

Thomas W. Sederberg, Associate Dean
The College of Physical and Mathematical Sciences

CONSTRUCTION OF A CALCIUM MATTER-WAVE INTERFEROMETER

Christopher Joseph Erickson

Department of Physics and Astronomy

Master of Science

ABSTRACT

I describe the construction of a calcium matter-wave interferometer. The interferometer is based on a Ramsey-Bordé scheme, and uses a thermal beam of atoms excited by an optical-frequency transition in calcium. In our experiment four $\frac{\pi}{2}$ pulses of light are delivered to the atoms, which split and recombine the wave functions of the atoms. Our experimental design minimizes first-order Doppler shifts, and allows for the cancellation of systematic errors including phase shifts due to rotation and acceleration. I describe the individual components of the interferometer and its assembly. The requirements for the electronics used in the experiment as well as their design and performance are described in great detail. I also give an overview of the techniques used to passively stabilize the laser and optical components. Finally, I report on the current status of the experiment as well as detail future work to be done on the apparatus.

ACKNOWLEDGMENTS

During the the time that I have lived in Provo, Utah I have grown to love it, and I am grateful for the many opportunities afforded me here. Most of these have come by virtue of attending Brigham Young University. BYU does a wonderful job of creating an environment conducive to the culture of learning and excelling. It goes out of the way to provide opportunities for scholarship, work, and extra-curricular growth at a fraction of the cost.

The Physics and Astronomy Department, in particular, goes out of its way to provide the best educational opportunities that it can. Having spent my entire pursuit of higher education within this department, I would like to express my gratitude to the faculty and staff who have spent many long hours to develop quality courses and learning materials. They are also mercifully understanding of the unique situations of individual students.

The excellence of my education in the Physics Department here at BYU has been taken note of by the Department of Defense of our country. They offered me a fellowship, sponsored by the Air Force, toward the completion of a Master degree. Since I had spent my first year entirely engaged in course work toward a Ph.D., this change of degree most definitely translated into a change of pace in the manner that research was being conducted. All in all, the Department of Defense was a large supporter and a large driving factor behind the research conducted for this thesis during the last year. I am honored to accept this award and the future opportunity it affords me to serve my country

with the skills I have developed pursuing this research.

I would like to acknowledge several of the people who have helped and inspired me throughout my education. Among them are Dr. Allred, Dr. Davis, Dr. Hirschmann, and Dr. Nielsen all of whom exhibited an interest and excitement in physics that I found contagious. I am grateful to Wes Lifferth, John Ellsworth, and Scott Daniels whose knowledge and skills have been integral to completing many of the projects described in this thesis. I would like to thank Nan for solving any and every problem I could throw at her. If it was not for the demands of her job, I am sure she would have been in the lab helping me conduct research as well. I would especially like to thank Scott Bergeson who possesses great energy, a willingness to be of help, a sense of humor, and the incredible ability to explain physical phenomena to me in a manner that I can understand.

Above all in the physics department, I owe many thanks to my advisor, Dallin Durfee. His willingness to allow me to touch extremely expensive equipment with minimal training, and his understanding and forgiveness when I break them are impressive. Despite that fact more than one group of students has now come and gone in the lab, he continues to overcome the temptation to kick me out with one of them. With great patience, he has helped me to come to an understanding of the physical principles that govern this world. He stands as a great example to me of devotion to students and moral resolution. He has become a great friend, as well as a mentor.

I would like to thank my wonderful family. My parents are great examples of hard work who instilled in me a hunger for education, and taught me to ever pursue my dreams.

More important than any of the afore mentioned is my wonderful wife,

Katie, and my two beautiful children, Ben and Kallie. Without their support, love, patience, and forgiveness I would be an incompetent wreck. Whatever spell has been cast over my wife to inspire her to stay by me through such a demanding education, is beyond my comprehension. She is loving, kind, charitable, and understanding. She is my hero, and I love her.

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

Table of Contents	ix
List of Figures	xi
1 Introduction	1
1.1 A Short History of Atom Interferometry	1
1.2 Properties and Applications of Atom Interferometers	2
1.3 Motivation	3
1.4 Outline of the Thesis	4
2 Experimental Design of a Thermal Beam Atom Interferometer	5
2.1 Ramsey-Bordé Interferometry	6
2.1.1 Phase Sensitivity of Interferometer	8
2.2 Overview of the Apparatus	10
2.2.1 Importance of the design symmetry	12
3 Interferometer Components	15
3.1 Precision Optics	15
3.2 Thermal Beam Oven	16
3.3 423nm Laser	18
3.4 657nm Laser	18
3.4.1 657nm Diode	20
3.4.2 Pound-Drever-Hall Lock	20
3.4.3 Laser Locking Techniques	21
3.4.4 Tuning the Locked Laser Frequency to the Atomic Transition	23
3.4.5 Injection Locked Diode Laser	25
4 Assembly of the Experimental Apparatus	27
4.1 Alignment of the Atom Beam	27
4.2 Alignment of the Laser Beams	30
4.3 Alignment of the Fluorescence Probe	31

5	Low-noise, High-speed Electronics	33
5.1	Photodiode Homodyne Detector	34
5.1.1	Performance of the Homodyne Detector	36
5.1.2	High Speed Op-amps	37
5.2	Future Homodyne Detector	38
5.3	Lock-circuit PID Controller	40
5.3.1	PID Processing	41
5.3.2	Laser Frequency Scan Function	43
5.3.3	Scan-balancer Function	43
5.4	Diode Laser Current Controller	44
5.4.1	Design and Construction	45
5.4.2	Results and Discussion	54
5.4.3	Conclusion	59
5.5	Microprocessor	60
6	Passive Stabilization of the Optical Cavity and the 657 nm Laser	63
6.1	High-finesse Optical Cavity	63
6.2	Isolation of Electronics and Mirror Mounts	64
7	Conclusions and Future Work	67
7.1	Conclusions	67
7.2	Current Status of the Interferometer	68
7.3	Optimization of the Interferometer	68
7.4	Construction of a Dual Species Interferometer	70
7.5	Future Measurements and Applications	70
A	Surface Mount Soldering in a Toaster Oven	73
	Bibliography	77
	Index	83

List of Figures

2.1	Recoil diagram of a Ramsey-Bordé interferometer.	7
2.2	Diagram of selected calcium transitions	8
2.3	Diagram of the experimental design	10
3.1	Calcium thermal beam oven	17
3.2	Schematic of the 657 nm laser.	19
3.3	Comparison of optical cavity frequencies to the intercombination line	24
4.1	Layout of the experiment.	28
4.2	Atom beam skimmers.	29
5.1	Functional diagram of the laser electronics	34
5.2	Schematic of current homodyne detector	35
5.3	Homodyne detector PCB layout	36
5.4	Schematic of next generation homodyne detector	38
5.5	Future homodyne detector PCB layout	39
5.6	Schematic of lock-circuit.	40
5.7	Lock-circuit PCB	42
5.8	Schematic of current driver circuit.	46
5.9	Current driver PCB	47
5.10	Accuracy and repeatability of current set point.	54
5.11	Spectral density of current driver noise.	55
5.12	Temperature dependence of current driver output.	57
5.13	Transient response of current driver.	58
5.14	Bode plot of current driver modulation circuit.	59
5.15	Schematic for microcontroller board	60
5.16	Top layer of microcontroller PCB	61
5.17	Bottom layer of microcontroller PCB	62

Chapter 1

Introduction

A matter-wave interferometer may best be described by analogy to an optical interferometer. An optical interferometer, such as a Michelson-Morley interferometer [1], splits the wave function of light into two distinct parts that travel different paths, and then recombines them. Likewise, a matter wave interferometer is a device which spatially splits and recombines the wave function of an atom. In an optical interferometer matter, in the forms of mirrors and beam splitters, is used to split and recombine light. However, in a matter wave interferometer light is typically used to split and recombine the wave function of an atom.

1.1 A Short History of Atom 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 [2]. 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 [3]. The first atom interferometer was developed by Norman Ram-

sey as an atomic clock [4–7]. Even though the atom beam never split spatially, its wave function was nonetheless split and recombined. The first true separated beam atom interferometer was published in 1991 by Carnal and Mlynek, who used narrow slits to split and recombine metastable helium atoms [8]. This was followed almost immediately by Pritchard, who used micro-fabricated diffraction gratings to realize a sodium interferometer [9]. Shortly thereafter, Bordé published his realization of a Ramsey spectrometer using beams of light to split and recombine a beam of thermal atoms [10].

Since then, interferometers of varying design have been developed and used to contribute to the field of precision measurement [11]. Different ways of splitting and recombining an atom’s wave function have been used including: slits [8], diffraction gratings [9], Raman transitions [12–15], traveling light waves [10], and the diffraction of atoms off of standing light waves [16, 17]. These varying devices have been used to perform a variety of experiments including measuring the photon recoil of an atom [18], measuring the local acceleration g due to gravity [19, 20], measuring the Earth’s gravity gradient [13, 15], making precision rotation measurements [12, 14, 21, 22], measuring the electric polarizability of atoms [23], and measuring time variance in the fine structure constant [24–28].

1.2 Properties and Applications of Atom Interferometers

Atom interferometers have several advantages over optical 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. Whereas 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. Atom interferometers can perform precision measurements of time variance in the fine structure constant [29, 30], time dilation and Kennedy-Thorndike experiments [31–34], Lorentz violation studies [35, 36], and the afore-mentioned inertial force measurements. We have also proposed experiments to use an ion interferometer in search of a photon rest mass and violations of Coulomb’s Law [37, 38].

The interferometer described in this thesis is similar to the experiments described in [10, 39–42]. It is based on beams of visible light that split and recombine the wave functions of calcium atoms traveling in a thermal beam. The interferometer will be used as an atomic clock and in the future will perform measurements of fundamental constants and tests of relativity. Initial work with this interferometer will also be a proof of principle for the future development of a transportable frequency standard, an atom interferometer accelerometer, and an atom interferometer gyroscope.

1.3 Motivation

Large amounts of money and time have been spent in recent years to develop the next generation of frequency standards with fractional stability at or above 10^{-16} . The idea that such precision could be obtained with clever engineering using the means of a smaller laboratory led us to layout the design for the interferometer described in this thesis. We calculated that such precision could be obtained with a thermal interferometer by using our design. Our design allows for the experiment to be used as a clock and a gyroscope, with the ability to optimize the experiment for one or the

other. While the interferometer we describe in this thesis is a single thermal beam calcium interferometer, the experimental design allows for the addition of a strontium interferometer within the same space in order to perform measurements of the time rate-of-change of fundamental constants (see section 7.4).

1.4 Outline of the Thesis

The following chapter describes in detail the design and theory of operation of our interferometer. Chapter 3 then breaks down the individual components of the interferometer and describes their design, construction, and performance characteristics. Physical assembly and alignment of the interferometer is described in chapter 4. This is followed by an in-depth account of the electronics developed for the laser and servo systems in chapter 5. Chapter 6 details the passive stabilization techniques we implemented for the laser system. The last chapter draws conclusions, describes the current state of the interferometer, and lays out future work to be done on and with the apparatus. The appendix describes time-saving and cost-effective surface mount soldering techniques we developed over the course of this project. While time constraints prevented us from taking data with the interferometer, we were able to obtain the necessary optics, construct the laser and electronics systems, and assemble the apparatus.

Chapter 2

Experimental Design of a Thermal Beam Atom Interferometer

Thermal atom beams have been used in many interferometer experiments including atomic clocks (such as the now retired NIST-7 national time standard) [39,43], inertial force experiments [12,44], and experiments for fundamental studies of matter-waves [8,40,41,45,46].

Thermal beam interferometers are widely used because they have certain advantages over trapped atom devices. The process of trapping and cooling atoms rejects the majority of the atoms available for the experiment. Thermal beam experiments, however, can be designed to use a majority of the atoms available resulting in greater signal-to-noise ratios and greater precision. Also, a thermal beam apparatus has continuous rather than pulsed read-out, and its construction can be simpler, more robust, smaller, and cheaper than trapped atom devices. Plus, thermal beam devices can more easily be converted to run experiments on a range of different elements other than the one they were initially designed for.

Thermal beam devices also have their drawbacks. First and second order Doppler

shifts are the primary limiting factors in the most recent thermal-beam optical-frequency atomic clocks [39, 42]. Short observation times can also be a limitation. In this case atomic-fountain microwave standards [47–49] and trapped-atom primary standards [50, 51] that slow atoms down by four orders of magnitude from thermal beam velocities have the advantage.

Our interferometer will be a secondary standard with stability to rival the primary standards while still using a thermal beam. This is done by exploiting the experimental design to cancel Doppler shifts and systematic errors. Further on we describe how Doppler shifts are minimized by our experimental design, and why short observation times are not an issue in our interferometer.

2.1 Ramsey-Bordé Interferometry

The interferometer we have constructed is a Ramsey-Bordé interferometer. A Ramsey-Bordé interferometer is, most simply, an optical frequency version of the microwave Ramsey spectrometer [4, 52].

The basic functional diagram of a Ramsey-Bordé interferometer is shown in figure 2.1. Four laser beams, shown by dotted arrows, cross a thermal atom beam at right angles. Even though the laser beams are continuous wave, the atoms effectively “see” a pulse of light as they travel through the beam. The laser intensity and time of flight of the atoms through the laser are controlled such that each laser beam delivers a $\frac{\pi}{2}$ “pulse” of light to the atoms.

As the atoms intersect the first of the laser beams, this $\frac{\pi}{2}$ pulse of light puts the atoms into an equal superposition of excited and ground states. Photon absorption in the excited state part of the superposition causes it to recoil away from the ground state part of the superposition. The second laser beam causes the excited state part




Figure 2.1 Recoil diagram of Ramsey-Bordé interferometry. Laser beams are shown by thin dotted lines while thick solid and dashed lines indicate paths for the two closed interferometer loops. Photon absorption by the atoms in the first laser beam causes the excited-state part of the atoms to recoil at 1.5 cm/s. The second laser beam creates parallel ground state beams ($20\mu\text{m}$ separation) via stimulated emission, and the remaining laser beams serve to recombine the atoms. The separation of the excited and ground states is greatly exaggerated and only closed loop interferometer paths are shown.

of the superposition to fall back to the ground state via stimulated emission. This creates two parallel ground state atom beams. The third and fourth laser beams are anti-parallel to the first two and recombine the atom beam using the same process. A second closed interferometer loop also exists as shown by thick dashed arrows in figure 2.1.

Our implementation of Ramsey-Bordé interferometry has several advantages. First, Ramsey-Bordé interferometry itself allows the use of a thermal atom beam which can yield large signal to noise ratios [39]. Second, Ramsey-Bordé interferometry also results in the cancellation of first order Doppler shifts as far as the pairs of laser beams in figure 2.1 are anti-parallel [10]. Our own experimental design ensures that




Figure 2.2 A diagram of calcium transitions showing only relevant transitions.

this cancellation is very precise. This design also allows for the reduction of other systematic errors such as second order Doppler shifts (see section 2.2.1).

We based the interferometer on the 1S_0 to 3P_1 intercombination line in calcium shown in figure 2.2. This line has a couple of advantages. The lifetime of the upper state (.38 ms) is long enough to make precise measurements, but short enough that we can still drive the transition with our laser. A side effect of the lifetime is that we can enjoy maximum observation time for the atom using a thermal beam. Also, the use of a single photon transition prevents the interferometer from being limited by AC Stark shifts, which can be a major limitation in interferometers based on a two-photon Raman transition [12, 14].

2.1.1 Phase Sensitivity of Interferometer

The interferometer phase is sensitive to the frequency of the interaction laser beams shown in figure 2.1. The phase of the interferometer is related to the frequency difference between the laser and the atomic transition by

$$\Delta\Phi = \Delta\omega\Delta t$$

where $\Delta\omega$ is the frequency difference between the laser and the transition and Δt is the time over which the upper and lower paths shown in figure 2.1 are in different states. For our interferometer a frequency difference of 10 Hz between the laser and the atomic transition will produce a phase shift of 1.67×10^{-2} radians, over the time it takes an atom to propagate through the interferometer.

The interferometer phase is also sensitive to Doppler shifts. At the velocity of the thermal beam in the experiment (~ 750 m/s), angles between the atoms and the laser deviating by milli-radians from orthogonal will produce frequency shifts on the kilo-Hertz level. However, these Doppler shifts are constant during the course of the experiment and are reduced by the experimental design as described in section 2.2.1.

The interferometer is also sensitive to acceleration and rotational phase shifts. The rotational phase shift is given by

$$\Delta\Phi_{rot} = \frac{4\pi\vec{\Omega} \cdot \vec{A}}{\lambda v}$$

where $\vec{\Omega}$ is the rotation rate vector, \vec{A} is the vector area enclosed by the interferometer, λ is the deBroglie wavelength of the atoms, and v is the velocity of the atoms. The acceleration phase shift is given by

$$\Delta\Phi_{acc} = \frac{-2\pi A}{\lambda v^2} a$$

Where A is the area enclosed by the interferometer, a is the acceleration, and λ and v are as defined for the rotational phase shift.

Our interferometer is 21 cm in total length with 10 cm separating the first and second as well as the third and fourth laser beams, and about 1 cm spacing between the second and third beams. Due to the 1.5 cm/s recoil velocity of the excited state part of the superposition caused by photon absorption, the maximum separation of the ground and excited state superpositions is $2 \mu\text{m}$. This means the interferometer




Figure 2.3 Atom interferometer layout. A thermal beam of atoms is generated in an oven and then orthogonally crosses four laser beams, which are numbered 1 through 4. The laser beams are made parallel by matching sets of prisms as discussed in section 3.1. The exit apertures and extra detectors serve for reversing the direction of both the atom and laser beams through the apparatus. Also shown is a strontium oven and lasers tuned to atomic transitions in strontium, as well as pump lasers for transverse cooling both calcium and strontium, which setup is described in chapter 7.

will produce a phase shift of 277.67 radians of interferometer phase per rad/s of rotation, and about 1.81 radians per g of acceleration.

2.2 Overview of the Apparatus

The experimental design of our interferometer is shown in figure 2.3. Atoms originate in a thermal oven at the left of the figure and travel to the right. The four numbered laser beams that intersect the atom beam in the central part of figure 2.3 are the four laser beams described earlier in figure 2.1. The phase of the interferometer is detected with the fluorescence probe on the right hand side of the figure.

The atom beam is produced in an oven at 792 °C giving the atoms a mean velocity of 750 m/s. The direction and divergence of the atom beam are set by a

pair of skimmers 2 mm in diameter spaced 10 cm apart. Based off of laser absorption measurements, the beam flux through the skimmers without transverse cooling is calculated to be 1.383×10^{13} atoms/s . By calculating the trajectories of the atoms under transverse cooling, we would expect the flux to increase to about 1×10^{17} atoms/s.

The interaction laser enters the apparatus through a spatial filter shown in the lower left of figure 2.3. The purpose of the filter is to create a nice spatial mode so that the interferometer is not affected by random intensity distribution in the laser beam. It also keeps the experiment precise by ensuring that the four laser beams will always be in the exact same position when the laser beam is properly aligned to the filter. If service is performed on the laser, or the experiment is performed on a different element, aligning the laser through the spatial filter precisely aligns the laser through the rest of the interferometer. Most importantly, the filter converts angular drift in the laser beam to intensity drift in the interferometer, which we can detect and remove from the signal.

The interferometer depends heavily on the quality of the optics that follow the spatial filter. The laser is split by a precision beam splitter seen near the center of figure 2.3 so that the light intensity in the two beams varies by less than 5%. These two beams are then reflected and retro-reflected by a series of prisms to create the four numbered beams seen in figure 2.3. The penta prisms, when used in matched pairs, create beams that are parallel to within 50 micro-radians or less. The right angle prisms also create beams that are parallel to better than 50 micro-radians. By using precision prisms rather than mirrors, Doppler shifts created by an imperfect alignment of the atom and laser beams are reduced.

The vertical alignment of the laser does not affect Doppler shifts and is not as critical. This alignment is done by hand with the four beams sitting in the same plane to within .5 mm over a distance of 20 cm. More detail on this alignment is given in

chapter 4.

Not shown in the schematic, but seen in figure 4.1, are a series of wire coils that span the interferometer concentric to the atom beam. These coils provide a magnetic field that pushes the σ^+ and σ^- transitions out of resonance with the driving light field. It also provides a quantization axis. Therefore, the atoms are only driven on one transition and no atoms are lost to optical pumping. Four different coils are spaced across the length of the interferometer with each coil consisting of 10 turns. With a current of about 5 amps, a magnetic field of a few Gauss is created in the direction of propagation of the atoms. The alignment of the coils does not have to be precise, with calculations showing that individual coils can be off center by as much as a centimeter before adverse effects on the interferometer's signal will be seen.

2.2.1 Importance of the design symmetry

One of the strengths of the interferometer's design lies in its symmetry. Seen in figure 2.3 are collimation apertures at the output of the interferometer, a spatial filter after the light has intersected the atoms, and an empty detector. The purpose of these items is to allow the light and the atoms to propagate in reverse through the interferometer. By running the interferometer in reverse, residual and higher-order Doppler shifts will be canceled. Also, running the interferometer in reverse will allow us to separate and remove rotational and acceleration phase shifts from the signal.

First-order Doppler shifts in the interferometer cancel to the extent that the two pairs of laser beams in figure 2.3 are anti-parallel. This cancellation is due to the symmetry of the Ramsey-Bordé interferometry and in practice it is limited by the precision of the prisms used to direct the laser beams. The $50 \mu\text{rad}$ accuracy of the prisms means that an atom traveling 750 m/s sees a maximum relative first-order Doppler shift of 57 kHz . First-order Doppler shifts should not shift the frequency

of the laser beams by more than a few parts in 10^{11} , with drifts in the frequency offset at less than a few parts in 10^{15} . By reversing the laser propagation through the interferometer it should be possible to remove long term drifts at the level of one part in 10^{16} .

Second-order Doppler shifts, or time dilation, will also be present in the interferometer. These shifts will result in offsets at the level of a part in 10^{12} . However, these offsets can be stabilized by precisely controlling the temperature of the oven as well as the intensity of the interaction laser. Currently, the oven is the main limiting factor and is stabilized to only a few degrees Celsius. This level of control only stabilizes the offsets from higher order Doppler shifts to about a part in 10^{14} . In the future, our goal is to stabilize these offsets to a part in 10^{16} . This will require the oven temperature to be controlled to $0.1\text{ }^\circ\text{C}$ and the laser intensity to a part in 10^4 .

Acceleration phase shifts and rotational phase shifts will also affect the interferometer signal [10, 52]. Passively stabilizing the interferometer will be critical to minimizing noise from these sources.

Since the interferometer is sensitive to acceleration and rotational phase shifts, it can also be used to measure them. With a predicted signal-to-noise ratio of about 3,000 in one second, we expect a rotational sensitivity of about 1.56×10^{-7} rad/s in one minute. The sensitivity to acceleration will be about $2.4 \times 10^{-5}g$ in one minute.

Chapter 3

Interferometer Components

3.1 Precision Optics

The optics which define and guide the laser are extremely important because the interferometer requires parallel laser beams in the horizontal plane that are equal in intensity and uniform in spatial mode. As stated in the previous chapter, we require the interaction laser beams to be parallel to within $50 \mu\text{rad}$. We also require optics with extremely good anti-reflective coatings to prevent loss of light as the laser propagates through the experiment.

We verified the accuracy of the prisms and their coatings ourselves in the laboratory. A previous member of the group, Rebecca Olson, developed an inexpensive means to measure the accuracy of the uncoated optics [53]. By interfering 1 cm diameter parallel beams generated with the optics the angle between the beams was directly measurable down to 0.13 mrad. However, by using curve-fitting techniques, angles down to 1 micro-radian were inferred. The wedge angles of the beam splitters all fell within 1 to 6 micro-radians. The penta-prisms, when matched, created parallel beams varying from nearly zero to 1.02 micro-radians in divergence. The right-angle

prisms were also found to create beams diverging by only a couple of micro-radians. For more information and greater detail on this technique see [53].

Only a few companies were willing to manufacture the optics to our specifications. In fact the penta prisms could not be manufactured to within spec and were made in pairs with the second prism correcting for the deflection error of the first. As manufactured, these prisms created parallel laser beams to within $5 \mu\text{rad}$. A separate company was hired to apply the anti-reflective and reflective coatings. Unfortunately, the coatings had a negative effect on the deflection angles of the optics. However, they still produce laser beams parallel to within $50 \mu\text{rad}$.

3.2 Thermal Beam Oven

The thermal beam oven for this interferometer is simple in design and is shown in figure 3.1. We constructed it after we found that an internal temperature gradient in the previous oven limited the atomic beam flux. The oven consists of an 8 inch long, $3/4$ inch diameter stainless steel tube mounted on a 6 inch Conflat (CF) vacuum flange.

A 1 inch long crucible holds the calcium at the end of the tube, and it is held there by a thin rod welded to it. The other end of this thin rod bolts to the CF flange and ensures that the crucible is always held in a fixed orientation. The crucible and rod are seen on the left hand side of figure 3.1.

A stainless steel cross capped with CF viewports is mounted on the oven at the output of the crucible. This allows for transverse cooling of the atom beam as soon as the atoms exit the crucible. A stainless steel baffle, shown on right hand side of figure 3.1, slides down each arm of the cross to prevent hot calcium from coating the viewports and blocking the cooling laser.




Figure 3.1 Calcium thermal beam oven. The calcium is housed in a crucible which slides down a length of stainless steel tube. The tube is capped with a copper slug for uniform heating. The cross on the tube allows transverse cooling of the calcium as it exits the oven. The baffle sleeve shown on the right prevents calcium from coating the viewports on the cross.

The crucible is heated external to the vacuum to make it easy to change and repair the heating element without opening the vacuum chamber. A 2 inch diameter copper slug encloses the end of the stainless steel tube housing the crucible to maintain a uniform temperature over the entire length of the crucible. Two high-temperature ceramic band heaters are installed on the slug, one of which is shown on the left hand side of figure 3.1. The oven and heating elements are wrapped with a ceramic insulation tape to help create a uniform temperature across the oven.

3.3 423nm Laser

The 423 nm laser serves as both the transverse cooling laser and the fluorescence probe. Since blue diodes at this wavelength are unavailable, and commercially available dye lasers are expensive, we will produce this light by frequency doubling an 846 nm diode laser. By amplifying an 846 nm laser diode with a tapered amplifier and then frequency doubling it with a BBO crystal, we anticipate that close to 100 mW of blue light could be produced. The design is simple enough and commonly done in other labs. Unfortunately, our tapered amplifiers turned out to be cut from a defective piece of silicon, causing the output of the amplifiers to not be single mode. New amplifiers have been ordered, but in the meantime we are borrowing 423 nm light from another experiment.

3.4 657nm Laser

The heart of the interferometer is a 657 nm grating-stabilized laser tuned the intercombination line in calcium. Due to the 410 Hz linewidth of the atomic transition, this laser requires linewidth on the kHz level in order for the interferometer to function. The goal, however, is to develop a laser with a sub-kHz linewidth. This narrow linewidth is achieved by locking the laser to a high finesse optical cavity [54]. As seen in figure 3.2, light from the laser is modulated by an electro-optic modulator (EOM) and locked to the high finesse cavity via the servo electronics (a functional diagram of the servo system is shown in figure 5.1). Once stabilized, its frequency is then shifted into resonance with the atoms by an acousto-optic modulator (AOM). This shifted light will be used to injection lock a second laser diode, which serves to boost the laser intensity.




Figure 3.2 Schematic of the 657 nm laser. The laser is stabilized by optical feedback from a diffraction grating. The output of the laser is then passed through two optical isolators. A wave plate and a polarizing beam splitter are then used to split off a fraction of the beam to lock to an optical cavity. The fraction of the beam to lock to the cavity is sent through an EOM. After the EOM, a lens focuses the light to match the TEM_{00} mode of the high finesse cavity. Light reflected off of the cavity is sent to a homodyne detector via a polarizing beam splitter. The other portion of the original laser beam is sent through an AOM to another optical isolator. A polarizing beam splitter sends the beam through the isolator to injection lock a second laser diode with higher output power. This laser is then sent to the interferometer. The AOM serves to shift the light locked to the optical cavity to be resonant with the atomic transition used in the interferometer.

3.4.1 657nm Diode

The laser consists of a 660 nm Circu-laser diode from Blue Sky Research in an extended cavity grating stabilized setup (ECDL) as seen at the top left of figure 3.2. Despite the fact that Circu-Laser diodes have a cylindrical lens in the diode package to reduce astigmatism, we found that when our diode is properly collimated the mode is still nearly one and a half times as long as it is wide. At low currents, the diode is stable enough to run in single mode. At high currents spatial hole burning occurs, resulting in the emission of multiple longitudinal modes. The diode will stably run single mode at around 70 mA, which only creates about 1.38 mW of output power for the laser.

3.4.2 Pound-Drever-Hall Lock

The ECDL is locked to the cavity using the Pound-Drever-Hall technique [54]. As seen in figure 3.2, the laser is sent through an EOM before passing to the optical cavity. The EOM dithers the phase of the light passing through it. This creates frequency side bands on the light, which are spaced from the main carrier frequency by the driving frequency of the EOM. When the main carrier frequency is resonant with the optical cavity, it is mostly transmitted while the side bands are completely reflected off of the cavity.

The quarter wave plate between the polarizing beam cube and the optical cavity in figure 3.2 reverses the polarization of the light reflected by the cavity from that of the light incident on the cavity. The polarizing beam cube then diverts the reflected side bands to a homodyne detector for feedback in the servo system.

The homodyne detector amplifies the signal and then demodulates it to create an error signal. A lock-circuit (shown in figure 5.1) then uses the error signal to determine

the current feedback to the laser head. For more information on the function of the servo electronics see chapter 5.

This method of locking to the cavity greatly reduces the linewidth of the laser. The free running linewidth of the diode is on the order of a giga-Hertz. In the extended cavity grating stabilized setup, the diode laser still has a linewidth of hundreds of kilo-Hertz. Locking the laser to the ultra-high finesse cavity further reduces the linewidth to the kilo-Hertz level.

We expect that the laser linewidth could be narrowed even further by narrowing the linewidth of the optical cavity that the laser is locked to. The optical cavity is specified to have a finesse of 300,000, and a linewidth of 10 kHz. However, we measured a linewidth of 50 kHz for the cavity. This reduction in finesse could be due to the ion pump serving the vacuum chamber for the cavity. When the ion pump was turned on it may have deposited a thin film on the inside surfaces of the cavity mirrors. If this is the case, cleaning the mirrors could result in further reducing the laser linewidth.

3.4.3 Laser Locking Techniques

Locking a laser to a high finesse cavity is no small task and requires patience. Whereas in most applications the free-running laser linewidth is smaller than the cavity linewidth, in our case the cavity linewidth is at least an order of magnitude smaller than the linewidth of the ECDL. This means that even when the laser is on resonance with a cavity mode, only a small portion of the light is actually on resonance with the cavity. This makes it difficult to carefully match the spatial mode of the laser to a cavity mode.

While all the cavity modes have the same linewidth, it is best to lock to the TEM00 mode. This mode is most similar to the laser beam spatially; therefore, more

light can be coupled into the cavity using this mode than any other mode. Due to this, the TEM_{00} mode can provide the best error signal for the servo circuitry.

To couple to the TEM_{00} mode, the laser must not only be correctly aligned to the cavity face, but the light must be spatially controlled to match the size and divergence of the mode. When the light is poorly aligned to the cavity, the laser will couple to both even and odd cavity modes as its frequency is scanned. When the beam mode is irregular and/or is not properly focused, the laser will still couple to even cavity modes (higher than the TEM_{00}) though the laser may be perfectly aligned.

To achieve a rough alignment to the optical cavity we found it beneficial to start with a second laser. By looking down the bore of the cavity by eye and placing two irises concentric with the cavity, we created a rough alignment to the cavity center. We then aligned a laser close in wavelength, a 633 nm Helium-Neon (He Ne) laser, to the cavity. The wavelength of the He Ne laser sits in the wings of the cavity's reflective coating range, which effectively lowers the Q of the cavity for that laser and makes it easier to couple to the cavity. When roughly aligned, the He Ne can then be seen by eye coupling to various cavity modes. The 657 nm laser was then aligned to the center of the cavity by overlapping it with the He Ne laser over the longest possible distance.

Subsequent adjustments of the laser alignment can be achieved by putting the cavity throughput on a CCD camera and slowly scanning the laser frequency. Once a fairly strong TEM_{00} mode is present in the scan, the throughput can be put on a high-speed photodiode and the scan speed increased. Fine tuning the alignment can be done by putting the servo error signal on an oscilloscope and maximizing the error signal that corresponds to the TEM_{00} mode. Once the alignment is good enough to lock the laser to the cavity, final tweaking of the alignment is done by optimizing transmission through the cavity while the laser is locked.

The response from our servo circuitry indicates that the locked laser has a linewidth as narrow as a few kilo-Hertz. Locking a diode laser to this level is no small feat and is made possible only by the excellent quality of the electronic components developed and passive stabilization techniques, both of which are detailed in the following chapters.

3.4.4 Tuning the Locked Laser Frequency to the Atomic Transition

Unfortunately, the high finesse cavity does not have a TEM_{00} mode that exactly matches the frequency of the calcium intercombination line. By scanning the laser through the high finesse cavity and a high temperature vapor cell simultaneously, a difference of about 500 MHz was measured between the closest TEM_{00} and the atomic resonance as shown in figure 3.3. To compensate for this difference in frequency, we use an AOM to shift the frequency of the locked laser to the atomic resonance.

In order to tune the locked laser to the atomic transition, light shifted by the AOM was first passed to a wave meter and the unlocked ECDL was scanned to roughly tune the shifted light to the intercombination line. The laser frequency was then scanned while passing the laser through the high temperature vapor cell. In the vapor cell, the atomic transition becomes Doppler broadened by 2 GHz, which only allows us to accurately tune the laser to within 10 MHz of the atomic transition [55]. Since the AOM can be scanned over a nearly 200 MHz range, the interferometer signal can then be found by scanning the AOM's driving frequency once the laser is locked to the optical cavity.

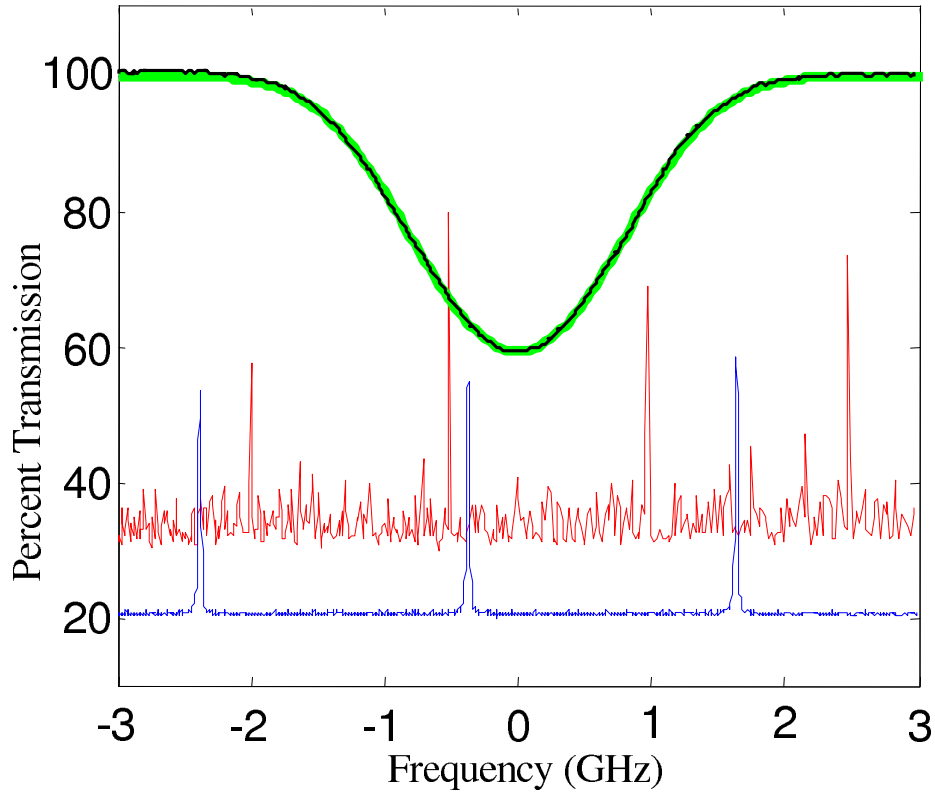


Figure 3.3 Comparison of optical cavity mode frequencies to the intercombination line in calcium. The thin black absorption curve is actual data, while the thicker green absorption curve is the calculated absorption for a 10 cm long column of Ca vapor at 700 °C [55]. The thin blue line is the transmission through a reference etalon. The thin red line is the transmission through the ultra-high finesse cavity. There is a 500 MHz difference between the nearest optical cavity mode and the center of the absorption curve.

3.4.5 Injection Locked Diode Laser

The use of so many optical components, coupled with the losses in the AOM, greatly decreases the laser intensity from the original 1.38 mW produced by the laser. After passing through the spatial filter and splitting the beam for the interferometer, less than .12 mW of light in remains in each of the four interferometer beams: a far cry from the .6 mW plus per beam that we anticipated needing to create a $\frac{\pi}{2}$ pulse. To remedy the situation we are going to injection lock a second laser with the light passed through the AOM as seen in figure 3.2. We purchased another Circu-laser diode at 660 nm with an output of 130 mW. By sending the light from the AOM into the reject port of an optical isolator to the diode we will force this diode to run in a stable wavelength mode with high output power. The light from this laser will be sent directly through the spatial filter of the interferometer and on to the atoms.

Chapter 4

Assembly of the Experimental Apparatus

Figure 4.1 shows the interferometer before it is placed under vacuum. The laser beams are labeled 1 through 4 to correspond with the labeling in figure 2.3. The spatial filter and collimation apertures are to the right of the picture, absent from view. After the spatial filter, the laser is split into two beams by the beam splitter at the top of the picture. These are the beams labeled “1” and “2”, and they are retro-reflected through a series of optics to create beams “3” and “4”. The atom beam crosses the lasers orthogonally, and is centered through the four magnetic coils. The collimation optics for the fluorescence are aligned with the probe laser at the left hand side of the picture.

4.1 Alignment of the Atom Beam

The atom beam is collimated by a pair of skimmers that define its direction and divergence. The alignment of the atom beam through the skimmers is critical. Pri-




Figure 4.1 Picture of the experiment just before being placed under vacuum. The beam splitter and edge mirror are at the top of the picture. The shield covering the skimmers is visible on the right hand side. The atom beam passes through the center of the four magnetic coils and the penta and right angle prisms direct the laser beams to intersect the atoms. The lens on the left is positioned to collimate the fluorescence of the atom beam.



Figure 4.2 Skimmers used to define the direction and divergence of the atom beam.

marily, if the skimmers, seen in figure 4.2, are not aligned well to the center and axis of the oven, then a significantly lower flux of atoms in the beam may result. Of less importance is the fact that keeping the atom beam parallel to the table facilitates alignment of the rest of the interferometer.

In order to align the skimmers to the center of the oven, a He Ne laser was centered on the calcium crucible in the oven. The laser's height and direction were adjusted until the laser exactly overlapped the desired path of the atom beam. The skimmers height and position was then set based on the path of the laser beam.

Since calcium exiting the oven is directed randomly over a solid 180° angle, an aluminum box was placed around the skimmers to prevent hot calcium, which missed the skimmers, from coating and destroying the optics. A hole in the box larger than the atom beam, but still smaller than the skimmers themselves, allows the beam to propagate through. The shield can be seen at the right hand edge of figure 4.1.

4.2 Alignment of the Laser Beams

The alignment of the 657 nm laser through the interferometer begins with the spatial filter. Though not seen in figure 4.1, the filter sits inside the vacuum on the same bread board as the rest of the optics. This ensures precision when realigning the laser to the interferometer once the chamber is brought down to vacuum. The fourth beam in figure 4.1 is passed back out of the vacuum chamber to aid in aligning the laser to the spatial filter when the chamber is under vacuum.

After the spatial filter, the laser is sent to the beam splitter seen at the top of figure 4.1. Careful alignment of the laser to the beam splitter is required because the relative intensities of the resultant beams depend on the angle and polarization of the incident light. This alignment was done by measuring the intensities of the resulting laser beams while scanning the angle of the incident beam. When the angle is correctly set, the two beams differ from each other by less than 5% in intensity.

An edge mirror reflects these two beams to intersect the atom beam. Since all remaining horizontal alignment is done by virtue of prisms, the angle of the edge mirror determines whether or not the laser beams intersect the atoms orthogonally. The mirror is currently aligned so that the laser beams intersect the atoms to within about $25 \mu\text{rad}$ from perpendicular, which corresponds to a Doppler shift of no more than 28 kHz. When compared with the 75 kHz Fourier transform broadening an atom at $\sim 750 \text{ m/s}$ sees when passing through a 1 cm diameter beam of light, 28 kHz is small enough to ensure that if an atom is resonant with one beam, it will be resonant with all four laser beams.

The vertical alignment and horizontal beam spacing was then done by hand. A ruler was made from aluminum with holes at the height of the atom beam spaced by 10 cm. The penta and right angle prisms were then placed to align the four laser

beams with the ruler. To ensure that the four beams lay within the same plane to within .5 mm over about a 20 cm distance, the ruler was placed at several points along the beam paths.

4.3 Alignment of the Fluorescence Probe

The fluorescence probe consists of a laser tuned to the resonance transition (see figure 2.2), and it intersects the atom beam in the same plane as the interaction laser. The probe enters through a viewport on the vacuum chamber and is shown to the left in figure 4.1. A short focal length lens (1 inch focal length) focuses the fluorescence from the atoms up and out of the vacuum chamber. To align the lens, an LED was mounted at the height of the atom beam. The lens height was then adjusted to collimate the light from the LED. Collimated fluorescence will be focused onto a photo-multiplier tube by a second lens external to the vacuum.

Chapter 5

Low-noise, High-speed Electronics

The required stability of the 657 nm laser could not be achieved with off the shelf laser diode drivers and lock-circuitry. However, we were able to find integrated chips with the desired requirements for our circuits. These low-noise high-speed electronic components allowed us to construct the three main circuits of the laser: the photodiode homodyne detector, the proportional-integral-differential (PID) lock-circuit, and the diode laser current driver.

Figure 5.1 shows a functional diagram of how the electronics are used to lock the laser to the high finesse cavity. Light produced by the 657 nm ECDL is modulated by an EOM to add frequency side bands to the light. The light is then coupled to an optical cavity and the light reflected from the cavity is detected by a photodiode. The beat note generated by the carrier sidebands is demodulated to generate a Pound-Drever-Hall [54] error signal, which can be used to precisely lock the laser to the cavity. A scan-balancing PID controller uses the error signal to adjust the current going to the ECDL, the ECDL's cavity length, and the diffraction grating's angle to keep the laser light resonant with the high finesse cavity.

Since large bandwidth, low noise, and phase shift are extremely important in the




Figure 5.1 Functional diagram of the laser electronics. Light produced by the 657 nm ECDL is modulated by an EOM to add frequency side bands to the light. The homodyne detector uses the EOM frequency to demodulate the light reflected off the cavity and create an error signal. This error signal is then sent to a scan-balancing PID controller which keeps the ECDL locked to the high finesse cavity.

laser servo system, these circuits were designed to work together to keep cable lengths short. All three systems are compact and sit inside the same mechanically isolated box as the laser head and optics. In fact, the PID controller and laser current driver are housed together and share a common power supply. The homodyne detector is separate for the purpose of alignment and is connected to the PID controller with only a 6 inch cable. Aside from designing all three circuits to minimize delay time between them, each circuit is individually optimized for speed and low noise.

5.1 Photodiode Homodyne Detector

A homodyne detector is a device that is used to detect a modulated signal. This is done by mixing the modulated signal with a sine wave, known as the local oscillator, of the same frequency as the modulation. The phase of the local oscillator is chosen such that, when mixed with the signal from the optical cavity, the direct current (DC) portion of the mixed signal is zero when the laser is on resonance with the optical cavity.

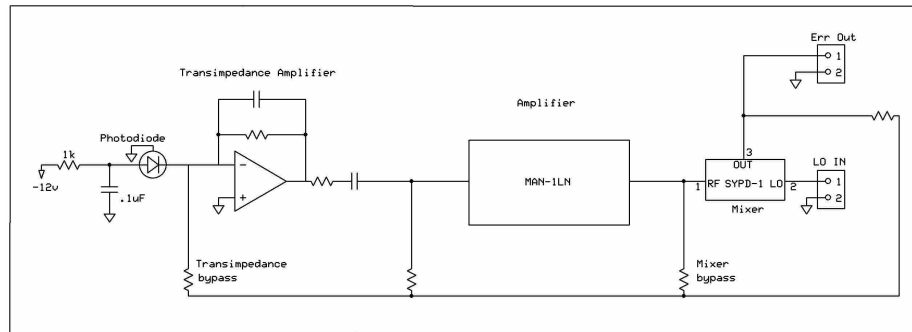


Figure 5.2 Schematic of the photodiode homodyne detector. A transimpedance amplifier converts the current to a voltage signal which is further amplified before it is sent to a mixer. The unit is designed for use as a simple photodiode, an amplified photodiode, or a homodyne detector which adds to its value as a lab standard.

Figure 5.2 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 speed. Then the signal from the photodiode is converted to a voltage by a transimpedance amplifier. The signal is further boosted by a monolithic amplifier before being sent to the mixer. A local oscillator signal from off the board is also sent the mixer, and the mixer output is sent to the lock-circuit.

At every stage in the circuit a pad for a zero ohm bypass resistor has been placed which allows for the board to be used as a photodiode, amplified photodiode, or a homodyne detector depending on the experiment. This is to allow us to use the same design for other purposes in other experiments. For this application the circuit was used only as a homodyne detector.

The physical layout of the printed circuit board (PCB) for the homodyne detector is shown in figure 5.3. It is entirely packaged on a 3.5 by 5 inch PCB. This image is the top side of the board, which contains most of the chips and all of the signal path. The signal originates at the photo-diode's pad in the upper left corner of the board. It is then passed through a transimpedance amplifier constructed with a single




Figure 5.3 Homodyne detector PCB layout. The pad for the photo-diode is shown in the upper left corner of the image circled for easy identification. The signal then travels to the right across the board through a transimpedance amplifier, an additional amplifier, and ends with a mixer. The mixer takes a local oscillator signal from the frequency generator driving the EOM to demodulate the light, and then sends the resulting signal off the board.

AD8099 op-amp, and then on to a monolithic amplifier (MAN-1LN). The MAN-1LN then sends the signal to the mixer (SYPD-1). The connector in the bottom right hand corner is for power, and the pads along the bottom of the board are for filtering and regulating the power.

5.1.1 Performance of the Homodyne Detector

The response of the homodyne detector is limited by performance of its individual chips. The photo-diode is a FSD-010 chosen for its small capacitance and high speed. It is reverse biased with -12 V to further decrease the capacitance and increase the speed of the diode. The photo-diode signal is converted to a voltage via an AD8099.

A MAN-1LN then boosts the signal to the level needed by the SYPD-1. Most of the components performed as specified. However, we found that the MAN-1LN's noise was greater than specified, and its gain was not flat across its frequency range in our circuit.

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. The total circuit response currently rolls off at 100 MHz, which sets the limit for the frequency of the incoming modulated signal and the local oscillator.

5.1.2 High Speed Op-amps

Only a few high-speed op-amps exist that meet the requirements for the homodyne detector. Currently, an AD8099 op-amp is being used because of its 3.2 GHz gain bandwidth product (GBP) and low noise figure. However, the AD8099 requires precise compensation capacitance, and it cannot drive capacitive loads larger than a few pico-Farads. It also has a centered ground pin that is particularly annoying if the chip needs to be replaced (when taking into account the time involved to do this it is more efficient to simply construct a new board). We recently started experimenting with a second op-amp, the OPA657. This chip still boasts a 1.6 GHz GBP, does not require compensating capacitance, can drive capacitive loads, has lower bias current, and it does not have a centered ground pin making it a better choice for our application.

When building a device using op-amps such as the AD8099 and OPA657, it is extremely important to give careful consideration to the load requirements for the chip. The low noise figure of these chips makes it tempting to use several gain stages of the same chip to amplify the signal instead of a monolithic amplifier. However, AC

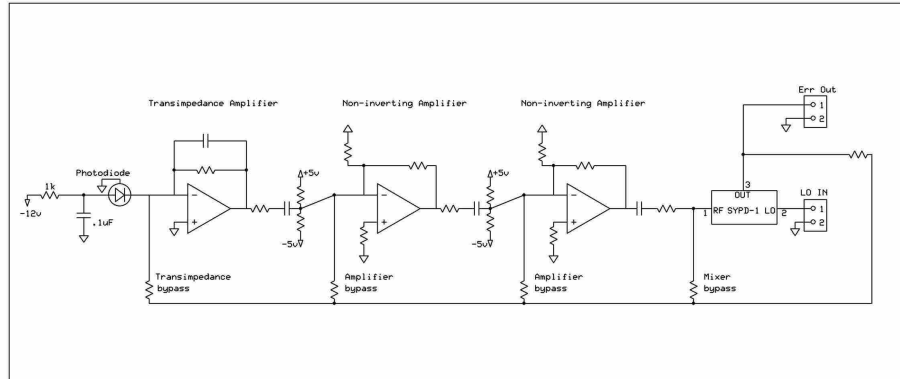


Figure 5.4 Schematic of the next generation photodiode homodyne detector which replaces the MAN-1LN monolithic amplifier with op-amps. This design has lower noise and higher bandwidth in accordance with the op-amps used.

coupling between chips must be carefully balanced, and the common and differential mode input impedances of the following op-amp must be taken into account to avoid undesirable oscillations of the circuit, making it more difficult to implement such a design.

5.2 Future Homodyne Detector

A new version of the board with higher bandwidth and a lower noise figure has nearly been completed. As shown in figure 5.4, the AD8099 is replaced by an AD8015 transimpedance amplifier with fixed internal gain and the MAN-1LN is replaced by two OPA657 op-amps daisy chained to the AD8015. The benefit of using the AD8015 is that it has an integrated feedback resistor, which reduces the size of the feedback loop and the effects of stray capacitance and inductance. Replacing the MAN-1LN with two OPA657's should lower the noise of the device; however, each op-amp requires a minimum gain in order to run stably, which puts the overall signal at a level that could damage the mixer. The difficulty in finishing this new detector seems to arise from an instability in the gain stages caused by a set of diodes placed on the output




Figure 5.5 Future homodyne detector PCB layout. The pad for the photo-diode is shown in the upper left corner of the image. The signal then travels to the right across the board through a transimpedance amplifier, two additional amplification stages, and ends with a mixer. The mixer takes a local oscillator signal from the frequency generator driving the EOM to demodulate the light.

in order to protect the mixer.

An image of the PCB for the new detector is shown in figure 5.5. In proto-typing and building the homodyne detector it was impossible to use an oscilloscope probe to monitor the signal. The capacitance and impedance of the probe caused the circuit to oscillate wildly at any stage. However, by testing each chip on the board in succession by hardwiring it through the bypass traces we were able to troubleshoot the circuit. This method was also difficult at times due to the sensitivity of the chips to load capacitance, as well as RF pick-up in the coaxial cable connecting the board to an oscilloscope.

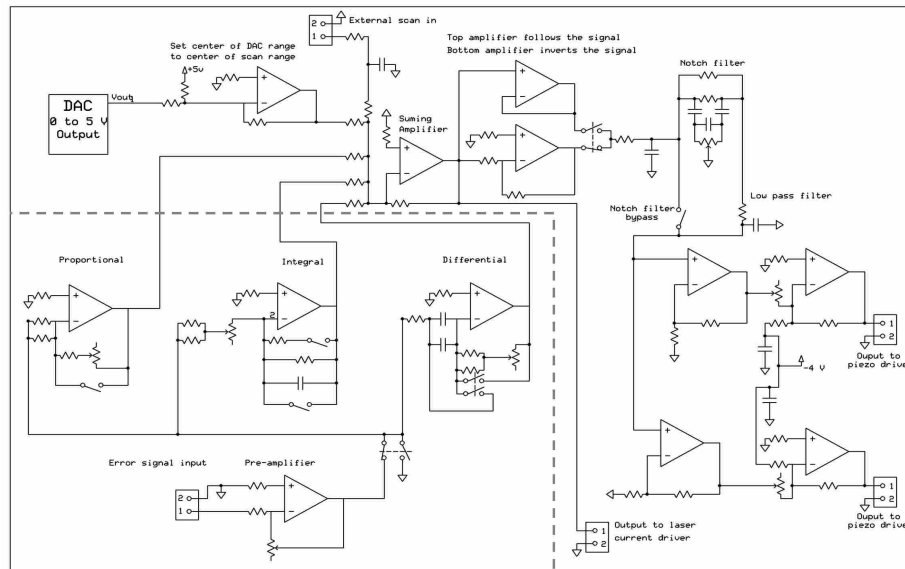


Figure 5.6 Schematic of the scan-balancing lock-circuit used in the experiment. The area enclosed by the dashed line represents the PID controller. The circuit can be scanned by use of a digital to analog converter or through the input of an external function generator.

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 whose linewidth is already smaller than the linewidth of the cavity it is being locked to, this works well. However, because our laser has a linewidth much larger than the optical cavity we desire as fast a lock as possible to keep the linewidth of the laser centered on the cavity resonance. To tighten the lock to our cavity, 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 lower left hand box of figure 5.6. It has a proportional, integral, and differential stage whose gains are tuned to optimize the response of the servo system. However, the bandwidth of the circuit is maximized by the use of high-speed low-noise surface mount chips and good PCB design. The feedback loop on each amplifier has been

minimized to reduce stray capacitance and inductance. To further minimize inductance, the gains on each stage are set with 1 turn trimpots. Short trace lengths between chips also helps to improve the circuit performance.

What makes this controller unique is that it can be digitally controlled. A separate unit consisting of a programmed microcontroller 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 relock the laser if the lock is lost.

Figure 5.7 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. The op-amps used in our application are AD8671 and AD8672. The DAC is an AD5541 and all the switches are contained within an ADG453.

5.3.1 PID Processing

The dashed box enclosing the lower left corner of figure 5.6 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




Figure 5.7 Image of the PCB used for the scan-balancing lock-circuit. The error signal enters through the connector labeled “ErrIn”. It is then amplified by the chip labeled “Trans” and passed to the proportional (“Prop”), the integral (“Int”), and the differential (“Dif”) stages. These stages along with signals from the DAC (“DAC”) and the connector “ScanIn” are then combined by the summing amplifier “Sum”. The output of this amplifier goes directly to the laser current driver through the connector “ModOut”. This signal is also buffered by the op-amp labeled “PInv” before passing through the notch filter (“filt”). The signal is then processed by the op-amps “PztGn” and “PztSum” before leaving through the connectors labeled “Piezo” to the piezo driver. All of the switches in figure 5.6 are contained in the chip labeled “DigitalSw”. The op-amps “DAC Quad” and “IsoSw” buffer the digital signals to the DAC and the digital switch which enter through the connectors “DAC” and “Switch” respectively.

separate resistors to a summing amplifier. The ratio of the respective input resistor after each stage to the feedback resistor for the summing amplifier determines the overall gain for each stage.

5.3.2 Laser Frequency Scan Function

The laser is scanned by either a DAC or through an external function generator as shown in the upper left portion of figure 5.6. The DAC puts out a value between 0 and 5 volts, which is shifted to -2.5 volts to 2.5 volts 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. It then sends the combined signal directly to the modulation input of the laser current driver.

5.3.3 Scan-balancer Function

The controller contains a scan-balancing circuit as shown on the right hand side of figure 5.6. This part of the circuit helps to keep the laser diode from mode-hopping while it is being scanned to the correct wavelength. 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. Therefore, slightly different gains are needed for controlling each side of the piezo mount. The scan-balancing circuit consists of two op-amps which control the two piezo-electric mounts on the laser's extended cavity diffraction grating. The

scan-balancer also contains a very selective notch filter which functions only when the circuit is locking the laser to prevent a mechanical resonance in the diffraction grating's mount from adding noise to the servo.

5.4 Diode Laser Current Controller

The goal to stabilize the 657 nm laser to a sub-kHz linewidth, as described in section 3.4, places some requirements on laser's current driver. First, the laser needs to have the current stabilized to within a few tens of nano-amperes rms over relevant time scales. This is the level at which the noise inherent to the driver does not widen the linewidth of the laser light beyond our desired linewidth. Second, a modulation input with a high bandwidth is needed. A high bandwidth modulation input prevents the current driver from limiting the bandwidth of the servo response, allowing us to avoid modulating the laser diode current directly. To satisfy the first requirement we have constructed a current driver with noise and stability as good or better than any commercially available driver we are aware of. The second requirement can be satisfied by a modulation input with a 3 dB point at or above 5 MHz.

Our design builds upon the Hall-Libbrecht design [56], 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. To satisfy the low noise requirement, however, 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.

The use of a DAC coupled with a microcontroller allows 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. The DAC provides more accurate control of the current by referencing each current set point with a 16-bit number. This also makes it possible to return to a previous current set point with extremely high precision.

The modulation input design of the original Hall-Libbrecht driver already provided a modulation input with enough bandwidth for our application. However, we implemented several smaller design changes. Most notably, our design uses surface mount technology to reduce stray capacitance and stray inductance.

5.4.1 Design and Construction

In figure 5.8 we have broken down the schematic of the circuit into several sections. Each section performs an important function and we will address each separately. While this design could be easily modified for an anode grounded supply, only the cathode grounded configuration is described here. An image of the PCB designed for the controller is shown in figure 5.9.

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 $I = (V_{reg} - V_{set})/R_s$.

As in the design in [56], an inductor is placed after the MOSFET to limit high-

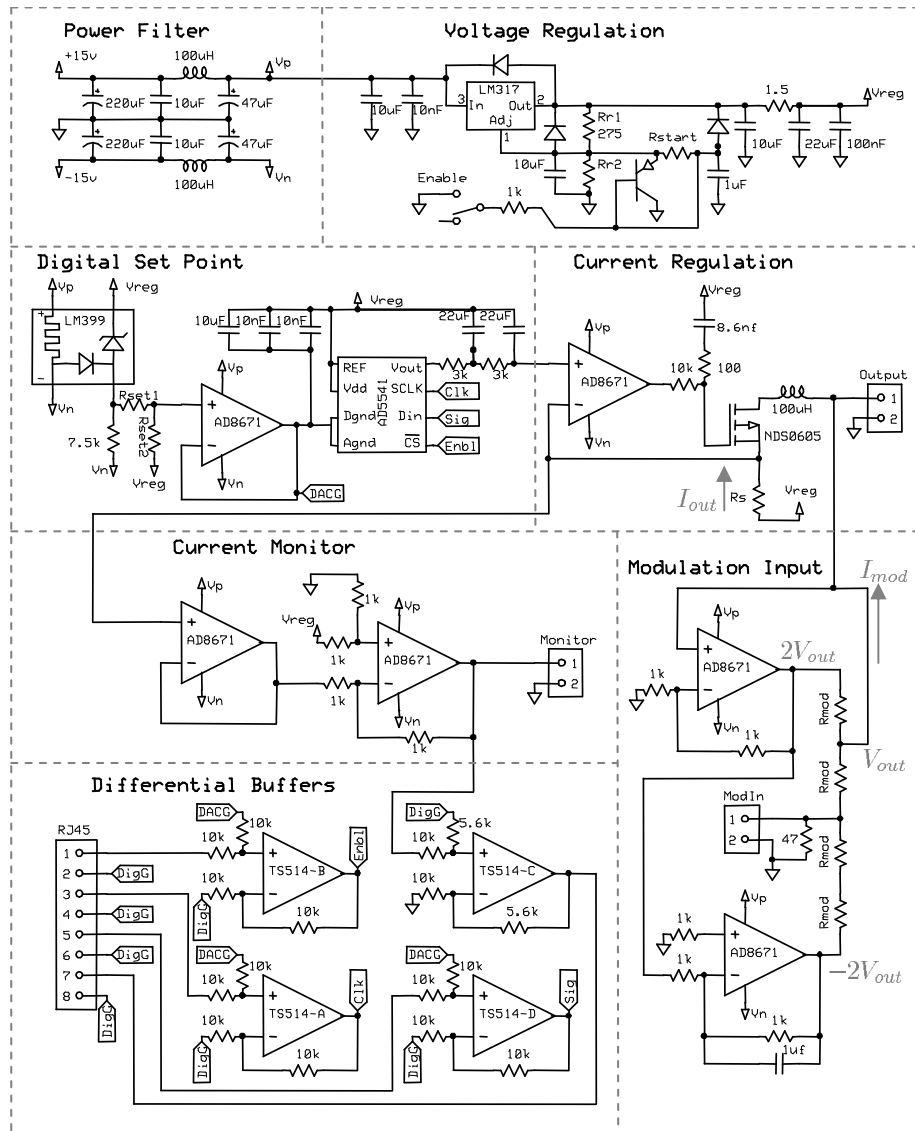


Figure 5.8 Schematic of the laser diode current driver circuit. The circuit is divided into separate functions for clarification and easy referencing.




Figure 5.9 Current driver PCB. 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 5.8. The center portion of the board is the “Modulation Input” section of figure 5.8. 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 figure 5.8. The resistor pads labeled “SMR3D 100R” comprise the precision resistor “ R_s ” in figure 5.8. The op-amp “Iset” controls the voltage on the gate of the MOSFET “Qfb”, which regulates the current through the resistors. The LM399 seen on the left hand side of the board and the op-amp labeled “DACref” create the 5 V reference for the DAC, seen just below the center of the board. 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.

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 resistor R_s is a key component that was considered carefully — stability of the current driver can be no better than the stability of this resistor. For this component we used a Vishay SMR3D precision resistor. These parts are surface-mount components for reduced noise pick-up, and have a low temperature coefficient of just 2 ppm/°C. The value for R_s is 50 ohms. This value for the resistance provides 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 the 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 [57]. Similar to the LT1028 used in the original Hall design [58], this op-amp has voltage noise levels comparable to the Johnson noise of a 50 ohm 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 decoupling 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 regulated with an LM317 adjustable voltage regulator to generate an extremely stable voltage, V_{reg} . The regulated output level of the LM317 is determined by a voltage divider created by the resistors labeled R_{r1} and R_{r2} in Fig. 5.8. The equation for calculating the output voltage of the regulator is [59]:

$$V_{reg} = 1.25 * (1 + R_{r2}/R_{r1}).$$

Soft Start Feature

The circuitry surrounding the LM317 in the “voltage regulation” section is part of a slow start mechanism shown in [59]. 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 μ F 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 $R_{start} = 1 \text{ k}\Omega$ 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 volts. 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 output voltage is V_{reg} .

To bias the DAC’s range, the chip’s “ground” was provided by a stabilized reference voltage, labeled V_{DACG} in figure 5.8, that is 5 Volts lower than V_{reg} . To stabilize this voltage difference, a LM399 zener diode was used to create an extremely stable 6.95 Volt drop from V_{reg} . Precision low temperature coefficient resistors, R_{set1} and R_{set2} , were then used to make a voltage divider to raise the voltage to the required 5 V difference.

Any relative drift between the voltage output of the DAC and V_{reg} will result in relative drift between the voltage output of the op-amp in the “current regulation” section of figure 5.8 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 and stability 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 5.8

would disappear. The sweeper of the potentiometer would be connected to the 3 k Ω 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 5.8 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 k Ω 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

$$I_{mod} = \frac{V_{out}}{R_{mod}} + \frac{V_{mod} - V_{out}}{R_{mod}} = \frac{V_{mod}}{R_{mod}}$$

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 5.8 below the first. It’s output is labeled “-2Vout”. This op-amp produces a voltage that exactly balances the output of the first op-amp, which is also separated from the “ModIn” connector by two precision 1 k Ω resistors labeled “Rmod”. This ensures that the “ModIn” connector voltage will be 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

$$I_{mod_{hf}} = \frac{V_{mod}}{R_{mod}} * \left(1 - \frac{2 * Z}{R_{mod}}\right) \quad (5.1)$$

where Z is the impedance of the laser diode. Note that in this limit the modulation current “ $I_{mod_{hf}}$ ” is not limited by the bandwidth of the op-amps.

Importance of Component Selection

Unlike the design in [56], 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 its 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 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 5.8 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 the microprocessor, a scope, or a multi-meter. The use of differential amplifiers effectively isolate the current driver from the ground and signals 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 5.8 serve to reference signals from the ground of the DAC (which is not board ground) to the ground of the microprocessor PCB. Three digital lines are passed to the DAC from the microprocessor to control the current, and one analog line is passed from the current driver to the microprocessor to reference the voltage drop across R_s . These op-amps are not directly connected to the driver’s output and, therefore, do not have to be as carefully selected.




Figure 5.10 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.

5.4.2 Results and Discussion

Accuracy and Repeatability

The accuracy and repeatability of setting the current digitally is shown in figure 5.10. When setting the current, the accuracy of the DAC is rated to within $2 \mu\text{A}$. As seen in the upper graph of figure 5.10, 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 spanning the range of the driver were returned to 10 times each from a specified point. As shown in figure 5.10, the current can be reset to the same point from a different value to within a few $n\text{A}$.

Noise and Stability

To characterize both the long term stability and the high frequency noise of the current driver we placed a 2.5 ohm precision resistor in series with the laser diode




Figure 5.11 Spectral density of the noise of the current driver. The black line represents the spectral density of the noise of the current signal when the diode is turned on at 74.5 mA. The thick gray line is the spectral density of the noise of the measurement apparatus when the current driver is not connected to the laser diode. The thin gray line is the spectral density of the noise of the original Hall-Libbrecht driver extracted from the data presented in the original paper.

and measured the voltage drop across the resistor. From this data we extracted the current spectral noise density shown in figure 5.11. For parts of the spectrum the noise of the driver measured at the noise floor of the test measurement setup. Overall, the spectral noise density of the driver is about an order of magnitude better than the original Hall-Libbrecht design.

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 cur-

rent was determined by measuring the voltage drop over a 10 ohm precision resistor.

Figure 5.12 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 about a few parts per million (1.7 ppm/C). Given that the set resistor R_s in figure 5.8 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.

Transient Response to Load Impedance Variation

The transient response of the current driver was evaluated by putting a 10 ohm resistor in series with a VZN1206 MOSFET and controlling the transistor gate to vary the overall impedance by $\sim 3\%$ with a step function. The black line in figure 5.13 is the current through the resistor measured by taking 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 50Ω resistor and comparing the amplitude and phase of the modulation signal on the current driver 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




Figure 5.12 (A) The current output of the driver through a laser diode. (B) Corresponding temperature measurements taken at the current set resistor R_s and the DAC. (C) Corresponding temperature measurements taken at the box housing the driver.




Figure 5.13 The transient response of the current driver to a total load change of roughly 3%. The black line is the current through a 10 ohm resistor, and the grey line is the impedance of the load.




Figure 5.14 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 high frequency response. The dot-dash line is the theoretical low frequency amplitude response, and the dashed line is the theoretical high frequency response.

Bode plot in figure 5.14. 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.

5.4.3 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 could easily be implemented as a lab standard. This performance is mainly due to the advances in technology of surface mount ICs and in their availability in small quantities. An inexpensive way to undertake surface mount soldering of PCBs in low quantities is described in appendix A.

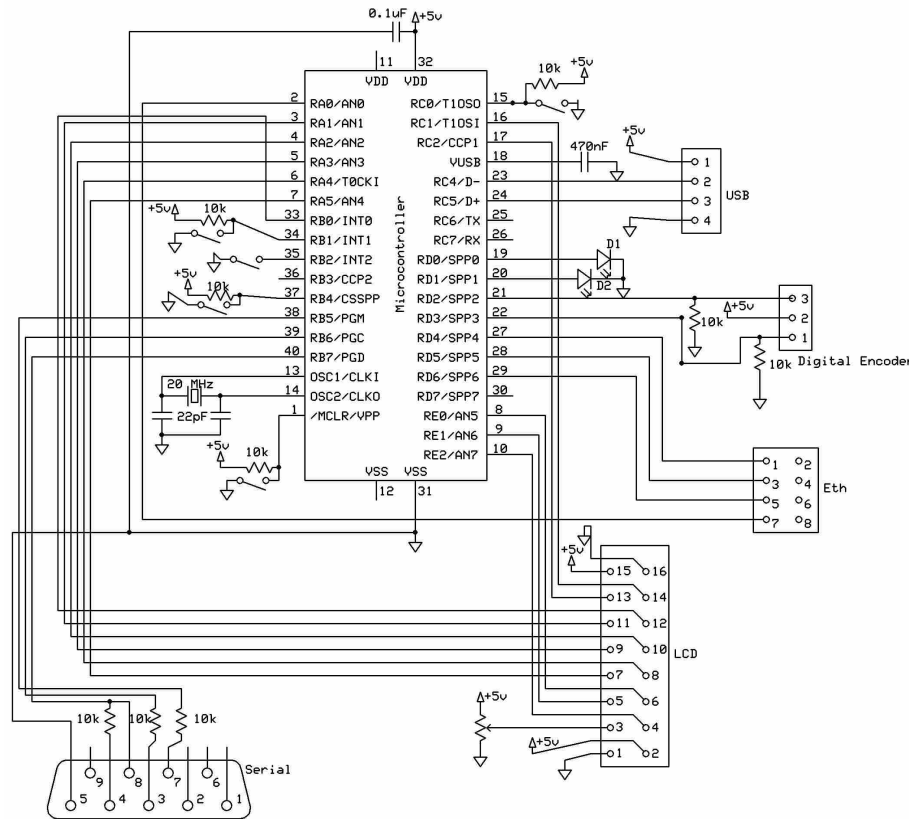


Figure 5.15 Schematic for using the microcontroller with the laser current driver. The serial connector is used to upload a small piece of software, called a bootloader, onto the microprocessor. Once the bootloader is installed, the USB connector is used for uploading the program we wish to run on the microprocessor. The connector labeled “Eth” is an RJ45 ethernet jack used for putting signals out to and reading signals back from the laser current driver. The LCD is for displaying information from the microcontroller.

5.5 Microprocessor

For control of both the current laser driver and the scan-balancing lock-circuit, we used a separate PIC18F4550 microcontroller from Microchip Technologies for each one. The PIC18F4550 is a USB programmable microprocessor with 35 pins which can be used as inputs or outputs. The PIC18F4550 also includes a built-in analog to digital converter (ADC). This chip is easy to update and can support many different functions, which makes it ideal for remotely controlling the current driver and the




Figure 5.16 Top layer of the PCB for the microcontroller.

lock-circuit. A previous member of the group, Marshall van Zijll, designed a circuit for the microcontroller shown in figure 5.15, as well as a PCB [60]. Images of the both sides of the PCB appear in figures 5.16 and 5.17.

Marshall also programmed the microcontroller [60]. The microcontroller reads in the signal from a digital encoder, and subsequently outputs a value to the DAC on either the lock-circuit or the current driver. To monitor the output on the current driver, the on-chip ADC reads back the voltage drop across R_s . The microprocessor was also programmed to display this signal on an LCD display along with settings being used by the microcontroller. For more information on the microprocessor and its circuit please consult [60].




Figure 5.17 Bottom layer of the PCB for the microcontroller.

Chapter 6

Passive Stabilization of the Optical Cavity and the 657 nm Laser

6.1 High-finesse Optical Cavity

Noise on the optical cavity to which the laser is locked will cause instability in the lock. With the best lock circuit in the world we would not be able to correct for thermal and mechanical noise on the optical cavity itself, which can arise from several sources. First, if any gas is present between the cavity mirrors, the intracavity laser intensity will create variations in the index of refraction. Also, if the incident laser intensity is too great the cavity mirrors will heat up, creating thermal instability. Finally, mechanical vibrations can cause noise that will destabilize the laser lock. By stabilizing the cavity passively; however, we can minimize such problems.

This passive stabilization begins with the construction of the cavity and its housing. Purchased from Research Electro Optics, the cavity was constructed of two ultra-low expansion (ULE) quartz mirrors which are set on an ULE quartz spacer. To support the cavity, we machined an Invar bar with grooves to hold four sapphire

balls on which the spacer rests. The main purpose of these measures is to minimize the effects of thermal expansion on the cavity. To isolate the cavity from atmospheric effects, these pieces sit inside a vacuum chamber.

Mechanical stabilization is achieved by placing the vacuum chamber in a 1 inch thick solid aluminum box lined with leaded foam. The box then rests on Sorbothane spacers on top of a floating optics table. This level of passive stabilization seems to be sufficient for the meantime. However, in the future this might be further stabilized with a temperature controller in an effort to further reduce the frequency drift of the laser.

6.2 Isolation of Electronics and Mirror Mounts

Mechanical and thermal vibrations in the optics preceding the cavity are written onto the laser light as noise, which then has to be dealt with by the servo system. To minimize this source of noise, the laser head, control and feedback electronics, and optics are housed in a second 1 inch thick solid aluminum box. This box also rests on Sorbothane spacers and it is mechanically joined to the box enclosing the optical cavity to prevent the optics from moving relative to the cavity. Putting the electronics inside the box helps isolate them from sources of electro-magnetic noise in the lab, and has the added benefit of keeping cable lengths as short as possible.

In an unfortunate twist of fate, our lab happens to lie directly below the transformers for the entire Eyring Science Center. The transformers couple noise into the lab very efficiently through the concrete floor and walls. It so happens that the lab itself has an acoustical resonance frequency at 540 Hz (a multiple of 60). Incredibly enough, the aluminum box housing the laser and electronics also has a resonance frequency near 540 Hz. Left as is, the servo electronics oscillate heavily at this frequency

when trying to lock the laser. Placing leaded foam inside the empty areas of the box absorbed some of the noise, but not enough.

Since we were dealing with an acoustical noise problem, we purchased insulation to place on the cement walls of the lab. Absorptive material should be placed to overlap the high pressure points of the acoustical wave where the velocity is greatest, rather than the nodes where the velocity is null. However, despite the fact that the wall is a nodal point, it is the easiest surface to mount to. Since a quarter wavelength is about $6 \frac{1}{5}$ inches for our resonance frequency, placing $8 \frac{1}{2}$ inch thick fiberglass insulation on the walls of the lab effectively killed the resonance.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

While we were not able to perform any measurements with the interferometer before the deadline of this thesis, we have constructed the apparatus and are waiting for favorable conditions to run the experiment (see section 7.2). We were able to obtain the necessary optics to implement our experimental design, as well as construct the necessary electronics and laser system for using the calcium intercombination line as a clock transition in an interferometer. Having the optics anti-reflective coated actually changed their specifications to outside the desired range, but they are still precise enough to run the interferometer as outlined. Also, the laser driving the atomic transition in the atoms is not as narrow as planned due to the properties of the optical cavity it is locked to. Still, the laser linewidth should be narrow enough for the interferometer to function properly. While these setbacks somewhat diminish the precision of the interferometer, they should not prevent its operation. They also allow us the opportunity to use what we have learned in optimizing the experiment in the future.

7.2 Current Status of the Interferometer

Data from the interferometer would have been taken by this point if the pressure in the vacuum chamber had reached the appropriate level. The chamber was initially evacuated and baked for several weeks without reaching the necessary level, and it was feared that the lid to the chamber had been scratched. However, it was instead discovered that beneath the Viton gasket sealing the chamber grease, dust, and metal shavings left over from the initial manufacture had collected. We also discovered wood splinters, plastic pieces, tissue scrap, posts, screws, and dust under the bread board inside the chamber. The entire chamber was cleaned thoroughly with methanol and reassembled. Despite this, the vacuum has still not reached the necessary pressure level several weeks later, and we are currently troubleshooting the problem.

7.3 Optimization of the Interferometer

Now that the first version of the interferometer has been constructed, several things can be done to improve and optimize it. The addition of our own blue laser will be a major improvement to the interferometer. This will allow us to transversely cool the atom beam, adding a couple hundred times more atoms to the beam density. Having a dedicated laser for the fluorescence probe will also alleviate any problems that might arise from running two experiments at once.

We will also construct a new oven and a precision temperature controller for generating the atom beam. The new oven will be more compact and will be designed to create a dual-species thermal beam in preparation for future experiments. The real benefit of the new oven will be the temperature controller that we will design for it. Inherent to the nature of a thermal beam, the flux is temperature dependent. Variations or drifts in the temperature can cause variations in the atomic density of

the beam which will affect the interferometer read out. Our current controller only holds the temperature steady to within a couple of degrees Kelvin. The new controller will hold the temperature steady to within a few hundred milli-Kelvin.

Other improvements to the thermal beam include upgrading the collimation apertures. The current apertures are 2 mm wide unheated skimmers placed roughly 10 cm apart and aligned to the oven with a laser. To replace these, we have started constructing a small stage which will house two heated 5 μm slits. The stage will keep the slit spacing at exactly 10 cm and allow for easy alignment of the slits to each other and also to gravity. Whereas the current beam divergence is 9.8 milli-radians, the beam divergence with the new slits will be about 100 micro-radians. The new slits will decrease the beam flux, however, by more precisely defining the beam divergence and direction. This will further reduce Doppler errors and improve the resolution of the interferometer's fringes.

Better reduction of rotational and acceleration sensitivity will also be achieved. Both of these shifts can be better stabilized with better mechanical isolation. Also, by aligning the interaction laser beams perpendicularly to gravity at the milli-radian level, frequency shifts due to acceleration should be further stabilized to a part in 10^{16} .

The 657 nm laser will also be optimized by taking measures to further narrow the linewidth. This can be done with temperature control of the optical cavity, as well as by completing the next generation photodiode homodyne detector, which may allow us to modulate the laser light as fast as 100 MHz. Since measurement of the optical cavity linewidth revealed it to be greater than expected, we believe that the cavity mirrors may have been contaminated by the ion pump serving the chamber. To restore the Q of the cavity we will have these mirrors professionally cleaned. The end goal is to create a Hertz-level laser well below the linewidth of the atomic transition.

Finally, we will double the length of the interferometer. Instead of the current 20 cm length, we will instead have a 40 cm length. This length will be optimal for the upper state lifetime of the calcium. Lengthening the interferometer will narrow the interference fringes [11] allowing for better resolution and frequency stability.

7.4 Construction of a Dual Species Interferometer

In the future, strontium will be added to the apparatus as seen in figure 2.3. The Strontium interferometer will exactly overlap the calcium interferometer realizing two interferometers in the space of one. The driving light field for the strontium, as well as the transverse cooling laser and fluorescence probe will be constructed in a similar manner to the lasers used for the calcium interferometer. A separate optical cavity will be used to stabilize the driving laser for the strontium as well.

The addition of strontium will allow the interferometer to make precision measurements of fundamental constants. The intercombination lines for strontium and calcium are 689 nm and 657 nm respectively. They are close enough that we will be able to compare the signals from each interferometer to make measurements of the fine structure constant in particular.

7.5 Future Measurements and Applications

The interferometer described in this thesis will be tested for sensitivity as a gyroscope and accelerometer, which measurements will provide a benchmark for optimizing and improving the apparatus. While the use of calcium will prevent it from competing with current atom interferometer gyroscopes, future versions modeled after this device will be constructed with a more appropriate choice of elements for those applications.

The dual species version of the apparatus will mainly be used to perform fundamental studies. It should be able to measure time variance in the fine structure constant to one part in 10^{16} per 8 months while the current measurement of the fractional variation of α stands at a couple parts in 10^{-15} per year [26–28, 61, 62]. The level of sensitivity of this device should also allow it to perform measurements of relativity.

Appendix A

Surface Mount Soldering in a Toaster Oven

Thanks to consumer markets like cellular phones and hand-held PDA's, integrated chip makers have developed many new high-speed, high band-width amplifiers to fuel the push for new technology. Due to the constant miniaturization of such devices most of these new chips are surface mount. While surface mount technology aids in the performance of such chips it also makes prototyping electronics for small laboratories time consuming and expensive. Table top wave-flow soldering stations start at \$6,000, and hand soldering surface mount parts is time consuming, messy, and can result in damage to chips. It is also true that building lab equipment from surface mount parts prevents the use of a standard prototyping board for easily changing chips and passive components of different values in and out during the development stage of building a component.

The problems and mistakes inherent to developing prototype boards for electronic components can be minimized by very careful review of circuit schematics and the use of SPICE models. SPICE programs are available for free from several chip man-

ufacturers and 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.

Several resources also exist for the design of the board layout and its production. We found ExpressPCB (<http://expresspcb.com>) well suited for our purposes. The website provides a free CAD program for design of both the circuit schematic as well as the board layout. The schematic can be linked to the board design file to help ensure traces connect the correct surface mount pads. Prices for the boards are reasonable and depend on the size of the board as well as the number of through holes that need to be drilled in the board.

The problematic issue of soldering surface mount chips without commercial technology is taken care of by the use of a common toaster oven. As far as we are aware, the idea of using a common table top oven for surface mount soldering was first put forth in the newsletter of the Seattle Robotics Society [63]. Their experience can be reviewed at http://www.seattlerobotics.org/encoder/200006/oven_art.htm. We discovered by our own experimentation that their method works without damage to components from heat. However, while the previous group insisted that this method is only made possible by the use of an expensive, water-soluble solder paste that requires refrigeration and has a short shelf life, we found that it can also be done by using ChipQuik's 63/37 w/flux solder paste, which is cheap, requires no refrigeration, and has a very long shelf life.

To solder surface mount chips to a PCB using solder paste, chips should be laid out on their respective pads first to ensure that all the chips are on the right pads in the right orientation. Then, using a 20 GA tip on the solder paste syringe, solder paste can be applied to individual pads while the component is momentarily set aside. Replacing the component on the solder paste is most easily done using tweezers for smaller components and fingers for larger components. Too much solder paste on a

single pad will create a solder bridge to an adjacent pad or the ground plane during the cooking process.¹This is especially true for chips with multiple adjacent pins and takes some practice. Components that are not surface mount should be hand soldered after the surface mount soldering process is complete.

Before placing the board in the toaster oven, make sure the oven is level and that the cooking rack inside the oven is also level. Using an Oster model 6232-015 1500 Watt toaster oven we found the following cooking times effective (no pre-heating):

- 250 °F for 4 minutes to bring the oven and board to a uniform base temperature
- 375 °F for 2 minutes to bring solder to it's melting point
- 450 °F for 30+ seconds to melt and bead solder
- Immediately turn off the oven and open the door when solder has beaded around **all** the pins.

As the board cooks, the solder should start to bead around the individual pins and it will become shiny near the end of the process. Monitor the board constantly at the 450° stage. When the solder has formed beads on all the components, turn off the oven and open the door. During the end of the process and when the board is finished cooking, one should be very careful not to bump the oven or tip the board until the solder has set to prevent chips from sliding off their pads and creating solder bridges to other chips, pads, and ground planes.

When the board has cooled, solder bridges and loose pins can be checked for under a bright light. Solder bridges and loose pins can be carefully corrected by a desoldering pump or surface mount solder wire and a hand-held iron respectively.

¹There is an exception for chips with multiple pins below .050 pitch. In this case, it is easier to lay down a bead of solder across all the pins. However, the bead should be thin since too much solder will still lead to solder bridges or even connecting all adjacent pins together.

Through practice of careful placement of the solder paste, both of these problems can be minimized.

It is possible that a small amount of flux from the solder may coat the board around the solder joints. This can be cleaned off with a commercially available flux cleaning solution or even just methanol, but it does not affect the electrical properties of the circuit as far as we have been able to tell. If a mistake has been made or a chip does need to be replaced, the use of ChipQuik desoldering alloy provides an easy and clean way to remove surface mount soldered chips by hand. We did not experiment with the effect of re-baking a board once surface mount chips had been soldered to it for the purpose of adding more components.

Bibliography

- [1] 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).
- [2] L. de Broglie, Ph.D. thesis, Université de Paris, 1924.
- [3] E. Esterman and O. Stern, “Diffraction of a sodium atom beam,” *Zeits. Phys.* **61**, 95 (1930).
- [4] N. F. Ramsey, “A molecular beam resonance method with separated oscillating fields,” *Phys. Rev.* **78**, 6 (1950).
- [5] N. F. Ramsey, “History of atomic and molecular standards of frequency and time,” *IEEE Transactions on Instrumentation and Measurement* **21**, 90–99 (1972).
- [6] N. Ramsey, “Precise measurements of time,” *Am. Sci.* **76**, 42 (1988).
- [7] W. M. Itano and N. F. Ramsey, “Accurate measurement of time,” *Scientific American* **269**, 56–65 (1993).
- [8] O. Carnal and J. Mlynek, “Young’s double slit experiment with atoms: a simple atom interferometer,” *Phys. Rev. Lett.* **66**, 2689 (1991).

-
- [9] 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).
- [10] 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).
- [11] *Atom Interferometry*, P. R. Berman, ed., (Academic Press, 1997).
- [12] T. L. Gustavson, P. Bouyer, and M. A. Kasevich, “Precision rotation measurements with an atom interferometer gyroscope,” *Phys. Rev. Lett.* **78**, 2046 (1997).
- [13] 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).
- [14] T. L. Gustavson, A. Landragin, and M. A. Kasevich, “Rotation sensing with a dual atom interferometer Sagnac gyroscope,” *Class. Quantum Grav.* **17**, 2385 (2000).
- [15] 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).
- [16] 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).
- [17] E. M. Rasel, M. K. Oberthaler, H. Batelaan, J. Schmiedmayer, and A. Zeilinger, “Atom wave interferometry with diffraction gratings of light,” *Phys. Rev. Lett.* **75**, 2633 (1995).

-
- [18] 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).
- [19] A. Peters, K. Y. Chung, and S. Chu, "Measurement of gravitational acceleration by dropping atoms," *Nature* **400**, 849 (1999).
- [20] A. Peters, K. Y. Chung, and S. Chu, "High-precision gravity measurements using atom interferometry," *Metrologia* **38**, 25 (2001).
- [21] 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).
- [22] 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).
- [23] 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).
- [24] J. D. Prestage, R. J. Tjoelker, and L. Maleki, "Atomic clocks and variations of the fine structure constant," *Phys. Rev. Lett.* **74**, 3511 (1995).
- [25] H. Marion *et al.*, "Search for variations of fundamental constants using atomic fountain clocks," *Phys. Rev. Lett.* **90**, 150801 (2003).
- [26] 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).
- [27] M. Fischer *et al.*, "New limits on the drift of fundamental constants from laboratory measurements," *Phys. Rev. Lett.* **92**, 2308021 (2004).

-
- [28] E. Peik, B. Lipphardt, H. Schnatz, T. Schneider, C. Tamm, and S. G. Karshenboim, “Frequency comparisons and absolute frequency measurements of $^{171}\text{Yb}^+$ single-ion optical frequency standards,” ArXiv:physics/0504101v1 (2005).
- [29] V. A. Dzuba and A. M. V. M. 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).
- [30] J. R. Torgerson, “Method for precision test of fine structure constant variation with optical frequency references,” ArXiv:physics/0012054v3 (2001).
- [31] R. M. Mansouri and R. U. Sexl, “A test of special relativity,” Gen. Relativ. Gravit. **8**, 497 (1977).
- [32] H. P. Robertson, “Postulate versus observation in the special theory of relativity,” Rev. Mod. Phys. **21**, 378 (1949).
- [33] R. J. Kennedy and E. M. Thorndike, “Experimental establishment of the relativity of time,” Phys. Rev. **42**, 400 (1932).
- [34] 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).
- [35] 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).
- [36] 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).

- [37] B. Neyenhuis, D. Christensen, and D. Durfee, “Testing non-classical theories of electromagnetism with ion interferometry,” *Phys. Rev. Lett.* **99**, 200401 (2007).
- [38] D. Christensen, B. Neyenhuis, and D. S. Durfee, “Numerical calculations of classical and non-classical electrostatic potentials,” *ArXiv:physics/0609128v1* (2006).
- [39] P. Kersten, F. Mensing, U. Sterr, and F. Riehle, “A transportable optical calcium frequency standard,” *Appl. Phys. B* **68**, 27 (1999).
- [40] R. L. Barger, J. C. Bergquist, T. C. English, and D. J. Glaze, “Resolution of photon-recoil structure of the 6573-Å calcium line in an atomic beam with optical Ramsey fringes,” *Appl. Phys. Lett.* **34**, 850 (1979).
- [41] F. Riehle, J. Ishikawa, and J. Helmcke, “Suppression of a recoil component in nonlinear Doppler-free spectroscopy,” *Phys. Rev. Lett.* **61**, 2092 (1988).
- [42] H. Schnatz, B. Lipphardt, J. Helmcke, F. Riehle, and G. Zinner, “First phase-coherent frequency measurement of visible radiation,” *Phys. Rev. Lett.* **76**, 18 (1996).
- [43] J. H. Shirley, W. D. Lee, and R. E. Drullinger, “Accuracy evaluation of the primary frequency standard NIST-7,” *Metrologia* **38**, 427 (2001).
- [44] D. S. Durfee, Y. Shaham, and M. Kasevich, “Long-term stability of an area-reversible atom-interferometer sagnac gyroscope,” *Phys. Rev. Lett.* **97**, 240801 (2006).
- [45] A. E. Leanhardt, A. P. Chikkatur, D. Kielpinski, Y. Shin, T. L. Gustavson, W. Ketterle, and D. E. Pritchard, “Propagation of Bose-Einstein condensates in a double-well potential,” *Phys. Rev. Lett.* **89**, 040401 (2002).

-
- [46] S. Yanagimach, K. Mizobuchi, and A. Morinaga, “Ramsey-Bordé atom interferometer having two arms with different Zeeman sublevels,” *Phys. Rev. A* **64**, 041601 (2001).
- [47] S. R. Jefferts *et al.*, “Accuracy evaluation of NIST-F1,” *Metrologia* **39**, 321 (2002).
- [48] Y. Sortais, S. Bize, C. Nicolas, A. Clairon, C. Salomon, and C. Williams, “Cold collision frequency shifts in a ^{87}Rb atomic fountain,” *Phys. Rev. Lett.* **85**, 3117 (2000).
- [49] R. Legere and K. Gibble, “Quantum scattering in a juggling atomic fountain,” *Phys. Rev. Lett.* **81**, 5780 (1998).
- [50] U. Sterr, T. Binnewies, C. Degenhardt, G. Wilpers, J. Helmcke, and F. Riehle, “Prospects of Doppler cooling on forbidden lines,” *J. Opt. Soc. Am. B* **20**, 985 (2003).
- [51] E. A. Curtis, C. W. Oates, and L. Hollberg, “Quenched narrow-line second- and third-stage laser cooling of ^{40}Ca ,” *J. Opt. Soc. Am. B* **20**, 977 (2003).
- [52] C. J. Bordé, “Atomic interferometry with internal state labeling,” *Phys. Lett. A* **140**, 10 (1989).
- [53] R. Olson, J. Paul, S. D. Bergeson, and D. S. Durfee, “Self-referenced prism deflection measurement schemes with micro-radian precision,” *Appl. Opt.* **44**, 4639 (2005).
- [54] 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).

- [55] 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).
- [56] K. G. Libbrecht and J. L. Hall, “A low-noise high-speed diode laser current controller,” *Rev. Sci. Inst.* **64**, 2133–2135 (1993).
- [57] A. D. Incorporated, “Precision, very low noise, low input bias current operation amplifiers: AD8671/AD8672/AD8674,” *Analog Devices Datasheet Catalog* (2005).
- [58] L. T. Corporation, “LT1028/LT1128 ultralow noise precision high speed op amps,” *Linear Technology Datasheet Catalog* (1992).
- [59] F. Semiconductor, “LM317 3-terminal adjustable regulator,” *Fairchild Semiconductor Datasheet Catalog* (2001).
- [60] M. van Zijll, Senior Thesis, Brigham Young University, 2007.
- [61] S. Bize *et al.*, “Testing the stability of fundamental constants with the $^{199}\text{Hg}^+$ single-ion optical clock,” *Phys. Rev. Lett.* **90**, 150802 (2003).
- [62] A. Cingöz, A. Lapierre, A.-T. Nguyen, N. Leefler, D. Budker, S. K. Lamoreaux, and J. R. Torgerson, “Limit on the temporal variation of the fine-structure constant using atomic dysprosium,” *ArXiv:physics/0609014v2* (2006).
- [63] K. Maxon, “Have you seen my new soldering iron?,” *Encoder: The Newsletter of the Seattle Robotics Society* 6 (2000).