DEVELOPING A NARROW LINEWIDTH 657 NM DIODE LASER FOR USE IN A CALCIUM ATOM INTERFEROMETER

by

Brian Neyenhuis

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

Bachelor of Science

Department of Physics and Astronomy
Brigham Young University
April 2006

BRIGHAM YOUNG UNIVERSITY

DEPARTMENT APPROVAL

of a senior thesis submitted by

Brian Neyenhuis

This thesis has been reviewed by the research advisor, research coordinator, and department chair and has been found to be satisfactory.

Date	Dallin S. Durfee, Advisor
Date	Jean-Francois Van Huele, Research Coordinator
Date	Scott D. Sommerfeldt, Chair

ABSTRACT

DEVELOPING A NARROW LINEWIDTH 657 NM DIODE LASER FOR USE IN A CALCIUM ATOM INTERFEROMETER

Brian Neyenhuis

Department of Physics and Astronomy

Bachelor of Science

I am designing a narrow linewidth 657 nm diode laser for use in an atom interferometer. I will discuss both passive and active stabilization of diode lasers as well as a new grating stabilization scheme developed in our lab with its advantages and disadvantages. I have constructed a high speed lock circuit with a bandwidth of 4 MHz that is used with the Pound Drever-Hall method to lock the diode laser to a cavity with a finesse of 30,000. The laser's current linewidth is approximately 3 kHz. Future work and a plan to achieve a Hz level linewidth are presented

ACKNOWLEDGMENTS

This work would not have been possible without the help of my advisor Dallin Durfee. Not only has he been a great resource of knowledge and guidance for this project, but he has been a great friend and mentor and has profoundly impacted my choice to continue on in atomic physics research.

I should also acknowledge help on various aspects of this project from Justin Paul, Rebecca Tang, Chris Erickson, Greg Doermann, John Ellsworth, Scott Daniel, Matt Washburn, and James Archibald.

Contents

List of Figures

Chapter 1

Introduction

1.1 Why Build an Ultra-stable Narrow Linewidth Laser?

The focus of this thesis is the construction, optimization, and testing of an ultra stable laser system. This laser is part of a larger project to construct an atom interferometer. Once constructed this interferometer will be used as an atomic clock, as well as an interferometer to test relativity at nonrelativistic speeds and measure changes in the fine structure constant.

Science depends on the ability to accurately test physical theories. As physical models become more refined, the predictions they make become increasingly subtle. To differentiate between conflicting theories and to gain further insight into how nature works we must find new ways to gather more detailed information about the world around us. For this reason precision measurement plays and will continue to play an important role in modern physics research.

Atom interferometry has played a key role in precision measurement. The small wavelengths of atoms allows atom interferometers to make measurements with very high precision and accuracy. Atom interferometers have uses in fundamental science: they can be used to measure the properties of atoms, test special relativity at non-relativistic speeds, measure the gravitational red shift, measure drifts in fundamental constants, test the Lense-Thirring effect (gravitational frame-dragging), and provide precision tests of physical theories in general. Atom interferometers can also be used as a ruler with which to measure time and force with unparalleled accuracy. They have applications in navigation as a gyroscope, and in the detection of underground structures by sensing changes in local gravity. Atom Interferometers can also serve as a high-precision frequency standard, which is of great importance in high speed communications and global positioning systems.

The higher frequencies of optical transitions allow us to obtain higher fractional precision in a given interrogation time. Ted Hänsch, 2005 Nobel Laureate, once said “If you want to find something that no one else has ever seen before, look in a place where no one else has looked.” [1] A new generation of optical frequency standards give us the ability to look where microwave frequency standards could not.

To realize these optical frequency standards and take advantage of their higher fractional precision, commercially-available laser systems must be improved considerably. The precision of an atom interferometer depends upon the linewidth of the laser driving the clock transition. Both the long and short term stability of the interferometer depend upon the stability of the laser. For these reason the laser can be thought of as the “heart” of the interferometer.

1.2 The BYU Calcium Atom Interferometer Project

To illustrate the importance of a stable narrow linewidth laser, I will briefly describe the proposed interferometer and some of its advantages. The proposed interferometer will use four $\pi/2$ laser pulses in a standard Ramsey-Bodé interferometry configuration [2] (see Fig. 1.1). These four pulses will be used to split and recombine the wavefunction of the atomic beam. The transition used is the calcium ^{1}S$_0$ to ^{3}P$_1$ intercombination line at 657 nm. After the two paths have recombined the percentage of the atoms in the ground state is measured to determine the relative phase between the two paths.

Our interferometer will use a thermal beam of atoms. The main advantage of a thermal beam is a continuous output and high signal levels. It also eliminates much of the complexity of a laser-cooled system allowing for smaller, more robust, and less expensive interferometers. Because the atoms in a thermal beam have velocities about 10^4 times larger than in laser-cooled experiments, Doppler shifts can be a problem. First-order Doppler shifts will be canceled with a unique precision prism alignment scheme insuring each pair of beams is exactly anti-parallel. The interferometer's completely symmetric beam path allows for a reversal of beam propagation to cancel residual Doppler shifts as well as other systematic errors. This will give us stability to match current primary standards.

1.3 An Ultra-stable Laser

When the interferometer is operating as a frequency standard, it is used to lock an extremely stable laser beam to the calcium ^{1}S$_0$ to ^{3}P$_1$ intercombination line. So the laser's long-term stability and frequency are determined by the atom interferometer, but feedback from the interferometer is slow. It takes a fraction of a second to read out the laser frequency relative to the atom transition. Because much faster feedback is needed to control the laser linewidth, the short-term stability and linewidth must be controlled independently of the interferometer before the laser light ever enters the

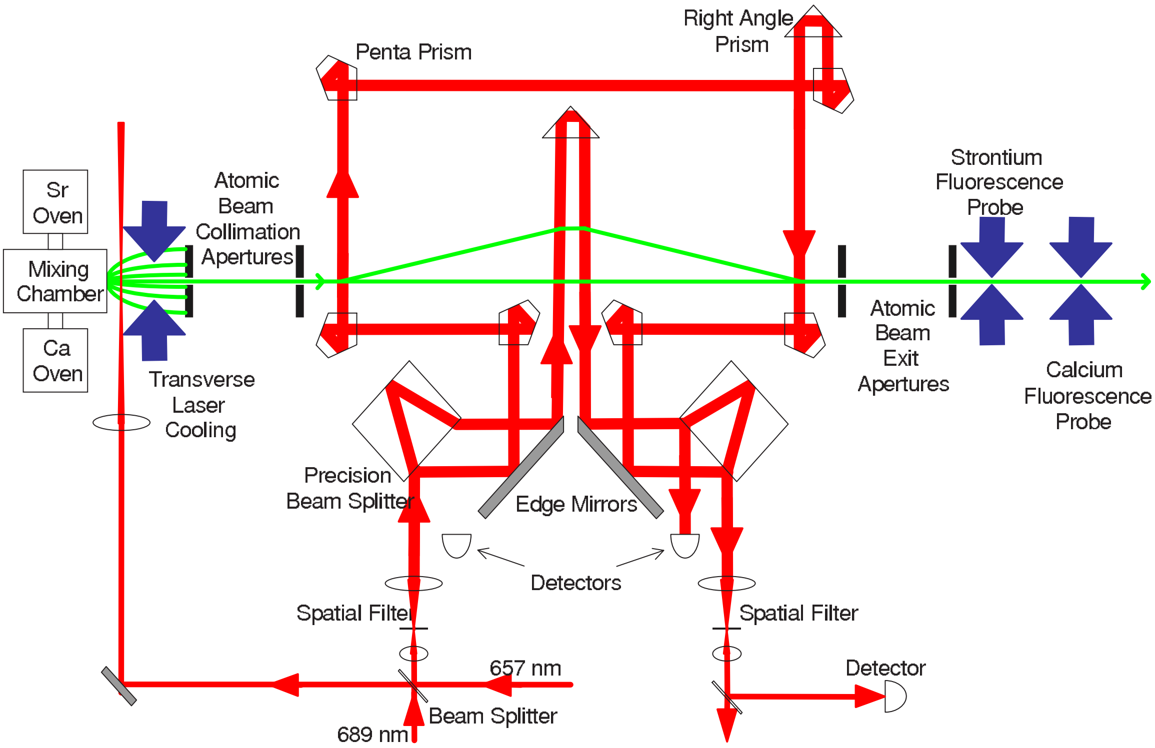

Figure 1.1 Interferometer layout

interferometer.

The laser must first be locked to a high finesse cavity to narrow its linewidth below 400 Hz (the linewidth of the calcium intercombination line). Since it is unlikely that the high finesse cavity will have a resonance at exactly the wavelength of the intercombination line, once locked to the high finesse cavity, the light from the laser must be shifted in frequency with an acoustic-optical modulator (AOM). This AOM will also be used to compensate for slow drifts in the cavity's frequency.

Although any linewidth below 400 Hz will provide a signal to lock to, the narrower the linewidth of the laser, the more accurate our frequency standard can be. Our goal (formed by what other groups have been able to achieve, and the requirements on linewidth to make our interferometer competitive) is a laser with a linewidth on the Hertz level. Assuming a symmetric lineshape, a Hz level linewidth laser should allow us to lock the central frequency of the laser to the calcium transition within 0.05 Hz. This is approximately one part in 10^{16} of the laser's frequency.

1.4 Linewidth and Stability

The quality of the laser is primarily determined by two properties: linewidth and stability. Linewidth is a measure of the width of the spread of frequencies emitted by a laser. But the linewidth depends on the time period over which it is observed. On the shortest of time periods (μs - ns) the linewidth is dominated by carrier density fluctuations, pump fluctuations, spontaneous emission, the quality factor (Q) of the cavity, and the coupling between the phase and amplitude of the lasing field [3]. On longer timescales (ms- μs), mechanical and thermal instabilities cause small drifts or jitter in the laser's frequency which are averaged into the linewidth, making it much wider. Over even longer periods of time (s - days), drifts in temperature and current can cause the laser frequency and amplitude to change. These changes are slow enough to no longer be included in the linewidth but looked at as stability.

Another problem that presents itself when using diode lasers is they often lase in more than one cavity mode simultaneously. Not only does multi-mode operation cause a large linewidth but it also makes the laser extremely unstable as power is continually transferred from one mode to another. This problem can often be fixed by adjusting the temperature and current. This changes the optical path length and electron-hole pair creation rate, hopefully making one cavity mode more favorable than the rest. The laser will then fall naturally into that single mode. But adjusting the temperature and current can be a tedious process and can not always be used to reach the desired frequency. More light must build up in the laser cavity either by using optical feedback or by increasing the Q of the cavity.

The optical cavity of a solitary diode laser is typically very short and the reflective facets of the diode that make up the optical cavity typically have relatively low reflectivity. This leaves us an optical cavity with low Q and the resulting linewidth is fairly large (tens of MHz) even if the laser is operating in a single cavity mode [3, 4]. The obvious way to decrease the linewidth of the laser and insure single mode operation is to increase the Q of the cavity. This is often done with a diffraction grating to direct some of the laser light back into the diode laser. This forms an Extended Cavity Diode Laser (ECDL). This not only increases the Q of the cavity, but also spectrally filters the optical feedback, introducing losses into undesired cavity modes. This method of optical feedback can narrow the fast linewidth of the laser by a factor of 100 and will be further discussed in chapter two.

A grating-stabilized diode laser will typically have a linewidth of several hundred kHz. To further narrow the linewidth and increase stability active feedback is most commonly used. The drift in frequency is measured relative to a source that is considered stable (in our case the resonance of a high-finesse cavity) and appropriate feedback is applied to correct for drifts and jitter. With the proper tricks it is possible to narrow the linewidth of a laser to the Hertz level and eliminate long term drifts entirely [5]. Active stabilization will be further discussed in chapters two and three.

1.5 Comparison with Other Lasers

The BYU atom interferometer group is not the first group to require a laser stable to the Hz level. There are several other groups working on frequency standards using the calcium $^1S_0$ to $^3P_1$ intercombination line. Although their interferometer designs are significantly different from ours, the same atomic transition is used. Regardless of the design, a narrower linewidth allows for greater precision. The frequency and time division at NIST in Boulder, Colorado has reported a 657 nm diode laser with a linewidth of 50 Hz [6]. A group at the Institut für Laserphysik in Hamburg Germany has claimed a linewidth as low as 30 Hz [5]. Both of these lasers use the Pound-Drever-Hall active stabilization method [7] and a lock circuit with a servo bandwidth of about 5 MHz. Passive stabilization of the high-finesse cavity plays an important role in both designs.

The design of our laser does not differ substantially from these. But while the basic design is similar, we have been able to apply several new tricks to improve upon their performance. We have integrated a scan balance circuit to allow for easier locking and scanning of the laser. We are also using high-speed surface mount electronics to increase our servo bandwidth. We are doing detailed analysis of the lock circuit system. We will use an optical cavity with higher finesse, will implement more passive stabilization, and will use a new ultra-stable current driver. With these improvements we hope to be able to create the most stable diode laser ever reported.

1.6 Outline of Thesis

In this thesis I focus on the development of the laser to drive the calcium clock transition in the atom interferometer. In Chapter 2 I discuss several diode laser grating stabilization systems including one developed and tested in our lab. The discussion includes advantages and disadvantages of our new design, as well as the techniques developed to test its theory of operation. Although this new grating stabilization scheme will not be used in the interferometer laser, it has many other potential uses, and many of the techniques learned through its testing are incorporated into the lock circuit that will control the Littrow laser used in the interferometer.

In Chapter 3 I discuss the passive and active stabilization that is used to further narrow the linewidth of the laser. The high-speed electronics used in the lock circuit are analyzed and discussed as well as a theoretical model of ideal feedback operation. A method for determining the laser linewidth from the error signal is discussed. We currently estimate a laser linewidth of 3 kHz. Since this thesis represents a work in progress, in Chapter 4 I offer a future outlook including important milestones that still need to be accomplished.

Chapter 2

Grating Stabilized Diode Lasers

2.1 Grating Stabilization

The first step in our ambitious goal of a Hertz linewidth laser is grating stabilization. To understand grating stabilization, some background on diode lasers is needed. Diode lasers are useful for many reasons: they are inexpensive, small, use little power, can be modulated at high frequencies, are tunable, and are available in a variety of wavelengths. But for them to be suitable for use in atomic physics they need a little help.

Just like any other laser, a diode laser has a gain medium (the active region of the semi-conductor) between two mirrors (the facets of the diode). Because diode lasers typically have very broad gain curves (compared to the small spacing between longitudinal modes) many cavity modes can lase simultaneously. This multi-mode operation leads to a large spread in wavelengths, fluctuations in power, and makes the laser practically useless for atomic physics. But with a small amount of optical feedback one mode can become stronger than the rest and through mode competition can virtually eliminate all other modes.

The most common way to make a diode laser run single-mode is with optical feedback. If light is fed into the laser diode at a frequency close to a cavity mode, the extra photons added to that mode will cause it to dominate stimulated emission, causing other modes to almost disappear entirely, leaving the laser to run single mode. Optical feedback also increases the Q of the cavity, narrowing the linewidth significantly.

2.2 Littrow Laser

One of the most commonly used ways to create an ECDL is the Littrow design [8]. This is perhaps the most simple of all schemes and is shown in fig 2.1. A diffraction grating is placed directly in front of the laser. The zeroth order reflection from the grating is used as the output beam and the first diffraction order is directed back into the laser. By changing the angle of the diffraction grating the frequency of light directed back into the laser can be adjusted. To tune the frequency of the laser without it hopping to another longitudinal mode, the angle of the diffraction grating and its distance from the cavity must be changed together, so that the number of wavelengths inside the cavity stays constant as the wavelength changes.

Perhaps the greatest advantage of this design is simplicity. There are only two parameters to adjust: the rotation and translation of the grating. It also has high output power, just enough light is coupled into the laser to cause it to lase single mode and the remaining light is all in the output. But it has one very serious disadvantage. As the grating is moved to direct different wavelengths of light back into the laser the angle of the zeroth order diffraction will also move. Even for very long scans this is a small change in angle, but it is large enough to cause problems in many applications.

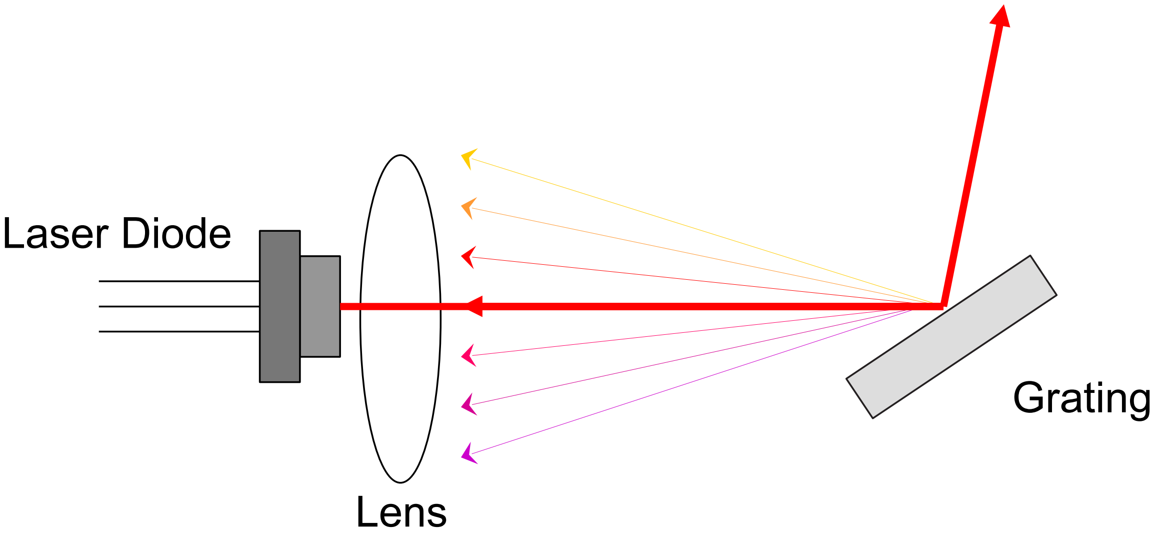

Figure 2.1 Littrow extended cavity diode laser stabilization scheme

2.3 Littman-Metcalf Laser

The Littman-Metcalf laser [9] solves the problem of the moving output beam. The output beam is still the zeroth order reflection from the diffraction grating, but instead of sending the first diffraction order directly back to the laser, it is sent to a mirror that directs the light back to the grating where it diffracts again and its first order is sent back into the laser. In addition to guiding the light back into the laser, the second diffraction further separates different frequencies of light giving increased side-mode suppression. The frequency of light directed back into the diode can be be adjusted by moving the mirror instead of the grating so the output beam stays stationary. The Littman-Metcalf design also has the advantage that if the geometry is properly chosen, a simple pivot point will keep the length of the cavity and the wavelength of light sent back into the diode perfectly matched, providing an “infinite” mode-hop-free scan range. Ideally this could allow one to scan mode-hop-free over the entire gain curve of the diode by simply pivoting the mirror about a fixed point. In practice the mode-hop-free scan range is limited by the extra boundary conditions introduced

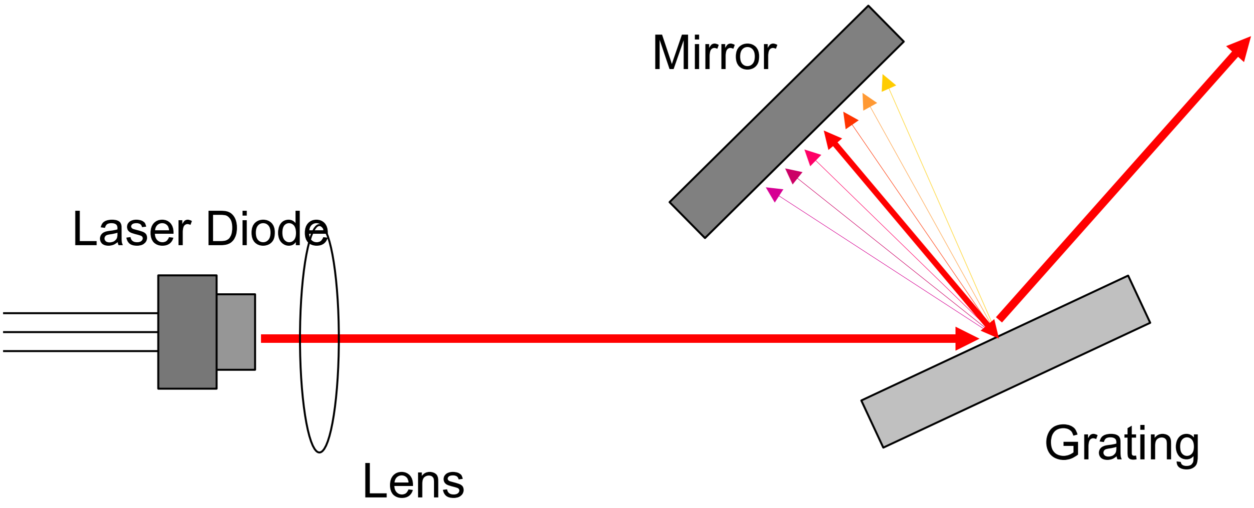

Figure 2.2 Littman-Metcalf extended cavity diode laser stabilization scheme

by the reflectance of the front facet of the diode.

The Littman-Metcalf design has one very large disadvantage. The light from the first diffraction order diffracts off the grating again as it is sent back into the laser. The first order beam is directed back into the laser, but the zeroth order diffraction is lost. This causes a loss of power that can be 20% or more. With visible diode lasers typically putting out just a few mW, this can be an unbearable loss.

2.4 Merrill-Durfee Laser

The Littrow and Littman-Metcalf designs each have their advantages and disadvantages. Littrow provides simplicity and power, but at the expense of a small angular displacement. The Littman-Metcalf design gives a stable pointing beam and a theoretically infinite mode-hop-free scan range with a simple pivot point, but throws away a large percentage of the power. It would be ideal if the two designs could somehow be combined to get both power and stability. It was noticed that in another experiment we injection lock a laser by sending light backwards through the rejection port

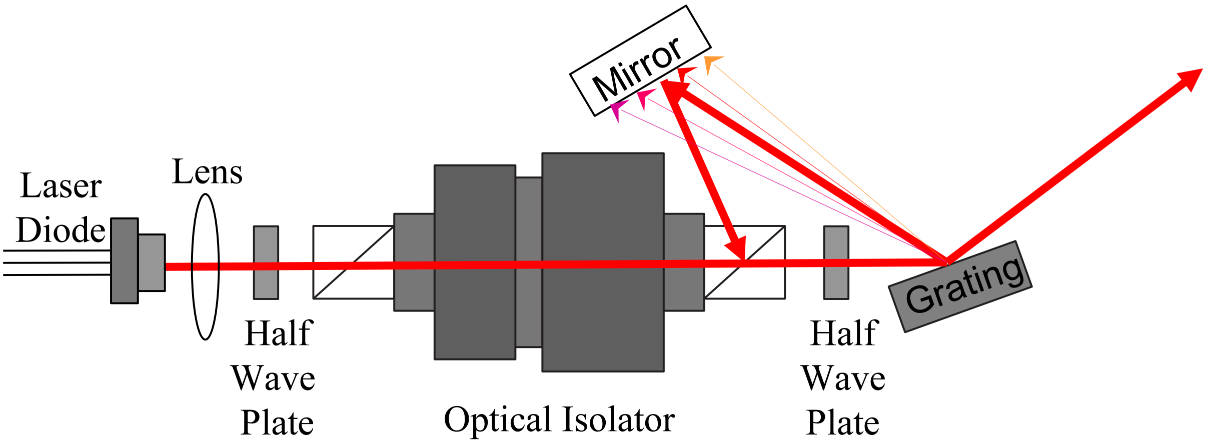

Figure 2.3 The Merrill-Durfee extended cavity diode laser stabilization scheme

of an optical isolator [10]. By similarly using an optical isolator inside the extended cavity of the diode laser, the light from the diffraction grating could be coupled into the laser through the rejection port, eliminating the second diffraction that kills the power output in the Littman-Metcalf design. But like the Littman-Metcalf design the frequency of the light sent back to the laser is tuned with a mirror allowing the angle of the output beam to remain stable during tuning.

The preliminary work of designing this new scheme was done by Rebecca Merrill and Dallin Durfee before my joining the group. They also developed a mathematical model to explain its operation and scan range. One of my first tasks upon joining the group was to validate this model with experimental results.

The mathematical model is so important because at first glance it appears that this design should not work. To tune the laser continuously without mode-hops, the length of the cavity must always be equal to an integer number of half wavelengths. As the frequency of the light is changed, the cavity length must also be changed to accommodate the new wavelength of light. But in this new design, as the mirror is moved to direct shorter wavelengths of light back into the laser, the path length

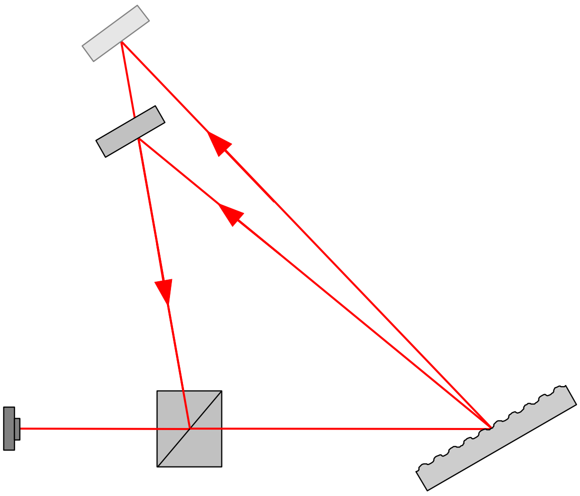

Figure 2.4 When the mirror is moved to direct shorter wavelengths into the diode laser the path length gets longer. This will cause the laser to mode-hop to another longitudinal mode. (figure reproduced from [11])

actually gets longer as shown in Fig 2.4. If it is assumed that the ray that is directed straight back into the laser always determines the mode that the laser will select, the laser should not be able to scan more than a few MHz before mode hopping.

But the laser is able to scan several GHz mode hop free with even a rough alignment. It seems reasonable that if a beam of light were to enter the laser with an angle less than the gaussian divergence of the beam, the light would couple into the laser almost as if it was coupled directly as illustrated in Fig. 2.5. With this reasonable assumption a model for its scanning behavior was developed. It predicted a pivot point around which the mirror can be rotated to obtain a large mode-hop-free scan range. The development of the Merrill-Durfee Laser and a potential mathematical model describing are explained in detail in Refs. [11, 12].

The testing of this model should have been relatively straight-forward. Scan the laser with different pivot points, and see if it scans as far as the model says it should. Unfortunately the common problem of front facet reflections had to be addressed

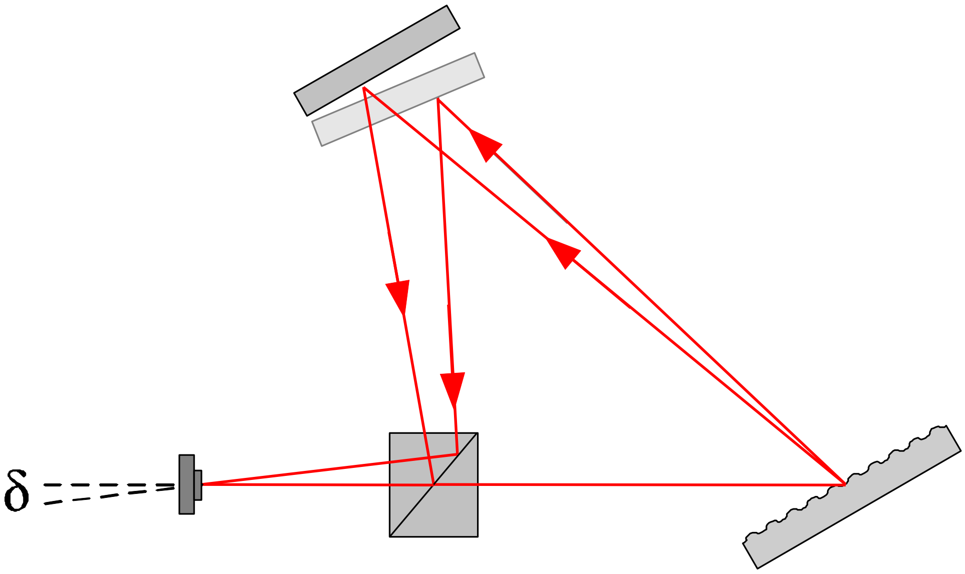

Figure 2.5 If the light is allowed to enter the diode at an angle less than the Gaussian divergence of the beam the mirror can be moved such that the cavity length decreases as the wavelength gets smaller (figure reproduced from [11])

before such measurements could be made.

2.5 Front-Facet Reflections and Their Solution

When running a diode laser in an extended cavity configuration reflections off the front facet of the diode introduce another boundary condition. Not only must there be an integer number of half-wavelengths in the cavity formed by the back facet of the diode and the diffraction grating, but there must also be an integer number of half wavelengths between the diffraction grating and the front facet of the diode and between the two facets of the diode. This additional boundary condition can be useful in suppressing unwanted modes, but it also causes a decrease in mode-hop-free scanning range.

To avoid these unwanted mode hops, an anti-reflective coating can be applied to the front facet of diode lasers used in external cavities. We bought a diode laser coated by Sacher Laser that quoted us an remarkably low reflectivity of one part in 10^8. These anti-reflective coatings are expensive and (as I found out when testing the Merrill-Durfee Laser) even with the best coatings do not entirely get rid of the unwanted boundary condition.

When testing the laser for its optimum scan range the laser could only scan a few GHz before hopping to another cavity mode. Using a wavemeter I was able to determine that the modes were consistently hopping to cavity modes 45 GHz away. This corresponds to the free spectral range of a mm size cavity. Several of the optical components in our setup had sizes comparable to the observed mode-hop. By systematically replacing waveplates, collimating lenses, and other optics in the system I determined that this mode-hop was caused by the reflection of the front facet of the laser diode. Even with the best anti-reflection coatings on the market the reflectance of the front facet was too high.

Since I could not get rid of the unwanted boundary condition from the reflection off the diode's front facet, I had to effectively change its position as I scanned the rest of my system. This boundary condition can be scanned by changing the current powering the diode as the frequency is scanned. This is known as current feed forward. Injection current changes the optical path length and the frequency that the laser will operate at through several different mechanisms. The dominant effect for relatively slow changes in current is a change in temperature which causes the diode to expand or contract, changing the optical length of the diode. Injection current also changes the electron-hole pair density which affects the index of refraction of the gain medium which changes the optical path length. A change in current also has slight effects on the gain curve of the diode, although this is usually insignificant compared to the other effects. Although current feed forward has several effects, the correct amount of current feed forward to "scan" the optical path length of the diode as it is tuned with a grating can be determined empirically.

The mathematical model that was tested predicts the mode-hop-free scan range of the laser with a particular pivot point. By using current feed forward I can assure that the mode-hop-free scan range is limited by the way the mirror moves around the pivot point and not by unwanted reflections off the front facet of the diode. I built a simple electronic circuit that applies a set ratio of voltages to the piezo-electric transducers (PZT) controlling the position of the mirror and the laser current driver. With the mirror mounted on a standard PZT mirror mount a pivot point can be simulated by adjusting the ratio between the PZT transducers on the two sides of the mount.

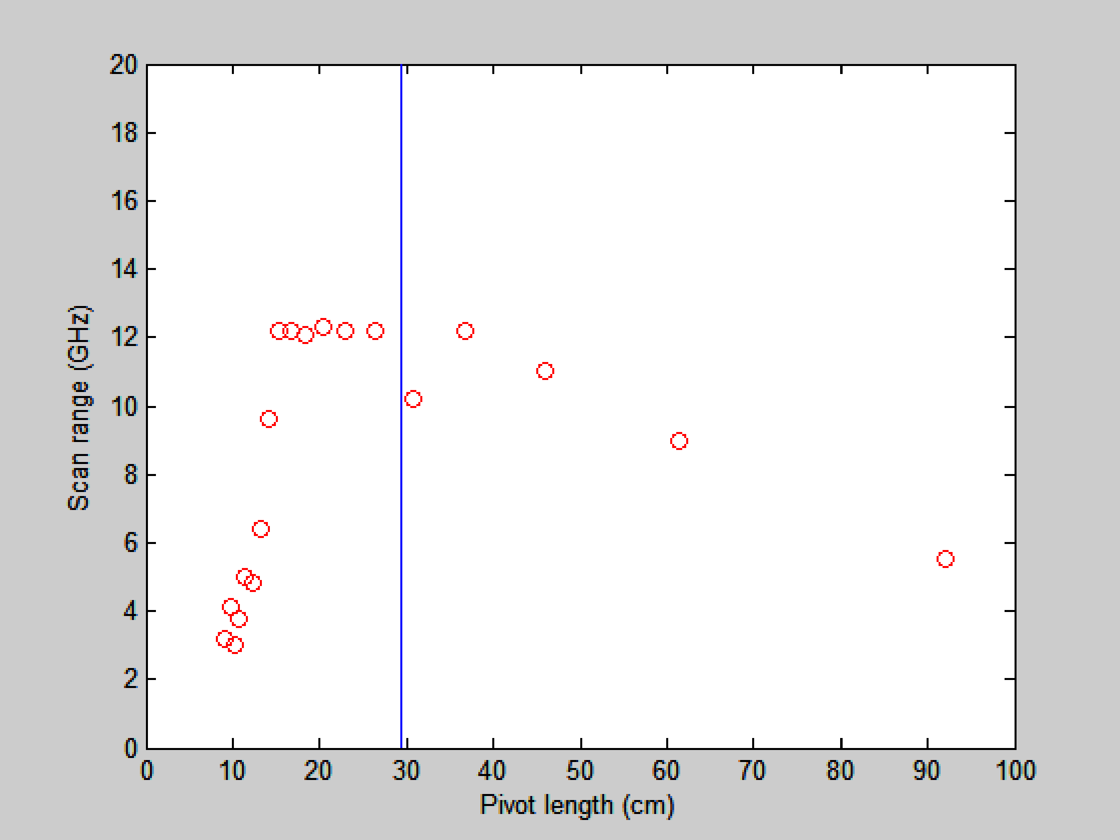

Figure 2.6 The results of the first test of optimum pivot length plotted together with a line representing the optimum pivot length predicted by the mathematical model

2.6 Testing for Optimum Scan Range

Originally I set the ratio on the PZT to give the optimum pivot point as predicted by the mathematical model. I adjusted the current feed forward to allow for the largest mode-hop-free scan range. Then to test the mode-hop-free scan range at other pivot points I changed the effective pivot point without changing the rate of current feed forward and measured the range (in frequency) the laser would scan without any mode hops. The results are illustrated in Fig. 2.6.

The proposed mathematical model predicted an optimum pivot point that would give a large mode-hop-free scan range (over 100 GHz). It also predicted that as the pivot point was varied from this optimum distance the possible mode-hop-free scan range would decrease drastically. A first glance at the results of my experiment seems to validate the model. The mode-hop-free scanning range goes up as the predicted optimum pivot point is approached. It hits a plateau. Then on the the other side of the predicted pivot point it goes back down again. The plateau occurs when the laser would scan the entire range of the PZT controlling the mirror without mode hopping. This seems to indicate that all that was needed was a larger PZT stack so the mirror could move farther, scans the full predicted mode-hop-free scan range and show that the model accurately predicts the optimum pivot point. I built a long mechanical pivot arm to increase the range of motion. As I tested the scan range I found that my results varied greatly with the amount of current feed forward. In fact the pivot point that gave the best scan could be changed with with the current feed forward.

I originally misinterpreted the data in Fig. 2.6, thinking that it supported the proposed mathematical model. As the laser is scanned three factors decide which mode it will lase in: the longitudinal modes of the extended cavity, the frequency of light directed back into the laser by the diffraction grating, and the longitudinal modes of the diode. The effect of these three factors is qualitativly illustrated in Fig. 2.7. The effect of the grating is very small compared to the cavity modes of the diode and the external cavity. The optimum pivot point was calculated by matching the change in wavelength from the grating and the length of the external cavity. But these two lengths don't have to be matched precisely if there is another strong boundary condition that keeps the laser from hopping to another mode. Since the modes of the diode cavity are spaced approximately 45 GHz apart, the mode of the laser will be determined almost entirely by the length of the two cavities as long as the peak of the diffraction grating and external cavity mode are mismatched by less than 45

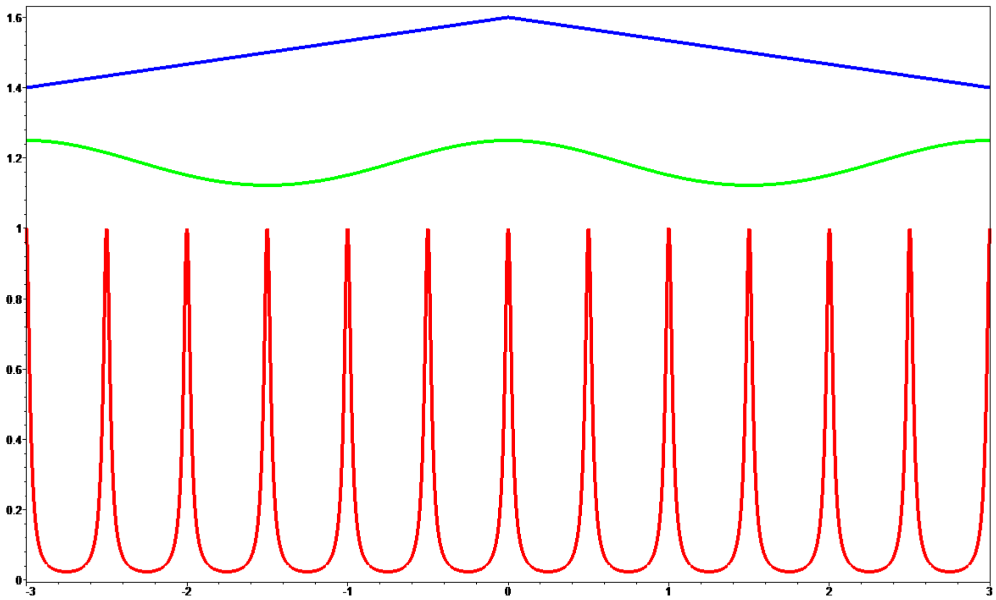

Figure 2.7 A qualitative illustration of mode selection - The laser will lase in whatever mode has the least losses. The red sharply peaking curve represents the cavity modes of the extended cavity, the blue triangle represents the light directed back into the laser by the diffraction grating, and the green curve represents the cavity modes of the diode. The laser will jump to wherever the sum of these three curves is the greatest.

GHZ. The interpretation of Fig 2.6 is now clear. If only the current feed forward and length of the external cavity are scanning together, many GHz can be scanned before the diffraction grating makes it preferable to jump to the next cavity mode of the diode 45 GHz away. After my original optimization I changed the pivot point without changing the current feed forward, this caused the external cavity to change at a different rate, but the diode cavity changed just as it did before and the laser mode hopped sooner, but to another external cavity mode 2 GHz away.

For this reason, testing just the pivot point is more complex than it might seem. With each pivot point tested, the current feed forward must be adjusted so that no small mode hops occur. This ensures that the current feed forward is lined up exactly with the external cavity length. The laser must then be scanned until a mode hop of 45 GHz is seen, indicating that the length of the external cavity and the wavelength of light from the diffraction grating are changing at different rates, and that the pivot point is incorrect. If the laser mode-hopped forward then the pivot point is too short; if the hop was backwards the pivot point is too long. With this method I was able to show that the pivot point predicted by the mathematical model was incorrect. With each pivot point the mode-hops indicated that optimum pivot point was larger. I eventually ran out of range on my lever arm so I was unable to find the optimum pivot point. But when using pure translation (an effective pivot point of infinity) the mode-hops pointed to a smaller pivot point. This shows that the optimum pivot point for the Merrill-Durfee Laser is larger that 30 cm but less than infinity.

These results definitely invalidate the mathematical model proposed in Ref. [11, 12]. It is possible that the light couples back into the laser with a slight offset rather than at an angle, or perhaps a combination of both. Whatever the reasoning, the scheme appears to be too complex for an exact prediction of its optimal operational parameters. But I have demonstrated mode-hop-free scans as large as 45 GHz using computer-controlled nonlinear current feed forward and shown that with a simple voltage divider used to linearly scan the current and PZT voltages together, the laser can easily be tuned to scan approximately 8 GHz mode hop free.

2.7 Application of the Merrill-Durfee Laser

Because of the Merrill-Durfee laser's high output power and ability to make long mode-hop-free scans without changing the angle of the output beam, it is a good choice for many applications. One such application is our recent measurement of the calcium intercombination line in a high temperature vapor cell [13].

A special high-temperature vapor cell was constructed to provide a high density vapor of calcium atoms so absorption on the $^1S_0$ to $^3P_1$ line can be measured. The light was sent through the high-temperature vapor cell and onto a photo diode. The laser's frequency was then scanned over approximately 8 GHz mode-hop-free. As the frequency of the laser was scanned across the atomic line the light was absorbed and spontaneously emitted in a random direction causing a decrease in the intensity of light hitting the photo diode. Five sweeps of the laser were averaged to increase our signal to noise. The light from the laser was also sent to a high finesse optical cavity and the position of the nearest cavity mode was measured relative to the atomic line.

Although the laser was able to scan many GHz mode hop free, it was difficult to prevent the laser from drifting drifting or mode-hopping between data sets. Frequent adjustment of laser current and PZT voltage was needed to keep the laser single mode over any period of time greater than a few minutes.

2.8 Advantages and Disadvantages of the Merrill-Durfee Laser

The Merrill-Durfee laser has many advantages. It is easily scanned many GHz mode-hop-free without computer control. It has the ability to scan the full free spectral range of the diode if the proper ratio for current feed forward is applied. As the laser scans the angle of the output beam stays fixed. And it has high output power.

Because the optical isolator must be placed inside the external cavity, the length of this cavity is much larger than most extended cavity diode lasers. Increasing the length of the cavity decreases the free spectral range, decreasing the spacing between cavity modes. This gives the laser a narrower linewidth, but it also makes the laser more likely to jump into another longitudinal mode when perturbed. So small drifts or vibrations are effectively magnified by the length of the cavity, making stability an issue in noisy environments. The passive stabilization that will be used in the interferometer solves the stability issue, but it is feared that a small change in temperature could change the optical path length inside the optical isolator even when the cavity itself remains stable. This change in optical path length could cause the laser to mode hop.

Many of the reasons for developing this new laser stabilization scheme were unrelated to the end goal of an atom interferometer. Higher output power and the ability to scan the laser many GHz mode-hop-free without moving the output beam are desirable for many applications. But while the atom interferometer is in operation the laser frequency is locked to the clock transition and is scanned only briefly to obtain the initial lock. So the moving output beam of the Littrow scheme is not a big problem. So to minimize complexity we decided to use the Littrow laser in the interferometer.

Chapter 3

Beyond Grating Stabilization

3.1 Passive Stabilization

Over very short measurement times diode lasers in an extended-cavity configuration have narrow linewidths. The problem is that mechanical vibrations, temperature fluctuations, fluctuations in power and other factors cause the narrow linewidth of the laser to constantly jitter and jump around so on any useful timescale the laser's linewidth appears much broader. This jitter can be actively canceled out by measuring the frequency of the laser and providing the appropriate feedback. But the less noise that is on the laser, the less active noise cancelation will be needed. For this reason great improvements in linewidth can be made just by mechanically and thermally stabilizing the laser.

The first step in mechanical stabilization is choosing the hardware carefully. I will not attempt to go into all the details of what makes a stable mount, but I will emphasize the importance of selecting each mount carefully and tightening each part carefully. In order to isolate our system from mechanical vibrations we have mounted all the optics onto a quarter-inch thick optical breadboard. We have isolated this breadboard from an inch-thick aluminium plate by four cubes of Sorbethane. The size of these Sorbethane cubes is carefully calculated to assure maximum vibration isolation. The larger aluminum plate is set upon four more cubes of Sorbethane which rests on a floating optics table. The purpose of this setup is to isolate the laser from any mechanical vibrations in the room. Other groups have employed similar techniques and have seen a dramatic suppression of mechanical noise above 100 Hz [14].

3.2 Active Stabilization

Even the best passive stabilization scheme leaves a residual amount of noise that has to be canceled actively. To do this the frequency of the laser must be measured relative to some stable reference. We choose to compare our laser's frequency with a transmission line of a high finesse cavity. Once aligned properly the amount of light reflected off of or transmitted through the cavity indicates how far the laser's frequency is from the cavity resonance. With a little electronic trickery this information provides us with an error signal with a voltage proportional to the laser's frequency drift. A high-speed electronic circuit uses this error signal to change the current of the laser diode and the PZT voltage to move the laser's frequency back on resonance with the high finesse cavity. If this is done fast enough the high frequency jitter that contributes to the laser's linewidth can be reduced significantly and the linewidth narrowed.

The best high-finesse cavities commercially available have a finesse of 300,000 corresponding to a linewidth of one kHz. This narrow linewidth is needed to get good frequency resolution so even small fluctuations of the laser's frequency can be detected and corrected. Because the laser is locked to the high-finesse optical cavity, the laser can only be as stable as the cavity. For this reason the two mirrors making up the cavity are held together with an ultra-low-expansion quartz tube preventing thermal expansion from changing the length (and resonant frequency) of the laser. This ultra-low-expansion quartz cavity rests upon sapphire balls (chosen because of their low thermal expansion rate and excellent vacuum qualities) which sit on an invar block. This assembly sits inside of a vacuum chamber evacuated to a pressure of 3×10^{-7} torr. This low pressure ensures that the index of refraction of the air inside the cavity will remain constant. The windows of the vacuum chamber are tilted 15° from vertical to ensure that any reflections from the windows miss the cavity entirely, thereby eliminating unwanted heating and the possibility of a standing wave forming unwanted boundary conditions. The vacuum chamber is placed inside a large box made of inch thick aluminium lined with lead foam sitting on four squares of Sorbehtane. This provides excellent mechanical and acoustic vibration isolation. The box will also be actively temperature stabilized.

3.3 Pound-Drever-Hall Method

The simplest way to compare a laser's frequency to the resonance of an optical cavity would be to detect the light transmitted through the cavity. The closer the laser's frequency is to the resonance, the more light would be transmitted. But this presents two problems. When the laser is right on the resonance the transmission is at a maximum, but when the laser's frequency drifts and the transmission drops there is no way to tell if the laser's frequency is too high or too low, making it difficult to provide the appropriate correction. High finesse cavities also have a long ring-down time. With a cavity of 300,000 finesse it takes light 150 μs to exit the cavity, limiting the speed at which error can be detected and corrected to 6 kHz.

Pound-Drever and Hall found a way to generate an error signal from the light reflected from the cavity [7]. By generating the error signal from the light reflected off the front of the cavity the problem of this long ring-down time is eliminated and the limitation on feedback speed is lifted.

To detect the frequency drift and create an error signal the light coming from the laser must be modulated so it has frequency sidebands. This can be done by modulating the laser's current (but this adds unwanted amplitude modulation) or by using an Electric Optic Modulator (EOM) to phase modulate the light. In our lock we are using an EOM which adds 30 MHz sidebands onto the laser. These sidebands each produce a beat note with the laser's central frequency peak, but with opposite phases. This way the beat note produced between the central peak and the upper sideband exactly cancels out with the beat note produced by the central peak and the lower sideband.

After sidebands have been added to the laser light, it is sent to the cavity. The reflected light is directed to a photo diode. The incident light interferes with the light leaking out of the cavity and provides an error signal encoded on a beat note at the modulation frequency which is chosen to be well above the noise that needs to be canceled. The amplitude and the phase of this beat note tell us how far the laser's frequency has drifted and in which direction.

The signal is sent through a high pass filter to get rid of the DC component of the photodiode signal and then through two radio frequency amplifiers. The signal is then mixed with a signal from the oscillator driving the EOM. The mixer demodulates the error signal, providing a DC voltage indicating how far the laser's frequency has drifted from the cavity resonance. By adjusting the length of the cable running between the oscillator and the mixer, the phase between the signal from the cavity and the oscillator can be adjusted to maximize the error signal magnitude and adjust it to have the proper slope. With our particular setup a positive voltage indicates

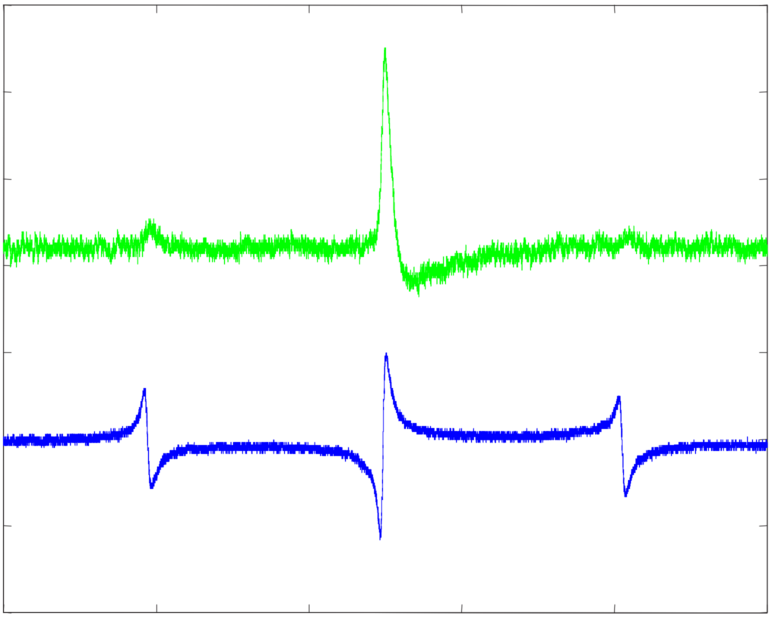

Figure 3.1 The error signal as the laser frequency is scanned over the cavity resonance.

that the laser frequency is too low and a decrease in laser current is needed to push the frequency back up to the desired value. The error signal obtained by scanning the laser's frequency through a cavity resonance is shown in Fig 3.1.

3.4 High Speed Lock Circuit

The Pound-Drever-Hall method provides a method for the deviation of the laser's frequency to be measured quickly. A lock circuit must be constructed to process the information from the error signal and provide appropriate feedback. It was mentioned earlier that there are several ways to control the wavelength of the laser. By controlling the voltage on the PZT the diffraction grating forming the extended cavity of the diode laser can be moved, changing the wavelength of the laser. Changing the current of the laser also causes a shift in the laser's wavelength.

The diffraction grating is important for long term stability. Over time the laser's wavelength may drift. By changing the position of the grating these drifts can be canceled. Unfortunately the mechanical resonances in the PZT and its mount prevent feedback at frequencies above a few hundred Hz.

Faster feedback and the majority of our locking ability comes from changing the current. Due to the many different mechanisms inside the laser that control the laser's wavelength, on different time scales the laser will respond to errors (as well as corrections for errors) in a different ways. For example at frequencies near DC a small increase in current causes the laser diode to heat up and expand, increasing the wavelength. But at high frequencies (several MHz) an increase in current causes changes in the temperature and the electron-hole pair density fluctuations dominate, decreasing the index of refraction which decreases the optical path length and decreases the wavelength.

The cavity and all the electronics used to detect and generate the error signal also effect how errors of different frequency and magnitude are seen by the lock circuit. For example the cavity's sensitivity at frequencies above the cavity's linewidth is thought to fall off at -20dB/decade [15]. The response of the combined laser and detection scheme to different frequencies could be expressed as a transfer function (ie at each frequency x amount of error corresponds to z volts of error signal).

An ideal controller is the "inverse" of the transfer function. Whenever an error was detected the controller would inject just the right amount of current at just the right frequency to exactly cancel out that error. Unfortunately we don't know the laser's transfer function. But a simple circuit that provides proportional, integral, and differential (PID) gain can do a reasonable job at canceling out errors in the laser.

Proportional gain provides a voltage (turned into a current in the laser driver) that is proportional to the error signal. The further away from the lock point the laser drifts, the harder the proportional gain pushes it back. Proportional gain gives even gain over all frequencies, but because the correction is proportional to the error,it can only demagnify the error.

To help the laser converge on the lock point integral gain is added. Integral gain provides feedback proportional to the integral of the error signal. The integral "sees" that the wavelength is consistently above (or below) the lock point and provides feedback to make it converge quickly.

Differential gain essentially predicts which way the laser's wavelength is changing to cancel out errors before they become large. It also helps expand the bandwidth of the circuit. It takes time to measure the linewidth, create the error signal, and provide feedback. This time delay sets a limit on the bandwidth of the lock. At frequencies where this time delay corresponds to a phase lag of 180° feedback that intended to cancel out an error will magnify it and the circuit will start to oscillate. The differential gain provides favorable phase shift, effectively predicting the future, expanding the bandwidth of the circuit. With our current PID controller our lock bandwidth is about 4 MHz (Fig 3.2) which we hope to further expand.

3.5 Linewidth Estimation

A standard way to measure a laser's linewidth is to measure its transmission through a high-finesse cavity with a known linewidth as the laser or cavity resonance is scanned in frequency. But the best optical cavities in the world only have linewidths around 1 kHz.

It is also possible that the laser could be scanned across a narrow transition; there is a very narrow transition in Co III less than a nm away that we could use to measure. But obtaining a dense vapor of Co III ions and doing Doppler-free spectroscopy seems like much more trouble than it is worth.

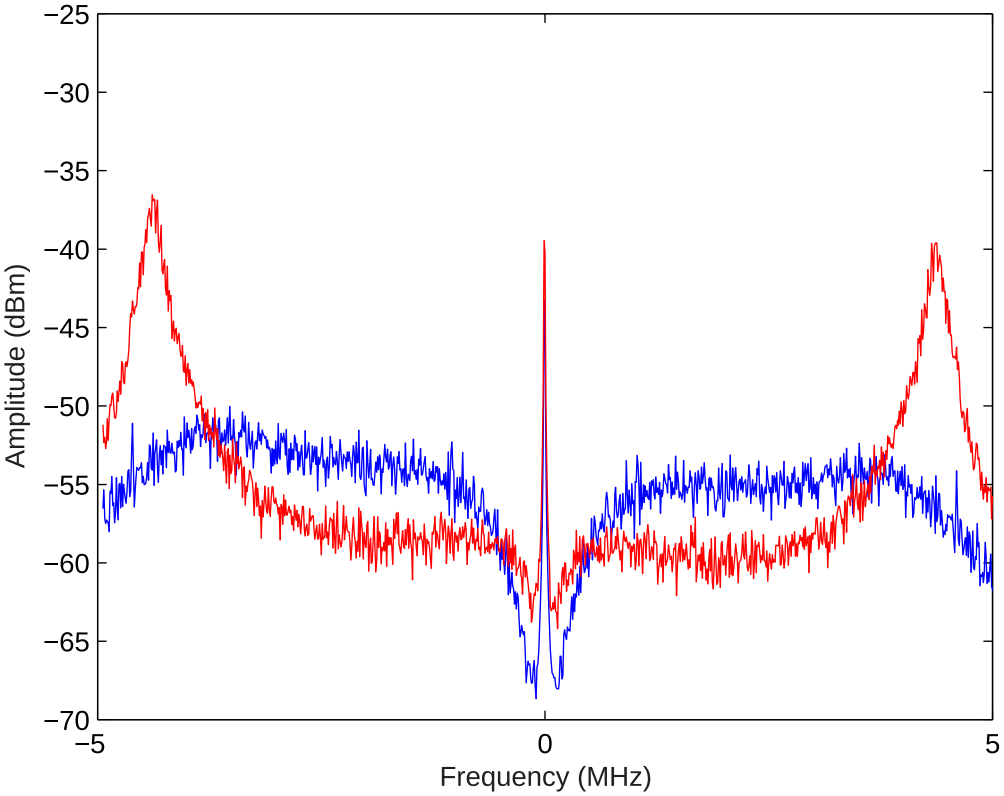

Figure 3.2 A fast fourier transform of the beat note recorded by the photodiode reveals the noise spectrum of the light. The blue trace is the noise spectrum with the laser locked. The lock circuit bandwidth is measured by tuning up the gain until oscillation peaks are first seen as shown in the red trace. This shows a bandwidth of about 4 MHz. The frequency shown measured relative to the modulation frequency of 30 MHz

Another way to measure the linewidth is to beat the laser together with an independent laser of comparable linewidth. This will be insensitive to any noise that is common to the two lasers, but if they are completely independent the method works. This may be the only way to accurately measure our linewidth. But until we are convinced that our laser really has a linewidth near the Hz level we don't want to invest the time and money constructing another laser that we are not sure will be good enough for our interferometer.

Because of the difficulties involved with an accurate measurement of linewidth, it is important to develop a good way to estimate the linewidth. We should be able to extract some linewidth information from the error signal generated by the high finesse cavity to which the laser is locked. This has some inherent problems. Our lock circuit is incapable of canceling out any error that the cavity is insensitive to. Any error the cavity can't see can't be canceled out by the lock circuit and will remain on the light. But since the cavity is insensitive to this noise, we cannot include this in our estimate of linewidth.

Even though estimating the linewidth from the error signal has many problems it is currently the best estimate tool we have. The laser is first locked very loosely using only proportional and integral gain. Then the proportional gain is turned up causing the laser to oscillate. As the proportional gain is turned up the magnitude oscillations will grow until they have reached the full peak-to-peak voltage of our error signal at which point the sine wave will become distorted as we "roll over" the peak of the error signal. This gives us a rough estimation of the voltage corresponding to the full linewidth of the cavity. Then the lock is optimized, adjusting the proportional, differential, and integral gain iteratively until the error signal is minimized. This peak-to-peak voltage can be measured and compared to the maximum voltage to estimate the laser's linewidth. As seen in 3.3 the error signal of the laser locked to a

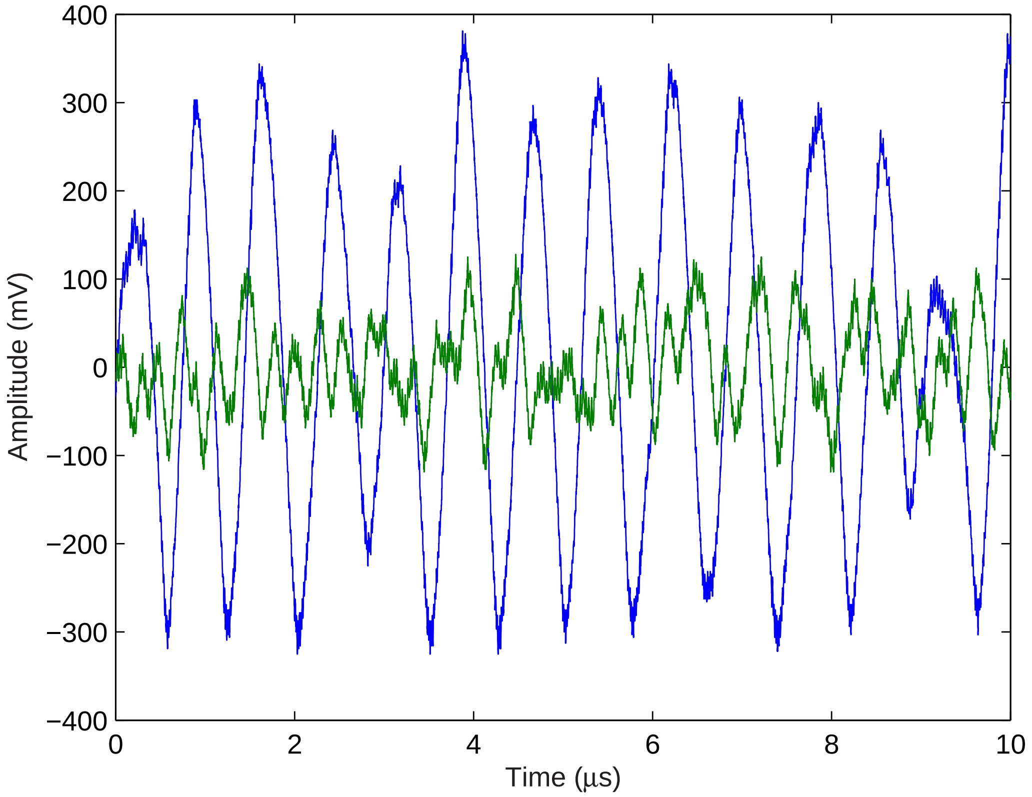

Figure 3.3 The gree trace is the error signal produced by the Pound-Drever-Hall method when the laser is locked. The gain is then increased until the laser oscillates the producing the blue trace. This comparison leads us to estimate our linewidth at 3 kHz

10 kHz linewidth cavity when the lock circuit is peaked up is less than one-third of the full-scale oscillations indicating that our laser's linewidth is about 3 kHz.

Unfortunately when the response of the mixer was measured it was found to be non-linear. As the amplitude of the beat note is increased there is a point at which the DC output "rolls over" and provides a smaller DC signal for a larger beat note. This probably means that the laser linewidth is actually less than estimated with this method. A precise measurement will be needed to know for sure. We are currently locked to a cavity with a linewidth of 10 kHz. With an AOM we can scan the frequency of the locked laser across the transition of the a higher-finesse cavity (which we plan to lock the laser to in the near future) to make a precise measurement of the laser linewidth. By measuring the linewidth with the laser at several stages of locking we can calibrate our estimation method so that when the laser is locked to the highest finesse cavity available we can make a good estimation of linewidth without having to build another laser. We are still waiting for the AOM driver to arrive so we can preform this test.

Chapter 4

Current Status and Future Outlook

I have locked the laser down to a linewidth of 3 kHz. This is far from the desired linewidth of 1 Hz, but it already makes it one of the more stable lasers in the world. I am still locking the laser to a cavity with a finesse of 30,000 with no isolation. This has been useful while testing various components of the lock circuit. Moving to the ultra-stable cavity with a finesse of 300,000 will not only improve our lock by a factor of ten due to the increase of finesse, but the stability of the cavity should also make large improvements.

We also have more passive stabilization planned. Temperature stabilization is very important. For this reason the entire laser system will eventually be placed in an inch-thick aluminum box. Cooling elements will be attached to the outside of the box and the temperature will be controlled actively. The thickness of the box gives it a large mass which not only reduces high frequency vibration, but also assures a uniform temperature inside the box. To prevent the heat generated by the laser diode from heating up the air inside the box, causing convection currents, the laser will be mounted directly to the aluminum box which will act as a heat sink.

We are also concerned with the stability of our laser current supply. Any noise on the current driver or on the power supply powering any part of our system can find its way onto the laser light. For this reason Greg Doermann, another student in the lab, is currently working on new laser current drivers that should provide greater stability with less noise. We are also considering replacing our lab standard switching power supply with a linear power supply. We are not yet sure how much of an improvement in stability this will yield.

If we can design a new tank circuit for our EOM the modulation frequency can be raised allowing us more usable bandwidth and greater insensitivity to noise. It would also push the residual modulation noise and its harmonics well outside of the range of our lock circuit making sure that residual modulation noise is not written onto our laser by the lock circuit.

We are also working with a controls group from the Department of Computer Science at BYU to design an optimal control loop. We are currently using PID because it is easy to understand and does a satisfactory job, but there is no reason to believe that it is optimal. With this help we hope to change our lock circuit substantially to match the transfer function corresponding to our laser diode.

Once we have learned a little more about the ideal lock circuit from our in-depth control analysis it will be time to "take the training wheels off". We will then lock to the stabilized 300,000 finesse cavity and will hopefully be able to achieve the goal of a 1 Hz linewidth laser.

Bibliography

[1] T. Hänsch, *Talk given in the plenary session OSA* (2005).

[2] C. J. Bordé, "Atomic interferometry with internal state labeling," Phys. Lett. A **140,** 10 (1989).

[3] R. W. Fox, L. Hollberg, and A. A. Zibrov, in *Atomic, Molecular, and Optical Physics: Electromagnetic Radiation*, F. B. Dunning and R. G. Hulet, eds., (Academic Press, 1997), pp. 77–102.

[4] C. E. Wieman and L. Hollberg, "Using diode lasers for atomic physics," Rev. Sci. Instrum. **62,** 1–20 (1991).

[5] A. Schoof, J. Grunert, and S. Ritter, "Reducing the linewidth of a diode laser below 30 Hz by stabilization to a reference cavity with a finesse above 10^5," Opt. Lett. **26,** 1562–1564 (2001).

[6] C. W. Oates, F. Bondu, and R. W. Fox, "A diode-laser optical frequency standard based on laser-cooled Ca atoms: Sub-kilohertz spectroscopy by optical shelving detection," Eur. Phys. J. D **7,** 449 (1999).

[7] 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,** 97–105 (1983).

[8] J. C. Camparo, "The diode laser in atomic physics," Contemp. Phys. **26,** 443–477 (1985).

[9] M. Littman and H. Metcalf, "Spectrally narrow pulsed dye laser without beam expander," Appl. Opt. **17,** 2224–2227 (1978).

[10] P. Bouyer, T. L. Gustavson, K. G. Haritos, and M. A. Kasevich, "Microwave signal generation with optical injection locking," Opt. Lett. **21,** 1502 (1996).

[11] R. Merrill, R. Olsen, and D. S. Durfee, "Increasing the Output of a Littman-Typer Laser by Use of an Intracavity Faraday Rotator," Appl. Opt. **43,** 3910–3914 (2004).

[12] R. Merrill, "A New Scheme for External-Cavity Diode Laser Stabilazation," (2003).

[13] C. J. Erickson, B. Neyenhuis, and D. S. Durfee, "High-temperature calcium vapor cell for spectroscopy on the 4s2 1s0 to 4s4p 3p1 intercombination line," Rev. Sci. Instrum. **76,** 123110 (2005).

[14] C. W. Oates, E. A. Curtis, and L. Hollberg, "Improved short-term stability of optical frequency standards: approaching 1 Hz in 1 s with the Ca standard at 657nm," Opt. Lett. **25,** 1603–1605 (2000).

[15] R. W. Fox, C. W. Oates, and L. W. Hollberg, "Stabilizing diode lasers to high-finesse cavities," Experimental Methods in Physical Sciences **40,** 284–90 (1998).