Increasing Stability in Extended Cavity Diode Lasers using Frequency Noise Feedback

McKinley Pugh

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

Bachelor of Science

Dallin Durfee, Advisor

Department of Physics and Astronomy

Brigham Young University

April 2016

ABSTRACT

Increasing Stability in Extended Cavity Diode Lasers using Frequency Noise Feedback

McKinley Pugh
Department of Physics and Astronomy, BYU
Bachelor of Science

Extended cavity diode lasers (ECDLs) have a number of useful applications, but they mode hop. We have observed an increase in frequency noise before mode hops in ECDLs. A feedback system using frequency noise instead of amplitude noise has been developed.

Keywords: diode laser, extended cavity diode laser, lock-in detection

ACKNOWLEDGMENTS

I would like to thank many people who have supported me and made this thesis possible. In particular, I would like to thank my advisor, Dr. Dallin Durfee. He offered me unending encouragement and patience as he mentored me throughout this project. Without his insights and willingness help me work through problems, the project would have been impossible for me. I would also like to thank Enoch Lambert, who started the project before me and who developed the original set up.

A special thanks goes to the many faculty of the BYU Physics and Astronomy Department who have taught and encouraged me during my time here. They always made me feel that undergraduate education was important to them, and that they wanted me to succeed. Thanks to them, my time at BYU was a true joy and felt immensely meaningful.

I would also extend my deep gratitude the BYU Physics and Astronomy Department, the BYU Office of Research and Creative Actives, and the NSF PHY-1205736 for funding this project.

Contents

List of Figures

Chapter 1

Introduction

1.1 Diode Lasers

Lasers are an important tool in atomic physics, particularly in applications such as laser spectroscopy, laser cooling, and precision measurement. To be used in these types of applications, a laser typically needs to have a stable frequency and a linewidth much narrower than an atomic resonance. Diode lasers are appealing because they are durable, compact, and relatively inexpensive [1]. They are available in a number of different wavelengths and are capable of quick modulations, making them viable for a number of applications from cancer therapy [2] to quantum computing [3]. However, diode lasers naturally have linewidths that are much wider than atomic transitions. This becomes a problem in applications like laser cooling where it is essential that the laser linewidth be much smaller than the transition it is exciting. Diode lasers also have a tendency to mode hop, causing abrupt changes in the frequency of the laser.

The linewidth of a laser is determined in large part by its laser cavity. In diode lasers, that cavity is short and has low finesse. These factors allow the laser to be compact and to modulate its frequency quickly. However, they come with the cost of giving the laser a wide linewidth.

In addition to widening the linewidth of the laser, the low finesse of the cavity also makes diode lasers very sensitive to optical feedback. Since the laser light escapes the cavity so quickly, any outside light that gets coupled into the laser cavity can have a big effect on the laser's frequency. This sensitivity to optical feedback can be used to control the laser, but it also can introduce a great deal of instability to the laser [4].

In order to understand how optical feedback affects the laser, we need to understand the laser's modes and mode competition. A mode is a three-dimensional description of the standing waves that occur inside an optical cavity. Laser cavities are basically an optical cavity with the gain medium inside, so they, too, have modes. Every mode in a cavity has a resonant frequency associated with it. The laser will lase at frequencies that match the resonant frequency of a mode but not at frequencies where there is no resonance.

The three-dimensional mode can be broken into two independent parts: transverse and longitudinal. Optical feedback can be used to control the longitudinal mode, but the transverse mode is very difficult to control. To get around this, labs typically purchase single mode lasers, i.e. lasers specifically designed to allow only one transverse mode. In these single mode lasers, the longitudinal modes remain unconstrained.

There are an infinite number of longitudinal modes in a laser cavity. However, the number of modes the laser can actually lase at are restricted by the gain curve of the gain medium. If a given mode has a frequency that is outside the energy of the electron transitions inside the gain medium, the laser cannot lase in that mode. The gain curves for diode lasers are very broad, though, so though they don't cover all the modes, they cover many, many modes.

If lasers were linear systems, they might operate in any combination of modes. However, lasers operate by stimulated emission, and this makes them non-linear. The process works like this: Lasers start with some spontaneous emission. Photons from these transitions can come from any frequency within that transition's linewidth, but if they do not match a mode of the cavity, they will not set up a standing wave in the cavity and the laser will not lase at those frequency. Other photons are emitted into a mode of the cavity, and they may excite many modes. This does not last long however, because of stimulated emission. As a photon moves through the gain medium, it will stimulate other photons to be emitted into the same mode it is in. The more photons there are in a mode, the more they stimulate other photons into that mode. Very quickly then, the mode with even marginally the most gain sucks up all the energy from the gain medium, and the side modes are suppressed. This process of one mode becoming more favorable and suppressing all the other modes is called mode competition.

The gain of each mode in the laser is not constant. If the environment—the temperature, the current, etc.—changes, the gain of the modes can change, and once a new mode starts getting more gain than the others, the laser will jump to that mode. This is called a mode hop. Mode hops cause sudden changes in laser frequency that can ruin experiments like laser cooling or precision measurement that depend on a stable frequency. Locking the laser, the usual way of stabilizing the frequency, can adjust for slow drift in the laser, but mode hops, with their sudden jumps in frequency, will throw the laser completely out of lock.

1.2 Extended Cavity Diode Lasers (ECDLs)

As noted before, diode lasers have wide linewidths because the laser cavity is short and has low finesse. Therefore, one way to decrease the linewidth is to effectively make the cavity longer and of a higher finesse. In extended cavity diode lasers (ECDLs), this is done by adding a reflective element outside the laser. The reflective element creates a longer cavity using the back of the laser cavity as the other mirror, and because it reflects light back into the laser, it decreases the loss in the cavity which raises the finesse. Together, these effects narrow the linewidth. However, adding the reflective element adds additional boundary conditions to the laser cavity, making the mode

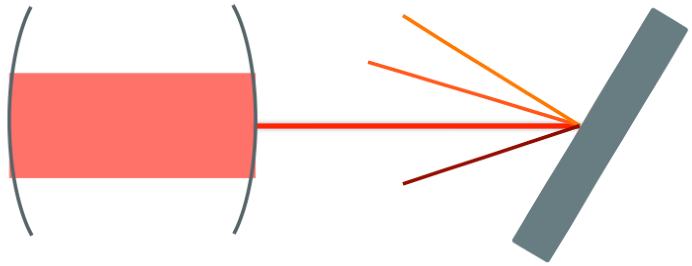

Figure 1.1 Light from the laser diode exits the laser cavity and hits a reflection grating. The reflection grating then reflects different wavelengths at different angles so that only one wavelength is reflected directly back into the laser.

structure more complicated and increasing the number of parameters that must be controlled to prevent mode hops.

Various designs for creating an extended cavity have been studied [1, 5, 6]. I used the Littrow method as shown in Figure 1.1. In this scheme, light from the laser exits the laser cavity and directly hits a reflection grating. The grating separates the light depending on wavelength. The 0^{th} order of the grating is directed away to the rest of the experiment. However, the 1^{st} order is reflected back to the laser, and the angle at which that 1^{st} order reflects depends on the wavelength.

The grating feedback to the laser is affected by both the grating's distance from the laser and the grating's angle relative to the incoming beam. For a narrow band of wavelengths, the reflection grating acts like a mirror. It creates the new cavity and so the position of the grating determines the length of the cavity which in turn determines the resonances. Therefore, fine changes in the position of the grating will change the wavelength of the laser.

Rotational changes to the reflection grating will change which wavelength will be fed directly back into the laser. This makes tuning the laser easier. With only a mirror creating the extended cavity, there are theoretically an infinite number of modes that resonate in the cavity. But the reflection grating reflects only one wavelength directly back into the laser. In this way, the position

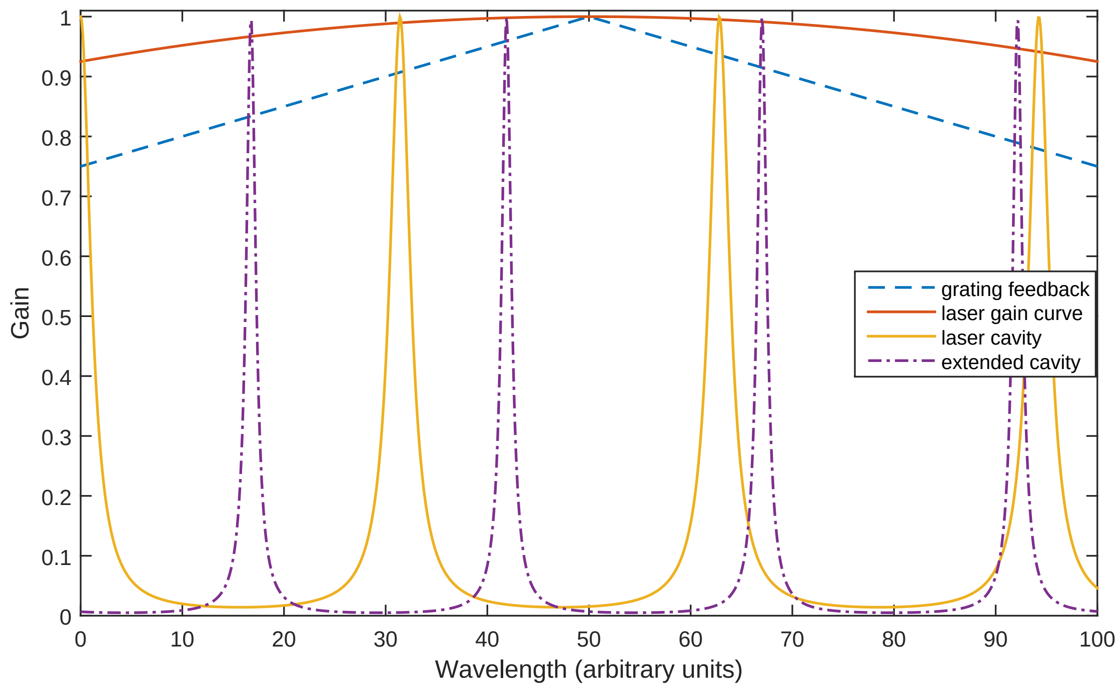

Figure 1.2 This an illustration of how the laser gain curve, the reflection grating angle, the laser cavity modes, and the modes of the extended cavity affect the gain inside the laser. To make an ECDL stable, the peaks of each of these curves must fall on top of each other. If one of the peaks begins to drift, the laser may go multimode or mode hop.

of the reflection grating determines the allowed modes in the extended cavity and the grating angle picks out which one of those modes the laser will lase in.

The extended cavity decreases the linewidth of the modes and allows one to tune the frequency of the laser. However, it also adds complexity to mode competition. A schematic in Figure 1.2 demonstrates this problem. When the laser is single mode and stable, the gain from all the factors—the gain medium gain curve, the reflection grating position and angle, and the laser cavity—are aligned and pushing the laser into the same mode. Over time, however these factors will drift, whether because of temperature change, mechanical drift, or some other change. Things do not have to drift very far before the gain in the laser changes and a new mode becomes more favorable. Then the laser mode hops.

1.3 Passive vs. Active Control

There are two types of modifications that can be made to an ECDL to increase its stability. These are called passive control and active control. Passive control attempts to increase the stability of the laser by decreasing changes in its environment. Some ECDLs have passive control stable enough that the laser can run a long time without mode hops. However, something—the current, the temperature, the reflection grating, etc.—will eventually drift. Once something has drifted too far, passive control can do nothing to correct it, and the laser will mode hop.

Passive control includes many factors. The most basic passive control will control the temperature of the laser, the injection current, and provide some sort of optical isolation. Commonly, the laser system is built on a floating optical table to decrease mechanical vibrations. In more advanced control, an anti-reflective coating is added to the laser diode to remove one of the boundary conditions and simplify the mode structure. These coatings are definitely effective, but they are also expensive [7, 8]. In other cases, the entire ECDL set up is machined out of one piece of metal and operated in vacuum to increase stability [6], and other set ups have been developed with significant passive stability. For instance, Ricci et al. obtained a linewidth of 100 kHz and had a scan range of 25 GHz [5]. Arnold et al. in their set up were to lock a 780 nm laser to a Doppler-free rudidium transition for several days [1].

Active control is different because it attempts to detect drift in these passive elements and correct for it. The major roadblock preventing active control is finding a usable measure for the drift that can be used to understand what drifted and how to adjust the laser to move it away from a mode hop.

1.4 Previous Work

In Steven Chu's lab at Stanford, amplitude noise on the laser was observed to increase before a mode hop. Using that noise, Chiow et al. were able build a feedback system that greatly increased the stability of their laser [8]. Previous students in my lab attempted to replicate their results, but found that the amplitude noise only predicted mode hops when the laser was locked to a frequency reference [9]. When the laser was free-running, the correlation between amplitude noise and mode hops was lost. Therefore, they reasoned that the true predictor of mode hops was not amplitude noise but frequency noise. As the laser approached a mode hop, the frequency noise increased and this noise appeared on the lock signal. Since the lock signal was directly used to correct the current, the noise was fed into the current and appeared as amplitude noise.

Noting that frequency noise, not amplitude noise, predicted mode hops presented an interesting opportunity. To measure amplitude noise, part of the laser beam had to be split off and extra optical components were required. However, since any lock system already has an error signal which measures frequency drift, frequency noise could easily be measured from the already present error signal. No additional optical components would have to be added. Instead, the system could be easily implemented wherever there was already a lock.

Previous work in our lab was able to use frequency noise to increase the scanning range of a laser [9]. However, the increases were not as large as expected from Chiow et al.'s results. We therefore believed that a great deal of improvement could be done to demonstrate the true power of using frequency noise to prevent mode hops.

1.5 Overview of Thesis

In Chapter 2, I discuss my experimental approach. I will describe my optical set up and the theory behind my approach, including how I predict a mode hop and how my frequency noise feedback will work together with a regular lock circuit. I will then explain how I measure frequency noise. In Chapter 3, I will show how frequency noise increases before a mode hop and will suggest the direction this project will take in the future.

Chapter 2

Experimental Approach

In this chapter, I describe my optical set up and how I how I determined a relationship between frequency noise and mode hops. I will then explain how I have taken that relationship and made a viable error signal (Section 2.3) and fed back to the laser (Section 2.4).

2.1 Optical Set Up

Figure 2.1 shows a schematic of my basic optical set up. In this section, I will explain each of the components in more detail, essentially following of the path of the light. First, I will describe in detail how I built my ECDL and established optical feedback. I will then describe the purpose and use of the optical isolator. A large section (Section 2.1.3) is devoted to semi-confocal optical cavities and their use in this set up. Finally, I describe the method I used to frequency lock the laser.

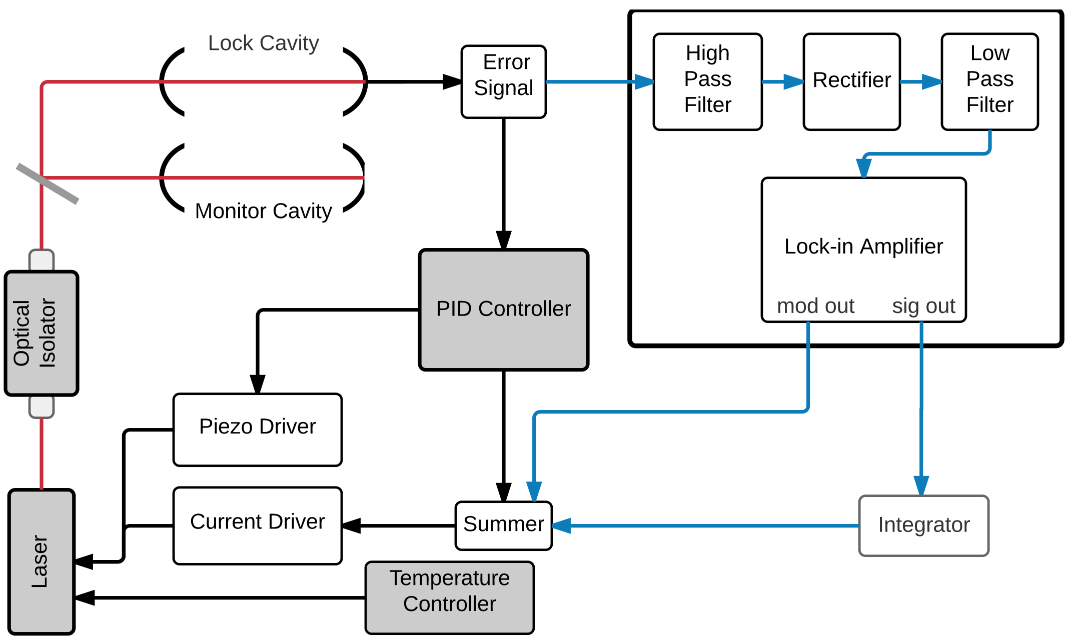

Figure 2.1 Light exits the laser and goes through an optical isolator. It is then split and enters two optical cavities. One cavity is used to monitor the laser's mode while the other is used to detect frequency noise. The black arrows indicate a typical lock circuit. An error signal comes from the lock cavity and is used by a PID controller to adjust to piezo controlling the reflection grating and to adjust the current. The blue lines indicate the new feedback I am implementing. The frequency noise is converted into a DC signal by a series of filters. A lock-in amplifier then uses lock-in detection to determine how to move the laser to keep it stable. Its output is put through an integrator. This signal as well as a modulation necessary for lock-in detection are added to the current signal from the PID controller, and together, all three adjust the current of the laser. A temperature controller stabilizes the temperature of the laser.

2.1.1 ECDL Set Up

The laser I used in my experiments was a US Lasers 650 nm 5 mW diode laser. This laser was certainly not the most stable I could have chosen, but since the goal of my project was only to *increase* the stability of a given laser, this laser was good enough. It was single mode laser, which was important. The laser could operate several minutes without a mode hop, and I was also able to scan the laser through several free spectral ranges of my optical cavities without a mode hop. I required no higher level of stability for the purposes of my project.

I created an extended cavity diode laser by using a few common optical components. The laser was housed in a Thorlabs TCLDM9 laser mount. The reflection grating was mounted on a piezo mount which allowed the grating to be moved and rotated electronically. The mounts were connected using aluminum rods.

The laser was temperature controlled using a built-in thermoelectric cooler (TEC), which I drove using lab built electronics. I controlled the current to the laser using a high stability, low noise current driver developed by Erickson et al. [10, 11] This current driver was also designed to allow a modulation to be added to the current.

2.1.2 Establishing Optical Feedback

When a reflection grating is put in an ECDL, it must be carefully aligned with the laser so that changes in the grating's angle and position affect the laser's frequency. This occurs when optical feedback is established meaning light reflecting off the grating couples into the laser cavity. To establish optical feedback, both the grating angle and the grating position need to be aligned. It is not enough for the grating to reflect light directly back into the laser; the wavelength it is reflecting back must also be an allowed wavelength in the extended cavity the grating has created. Otherwise, there are phase mismatches and no stable mode results. If there is no optical feedback, changes in the grating will not change the frequency of the laser.

Establishing optical feedback turns out to be rather simple. A laser becomes a laser when it reaches threshold—round trip gain is equal to round trip losses. As one increases the current, the gain increases until the laser passes threshold and lasing occurs. For this reason, diode lasers are often characterized by a threshold current or the current at which threshold occurs. Adding a reflection grating also increases the round trip gain by reflecting light back into the laser so losses are decreased. Therefore, if the threshold current is lowered, optical feedback has been established.

In practice, the process works like this: The current on the laser is increased until the laser hits threshold. This is sometimes described as the "flash" because the light coming from the laser suddenly gets brighter. The current is then turned down, just below threshold. Then, the grating is adjusted until the flash is seen again. The grating has fed back into the laser and increased the gain enough to reach threshold—even at this lower current level. Once this has occurred, optical feedback has been established. The current can be raised again as desired and adjusting the reflection grating will change the frequency.

2.1.3 Optical Isolator

Because diode lasers are so sensitive to optical feedback, it is important that the reflections making it back to the laser should only be those from the reflection grating. Other reflections from optics further on in the set up could disturb the mode of the laser. The optical isolator stops unwanted reflections from getting back to the laser.

Optical isolators allow light to only go through in one direction. They do this by taking advantage of the Faraday effect. Two polarizers are placed on either end of the optical isolator, set at 45° angles from each other. Between the two polarizers, the light is rotated 45° by the interaction between the light and a magnetic field in a medium, a phenomena called the Faraday effect. Though the Faraday effect is not the only way to rotate the polarization of light, it is unique in that it rotates all light the same direction regardless whether the light is propagating parallel or antiparallel to the magnetic field. This means that coming in one direction, the light enters through a polarizer, is rotated 45° and reaches the second polarizer polarized at exactly the right angle to make it through. Coming in the other direction, however, the light is polarized and again rotated 45°, but this time the rotation puts the polarization of the light exactly 90° off from the polarizer. No light comes through.

In my set up, the light coming from the laser is allowed through the optical isolator, but light hitting the optical isolator from the other direction does not make it to the laser.

2.1.4 Semi-Confocal Optical Cavities

Optical cavities are useful in this context because they can be used to observe the mode of the laser or to see small changes in the laser's frequency. I will first discuss the theory of optical cavities before going on to describe the particular qualities of semi-confocal cavities and their use in my optical set up.

Optical cavities at their most basic are two mirrors a distance L apart. Light bounces between the two mirrors in the cavity and at each bounce, some of the light is transmitted and some is reflected back as shown in Figure 2.2a. With each pass the light makes through the cavity, it picks up a phase shift. If the round-trip length is equal to an integer number of wavelengths, the transmitted light adds constructively. However, if the round trip length and the wavelength are not integer multiples, the phases of each pass do not match, and the light adds deconstructively.

In Appendix A, I go through the derivation of the transmission through an optical cavity. Here, I only quote the results. The transmission through an optical cavity is given by

$$T = \frac{T_{max}}{1 + F\sin^2(\phi)} \tag{2.1}$$

where T_{max} is the maximum possible transmission through the cavity, ϕ is the phase shift picked up after every round trip through the cavity, and F is the finesse. The shape of T as a function of ϕ

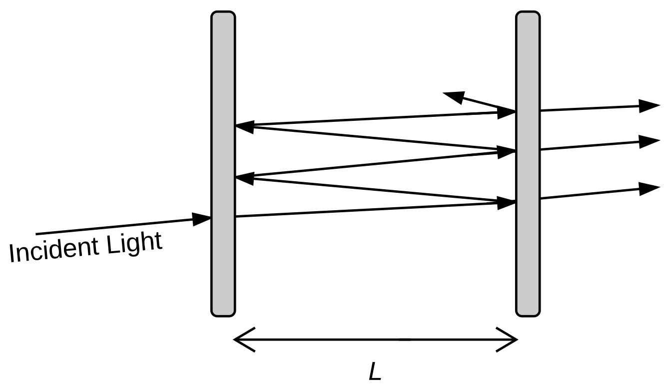

(a) Schematic of an optical cavity

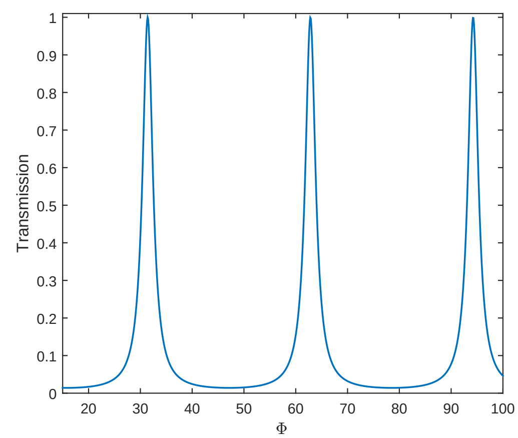

(b) Theroetical transmission as a function of ϕ

Figure 2.2 On the left, light enters the optical cavity and bounces between the mirrors, with some light being reflected and some transmitted every time it hits one of the mirrors. The relative phases of each of the passes determine if the transmitted fields add constructively or deconstructively. The graph on the right shows the theoretical transmission of light through the cavity as a function of the phase ϕ picked up after each pass through the cavity with $T_{max} = 1$ and $F = 200$.

is shown in Figure 2.2b. It is important to note that ϕ depends on both the wavelength of the light in the cavity and on the length of the cavity so changes in either will result in the peak-like pattern.

In order for the analysis done in Appendix A to hold true, we must have a stable optical cavity, a cavity where light bouncing between the mirrors does not "walk" away from the optical axis and out of the cavity. There are numerous possible configurations that create a stable optical cavity. In my experiment, I chose to use semi-confocal cavities.

Semi-confocal cavities are best understood by first explaining confocal cavities. In confocal cavities, two mirrors with the same radius of curvature are positioned so that the distance between them is equal to the radius of curvature. Therefore, their foci both hit at the same spot, halfway between the two. Semi-confocal cavities work similarly but instead of using two curved mirrors, they place one flat mirror at the focus of a curved mirror. This creates an optical cavity that is something like a confocal cavity folded in half. Two round trips through a semi-confocal cavity work out like one round trip through a confocal cavity. (See Appendix C.4 of [12] for more about semi-confocal cavities.)

Semi-confocal cavities are useful because of how the resonances of their modes line up. In other optical cavities, each of the modes has independent, possibly unrelated resonate frequencies. However, in a semi-confocal cavity, all the odd modes and all the even modes stack up, meaning that though many modes might be excited in the semi-confocal cavity, there will be only two resonant frequencies. This makes semi-confocal cavities especially useful for observing the mode of a laser. In other optical cavities, one must very carefully couple into the TE_{00} mode, and only the TE_{00}, to observe the mode of the laser since the different modes have different resonances. Not so for semi-confocal cavities. One simply has to couple to some of the modes of the cavity, and it does not matter which ones. This makes alignment much easier.

How does one use semi-confocal optical cavity to observe the mode of a laser? As we saw before, transmission through the cavity depends on ϕ, and we noted that both the length of the

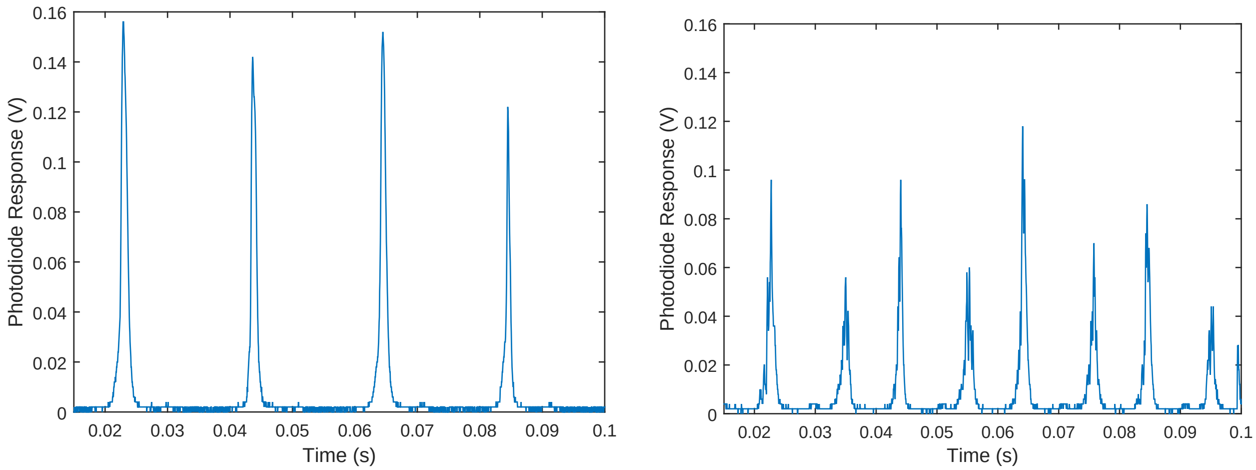

(a) Stable Laser **(b)** Unstable Laser

Figure 2.3 Data collected by scanning an optical cavity and looking at the response on a photodiode placed at the end of the cavity. Graph (a) shows the laser in single mode operation. In graph (b), the laser is multi-mode. Note how the peaks decrease in hight and the additional peaks that crop up in between peaks when the laser is multi-mode.

cavity and the wavelength of the light affect ϕ. As the length of the cavity is scanned, a photodiode can be placed on the end of the cavity to observe transmission. When the length of the optical cavity makes the round trip of the light an integer number of wavelengths long, the response on the photodiode is the greatest. Off resonance, there is little response.

When the laser is single mode, it has a very narrow linewidth centered on one wavelength. As one scans the length of the cavity, the signal on the photodiode resembles a number of peaks like those predicted Figure 2.2b. But should the laser begin to operate in multiple modes, light at new wavelengths would be produced. Those new wavelengths would resonate at different lengths of the cavities than before. On the signal from the photodiode, new peaks would appear. Figure 2.3 shows the output from a photodiode at the end of an optical cavity both when the laser is single mode and when other modes are starting to be excited. It is obvious at a glance to see the difference between the two.

Instead of changing the length of the cavity, one can also change the wavelength of the light through the cavity. Near resonance, very small changes in the wavelength cause big changes in the transmission. For this reason, optical cavities are frequently used as a frequency reference that lasers are locked to.

In my optical set up, I have two semi-confocal cavities, and I take advantage of both ways of using the optical cavities. With one cavity, which I call the monitor cavity, I use to observe the laser's mode. Using this cavity, I can see when the laser is single mode, when it is starting to go multimode, and, important for this project, I can see when the laser mode hops. The other cavity I call the lock cavity. I held the length of this cavity constant, and I locked the laser to this cavity. I was able to measure frequency noise on the laser by measuring the noise of the response of the photodiode at the end of the cavity.

Using both cavities together, I was able to observe both the frequency noise and the quality of the mode of the laser at the same time. That way I was able to establish whether there was a relationship between the frequency noise on the laser and the laser's tendency to mode hop.

2.1.5 Locking the Laser

To lock my laser, I used a side lock. Side locks are not the most effective ways of locking a laser, but they are one of the easiest to implement. To create a side lock, one takes the signal from a frequency reference—in my case, I used one of the peaks of my optical cavity—and subtracts an offset. This then creates a viable error signal (a signal that goes both positive and negative where the desired position is the zero-point) that can then be put into a PID control to feed back to the laser. Because it is so simple, a side lock can be implemented without an additional optics and only simple electronics. Side locks do have draw backs, however. For one thing, they cannot be used to lock to the peak of a resonance because otherwise the error signal would always be negative and never positive. They are also very susceptible to amplitude noise; since both changes in amplitude and changes in frequency change the response on the photodiode in the same way, the side lock cannot tell the difference between the two. In order to get a very tight lock, one would be better off using the Pound-Drever Hall lock [13]. However, my goal in this project was not to get the tightest lock possible. I only wanted to lock the laser so that I could observe its frequency noise before it mode hopped and went out of lock. If my method for increasing laser stability worked for a side lock, it would would for a Pound-Drever Hall lock, so I chose the simpler lock to avoid dealing with the unnecessary complications of a more sophisticated lock.

I used a PID circuit to generate feedback for the laser. The PID circuit took the error signal, processed it, and generated the feedback signal to control the lasers current and the reflection grating. By adjusting these parameters, the PID circuit moved the laser's frequency until the error signal was zero—and the laser was locked to the cavity.

In Section 1.2, I discussed how the reflection grating changes the mode in the laser cavity. The current also affects the modes by changing the index of refraction of the gain medium. Indeed, it is possible to lock the laser using only the current. However, the lock can tolerate more drift and the laser can be scanned further if the current and the grating are adjusted together.

It is also possible to adjust the grating's angle and position separately. This also allows the lock to compensate for more drift, and the laser can be scanned further. It is not necessary, however. Since I only had one piezo driver, I scanned the grating position and angle together. This limited my lock but was not a problem. My goal was to increase the stability of a given lock. A less-than perfect lock for a starting point was perfectly acceptable.

2.2 Predicting a Mode Hop

As noted before, amplitude noise predicted a mode hop for ECDLs, but only when the laser was locked. In order to test this, I used my lock cavity to look at the frequency of the laser. I used an oscilloscope to take the Fast Fourier Transform of the signal from the cavity to see the frequency noise on the laser. At the same time, I scanned the monitor cavity in order to observe the mode of the laser. By watching the monitor cavity, I was able to tell if the laser was single mode, multimode, or if it mode hopped. Therefore, I could observe both the frequency spectrum of the laser and its mode at the same time.

Rather than wait for the laser to mode hop on its own, I scanned the grating position and angle using a piezo driver. Small adjustments to the grating move the frequency of the laser, but the laser will not mode hop. Eventually though, the changes go too far, and the laser mode hops or goes multimode. Thus, scanning the laser essentially simulates the effect of letting the laser run for a long period of time but is more efficient. Scanning the laser also gives a certain level of control not available if the laser is free-running or locked because one can deliberately approach a mode hop and observe the laser's behavior as it comes near.

I observed the frequency noise both when the laser was locked and not locked in order to establish that the relationship to frequency noise was not a result of my lock. Unfortunately, due to the nature of optical cavities, I could see the noise much better when the frequency was near the cavity's resonance. When the laser was locked, the laser was always near resonance, but this was not the case when the laser was not locked. While this meant that the unlocked data were not as good as the locked data, they were still sufficiently accurate to be be used. Further discussion of the data will be described in Section 3.1.

2.3 Measuring and Using Frequency Noise

Once a correlation between frequency noise and mode hops is established, the problem is how to measure that noise and to feedback on to the laser. In this section, I will discuss both how I measured the level of frequency noise on my laser and how I used lock-in detection to feed back to the laser.

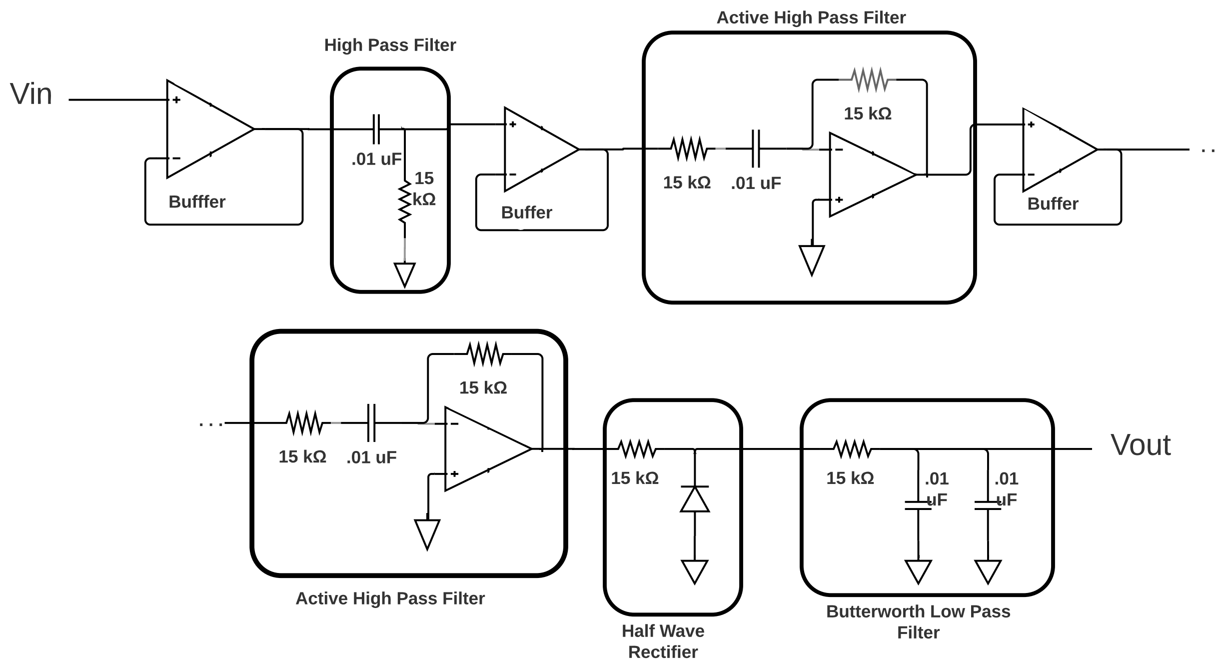

Figure 2.4 The signal passes through a series of filters in order to create a voltage proportional to the amplitude of the noise. The high pass filters need cut off frequencies significantly above the modulation frequency 90 Hz but also low enough to catch most the noise. The low pass filter has the same time constant so that the change in noise level due to the modulation is passed.

2.3.1 Converting Noise into a Usable Signal

In order to know how to feed back to the laser, we need a dc voltage that is proportional to the amount of frequency noise on the laser. The circuit I designed to do this is given in Figure 2.4. My original signal comes from the error signal for the laser's lock. As discussed in Section 2.1.4, my method of locking the laser includes adding an offset to the signal from the lock cavity. However, since I am only interested in the noise, all I need is a signal that changes with changes in frequency, which my error signal does. I do not want my signal to record the offset, so I pass the error signal through a high pass filter to remove the dc offset and ensure that the signal only contains the noise. I then put the signal through a half-wave rectifier and a low pass filter. The rectifier cuts the negative part of the signal, and the low pass filter smooths it out to create a voltage that is proportional to the amplitude of the noise.

In choosing time constants for my filters, I had to be conscious of two factors. The first was the shape of my noise. The frequency noise of my laser followed a $1/f$ shape, with most the noise below 10 kHz (see Section 3.1 for discussion). In some ways, this was good because I did not have to worry about getting high speed op amps that could capture very high frequency noise, say in the MHz or GHz range. I did have to make sure that I designed my high pass filter with a low enough cut-off frequency that I passed the band of frequencies I was interested it.

The other factor I had to take into consideration was the modulation of the laser. I will discuss lock-in detection in more detail in the next section (Section 2.3.2), but a key component is adding a small modulation to the current. For this application, I modulated the current at 90 Hz. This modulation in the current could potentially change the frequency of the laser. However, the PID circuit is fast enough to catch any drift and correct for it. The end result of the modulation, then, is to only make the laser a little more or a little less stable. So it should not show up on the error signal. However, a change in current also changes the amplitude of laser and since the optical cavity can't distinguish between a change in amplitude and a change in frequency, this modulation still shows up on the error signal.

When I pass the error signal through the series of filters, I want to see how the frequency noise is changing, not how the modulation is changing the amplitude. Therefore, I want my high-pass filter to have a cut-off frequency far above 90 Hz. Balancing these two demands, I settled for a cut-off frequency of 1 kHz. To further ensure that the modulation was completely killed, I used a second order high pass filter.

I used the same time constant for my low pass filter, this time to ensure that the change in noise level due to the modulation is passed through the filter.

Because the noise I was measuring had very small amplitudes, sometimes less than a mV, I also used active filters in the high pass stage. These filters produce gain, so the noise is big enough for the rectifier to work.

2.3.2 Lock-in Detection

A good error signal goes both positive and negative. Otherwise, it is impossible to tell by looking at the signal at a single point in time whether factors, like current to a laser, should be increased or decreased to move the laser away from a mode hop. But the amount of frequency noise is always positive. Adding an offset, like was done with the side lock, would make an error signal that went positive and negative, but it would lock the laser to a certain level of noise. My goal is to move the laser to the lowest possible level of noise. An offset won't do. Instead, I used lock-in detection.

Lock in detection is often used to extract a very small signal from a large amount of noise. It works like this: Suppose that the input to the lock-in amplifier is

$$V = V_i \sin(\omega_i t + \theta). \tag{2.2}$$

The lock-in amplifier generates a reference signal at ω_r

$$V = V_r \sin(\omega_r t). \tag{2.3}$$

The lock-in amplifier then multiplies the two signals together so that

$$V = V_i V_r \sin(\omega_i t + \theta)\sin(\omega_r t) \tag{2.4}$$

$$= \frac{1}{2} V_i V_r \cos((\omega_i - \omega_r)t + \theta) + \frac{1}{2} V_i V_r \cos((\omega_i - \omega_r)t + \theta). \tag{2.5}$$

Then the amplifier low-pass filters the signal, and this is where the interesting stuff happens. Low-pass filtering can also be thought of as time-averaging the signal. The average of $\cos(\omega t)$ is zero for any ω except in the special case when ω is zero. Therefore, the signal is always killed except for the case when $\omega_i = \omega_r$. In that case,

$$V = \frac{1}{2} V_i V_r \cos\theta. \tag{2.6}$$

This is fantastic because we have now generated a dc signal that is proportional to V_i and also $\cos\theta$—this depends on both the amplitude and the phase of the incoming signal.

Now suppose instead of a pure sine wave, the input is some function $V = f(t)$. If we expand $f(t)$ using a Fourier Series, we get

$$V = \frac{V_0}{2} + \sum_{n=0}^{\infty} V_n \sin\left(n\omega t + \theta\right). \tag{2.7}$$

When this goes through the lock-in amplifier, the only part that will survive will be the sine wave of frequency ω_r. The lock-in amplifier will be able to tell how much of the signal was that sine wave at ω_r and the phase of that wave.

When put in practice, lock-in detections pulls a very small signal out of noise by "marking" it with the frequency ω_r. The lock-in amplifier generates the reference sine wave, and it sends that wave to the device that will create the signal we are trying to detect. For instance, one might be trying to detect an LED on the other side of the room with a simple photodiode. The generated sine wave is sent to the LED and modulates its voltage. The intensity on the photodiode is then put into the amplifier. By multiplying by the frequency reference, the lock-in amplifier can throw out all the noise but keep the LED's modulated signal.

In my application, I am less interested in the amplitude of my signal but in the phase. I cannot tell which way to move the current if I only look at the noise level on my laser. However, if I move the current slightly up or down, I can see how that affects the noise. The lock-in amplifier does just that. The sine wave it generates dithers the current slightly. I then feed the measured noise into the lock-in amplifier, and it registers the phase of the signal. If the signal is in phase with the reference, that means that an increase in current increases the noise. If it is out of phase, then an increase of current decreases the noise. In phase corresponds to a positive $\cos\theta$; out of phase, a negative $\cos\theta$. I now have an error signal that tells me how changes in current change the noise, going both positive and negative.

I can then take the output from the lock-in amplifier and input it into a controller. In this case, I used a simple integrator.

2.4 Preventing a Mode Hop

With the whole system up and running, I have three sources of error detection adjusting the environment of the laser: the lock to the frequency reference, the frequency noise feedback, and the dithering from the lock-in detection. These all work together to keep the laser locked without a mode hop.

The lock to the frequency reference uses a PID controller and has the fastest time scale. The lock is unique in my set up because it adjusts the grating as well as the current. Within the lock, the current has a faster time scale than the grating. In my set up, I have one piezo driver controlling both the grating angle and the grating position.

The frequency noise feedback only adjusts the current. When the current is modulated, the PID controller compensates to keep the laser locked to the cavity. Because the PID controller changes both the current and the piezos, it will partially, but not entirely cancel the current modulation, doing the rest of the compensation with the piezos. That way, when the modulation increases or decreases the current, it causes the current and the piezos to change in such a way that the frequency of the laser remains the same, but the conditions which could lead to a mode hop change. If the frequency noise on the laser increases or decreases when the current is increased, the integral controller on the noise feedback will adjust the laser to reduce or increase the current such that the laser is moved to the point where the frequency noise is minimized.

Chapter 3

Results and Conclusions

3.1 Increased Frequency Noise Before Mode Hops

To test that frequency noise predicted a mode hop, I put the laser in single mode operation and then intentionally moved it towards a mode hop. I examined both the mode and the frequency noise of the laser at the same time by using both my optical cavities, one scanned and one held steady. Though mode hops are very fast occurrences, I could tell when the laser was approaching a mode hop because the peaks started to get smaller and wider. I was then able to look at the frequency noise of the laser when I knew it was coming close to a mode hop.

In Figure 3.1, I compare the frequency noise of the laser when it is stable and single mode to when the laser is approaching a mode hop. It is clear that the noise increases as the laser comes closer to a mode hop. The biggest changes are in the lower end of the spectrum. However, the noise baseline also rises. All of these data were taken with the laser unlocked. There was no chance, then, of noise being fed back from the lock as we suspect happened in Chu's experiment.

Because optical cavities are so sensitive to frequency changes, we believe this noise is due to frequency noise. However, it is possible that some of this is amplitude noise as well. The optical

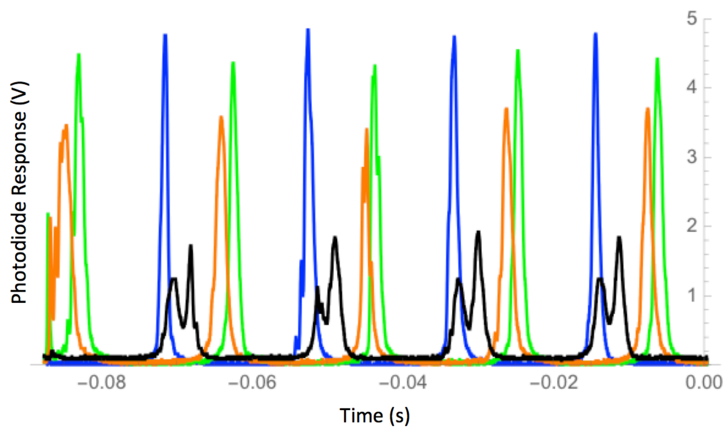

(a) Modes of the laser

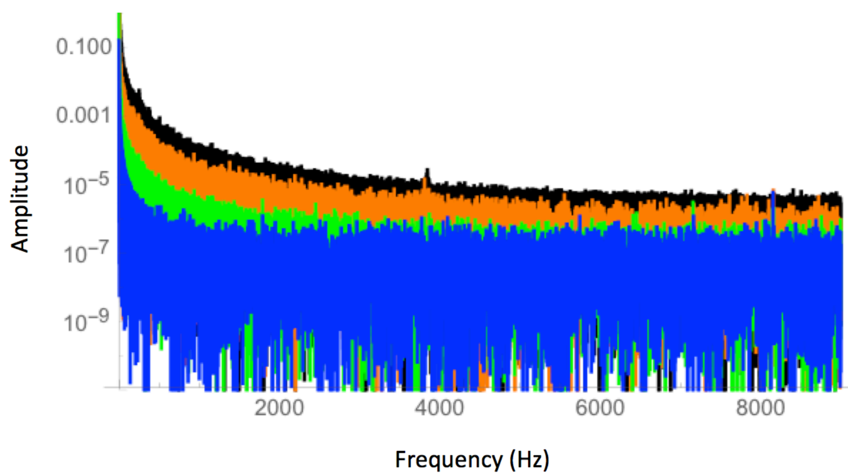

(b) Frequency Noise

Figure 3.1 The modes and frequency noise of the laser at four different times are shown. All were taken with the laser unlocked. In graph (a), the signal of the ramped cavity is plotted. In graph (b), the Discrete Fourier Transform of the un-ramped cavity is graphed on a log plot. The blue lines correspond to a very clean single mode. The black line corresponds to the worst mode, and the closest to a mode hop. It is clear that the cleanest modes have lowest noise. Note that as the laser approaches a mode hop, the noise increases the most in the lower end of the spectrum, though the baseline rises as well.

cavity as we have it set up cannot differentiate between the two. Further work will be required to completely separate the two kinds of noise.

3.2 Conclusions and Future Work

I have established that there is a correlation between mode hops and noise. I believe this noise is frequency noise. However, there is a chance that some of the noise I observed was amplitude noise as well.

Future work needs to more carefully pin down the frequency noise before a mode hop. For one thing, in my set up, I locked the laser to the cavity or didn't have a lock at all. This posed two problems. The first was that frequency noise far away from a resonance was hard to detect. The other was that changes in amplitude showed up exactly the same as changes in frequency. What I should have done was lock the cavity to the laser and monitored amplitude noise separately. This would have gotten rid of the frequency noise verses amplitude noise ambiguity. This would also keep the resonance of the cavity near the frequency of the laser, where frequency noise would show up better, without having to do any feedback to the laser. That way, any noise observed would certainly not be due to feedback from the lock. Unfortunately, I did not realize this until it was too late.

The next thing to do would be to plot the level of noise verses current. We believe that the noise should follow a U-like shape, with low noise in the middle where the laser is stable and higher noise on either side as the laser approaches a mode hop in either direction. I have observed that the noise does increase before a mode hop, but I do not know exactly what the noise as a function of current looks like.

Once the frequency noise of the laser is better understood and characterized, then it would be time to go back to the noise filtering circuit and the lock-in amplifier and start testing exactly how this feedback increases the stability of the laser. The stability could be tested both by seeing how long the laser could run without a mode hop and by seeing how far the laser could be scanned without a mode hop.

Previous work in our lab was able to increase the scanning range of the laser using frequency noise feedback [9] but not nearly as much as we expected, especially compared to the results from Chiow et al. [8]. I have made improvements to the optical set up and the electronics used in this previous work. In particular, I redesigned the circuit to measure noise so that it gave a much cleaner dc signal. Hopefully, with the improvements I have made, once the frequency noise is understood, setting up the feedback will go fairly smoothly.

Appendix A

Transmission Through Optical Cavities

The transmission through an optical cavity can be derived using complex analysis. We describe the incident light as

$$\mathbf{E}(z,t) = \mathbf{E_0} e^{i(kz-\omega t)}. \tag{A.1}$$

At each mirror, a fraction of the electric field t is transmitted and a fraction r is reflected. Therefore, after the light is transmitted through both mirrors, it has picked up a factor of t^2. With every round trip through the optical cavity, it also picks up a factor of r^2. Typically, when talking about transmission and reflection, we talk about the fractions of intensity, not electric field. These fractions of given the symbols T and R and are related to the portion of the electric field transmitted and reflected by the relations $T = t^2$ and $R = r^2$. Since the total fraction of light both transmitted and reflected is 1, T and R are related by $T + R = 1$. To avoid confusion with the total transmission through the optical cavity (which is also given the symbol T), we will replace r^2 with R and t^2 with $1 - R$ from here on on out.

In the cavity, we must also take into account that the electric field is a wave evolving in time as it travels through the cavity. This means that with each pass from mirror to mirror, the electric field also picks up a phase shift ϕ. This phase shift depends on the length of the cavity, the wavelength of the light, and the mode of the cavity that the light couples into.

Putting together the transmission, reflection, and the phase shift after n passes through the cavity, the transmitted electric field is given by

$$\mathbf{E}(z,t) = \mathbf{E_0}(1-R)R^n e^{i(kz-\omega t)+i(2n+1)\phi}. \tag{A.2}$$

The total electric field transmitted through the cavity is a sum of the transmitted electric field after every pass through the cavity,

$$\mathbf{E_{tot}}(z,t) = \sum_{n=0}^{\infty} \mathbf{E_0}(1-R)R^n e^{i(kz-\omega t)+i(2n+1)\phi} \tag{A.3}$$

$$= \mathbf{E_0}(1-R)e^{i(kz-\omega t+\phi)} \sum_{n=0}^{\infty} R^n e^{2in\phi}. \tag{A.4}$$

We recognize that the sum in Eq. 2.4 has the form of a geometric series $\sum_{n=0}^{\infty} a^n$ where $a = Re^{2i\phi}$. It can be proven that geometric series will converge to $\frac{1}{1-a}$ so

$$\mathbf{E_{tot}} = \frac{\mathbf{E_0}(1-R)e^{i(kz-\omega t+\phi)}}{1-Re^{2i\phi}}. \tag{A.5}$$

To see what we would measure on a photodiode at the end of the cavity, we need to calculate the intensity.

$$I \propto \mathbf{E}\mathbf{E}^* = \frac{(1-R)E_0^2}{(1-Re^{2i\phi})(1-Re^{-2i\phi})} \tag{A.6}$$

$$= \frac{(1-R)E_0^2}{1-2R\cos(2\phi)+R^2} \tag{A.7}$$

Using the double angle identity and some algebra, we can finally write the total transmitted intensity, T, in the form

$$T = \frac{T^{max}}{1+F\sin^2(\phi)} \tag{A.8}$$

where T_{max} is the maximum possible transmission and F is the finesse given by

$$F = \frac{4R}{(1-R)^2}. \tag{A.9}$$

Bibliography

[1] A. Arnold, J. Wilson, and M. Boshier, "A simple extended-cavity diode laser," Review of Scientific Instruments **69,** 1236–1239 (1998).

[2] S. Zakharov, I. Korochkin, A. Yusupov, V. Bezotosnyi, E. Cheshev, and F. Frantzen, "Application of diode lasers in light-oxygen cancer therapy," Semiconductors **48,** 123–128 (2014).

[3] M. T. Rakher, R. J. Warburton, and P. Treutlein, "Prospects for storage and retrieval of a quantum-dot single photon in an ultracold 87 Rb ensemble," Physical Review A **88,** 053834 (2013).

[4] C. E. Wieman and L. Hollberg, "Using diode lasers for atomic physics," Review of Scientific Instruments **62,** 1–20 (1991).

[5] L. Ricci, M. Weidemüller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. König, and T. W. Hänsch, "A compact grating-stabilized diode laser system for atomic physics," Optics Communications **117,** 541–549 (1995).

[6] E. C. Cook, P. J. Martin, T. L. Brown-Heft, J. C. Garman, and D. A. Steck, "High passive-stability diode-laser design for use in atomic-physics experiments," Review of Scientific Instruments **83,** 043101 (2012).

[7] D. Lonsdale, D. Andrews, and T. King, “Single mode operation and extended scanning of anti-reflection coated visible laser diodes in a Littrow cavity,” Measurement Science and Technology **15,** 933 (2004).

[8] S. Chiow, Q. Long, C. Vo, H. Müller, and S. Chu, “Extended-cavity diode lasers with tracked resonances,” Applied optics **46,** 7997–8001 (2007).

[9] E. Lambert, “Mode Stabilization of a Diode Laser Using Radio-frequency Lock Noise,” Senior Thesis (Brigham Young University, Provo UT, 2014).

[10] C. J. Erickson, M. Van Zijll, G. Doermann, and D. S. Durfee, “An ultrahigh stability, low-noise laser current driver with digital control,” Review of Scientific Instruments **79,** 073107 (2008).

[11] D. L. Troxel, C. J. Erickson, and D. S. Durfee, “Note: Updates to an ultra-low noise laser current driver,” Review of Scientific Instruments **82,** 096101 (2011).

[12] J. Archibald, “Construction of 408 nm Laser System for Use in Ion Interferometry,” Master of Science Thesis (Brigham Young University, Provo, UT, 2016).

[13] R. Drever, J. L. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, “Laser phase and frequency stabilization using an optical resonator,” Applied Physics B **31,** 97–105 (1983).

Index