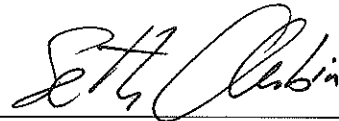


Thermal, Mechanical, and Electrical Engineering of an External Cavity Diode Laser

A thesis submitted in partial fulfillment of the requirement
for the degree of Bachelor of Science in
Physics from the College of William and Mary in Virginia,

by

Quinton Olson

A handwritten signature in cursive script, reading "Seth Aubin". The signature is written in dark ink and is positioned above a horizontal line.

Advisor: Prof. Seth Aubin

Williamsburg, Virginia
May 2019

Contents

Acknowledgments	iii
List of Figures	v
Abstract	v
1 Introduction	1
1.1 Motivation	1
1.1.1 Experiments with Ultracold Atoms	1
1.1.2 Laser Cooling Production of Ultracold Atoms	3
1.2 External Cavity Diode Laser	6
1.2.1 Narrow Linewidth	6
1.2.2 High Tunability	8
1.2.3 Low Cost	11
1.3 Thesis Structure	12
2 Existing ECDL	13
2.1 Laser Cavity	13
2.2 Control Boxes	14
2.3 Laser Housing	14
2.4 Method for Testing Effectiveness	14

3	Thermal Engineering	20
3.1	Diagnosis of Initial Problems	20
3.2	Implementing Solutions	22
3.3	Further Improvements	24
3.4	New Problems	27
4	Mechanical Engineering	37
4.1	Basic Problem and Solution	37
4.2	Implementing Solutions	37
5	Electrical Engineering	41
5.1	Ground Loop Supression	41
5.2	Integrating the Floating Power Supply	42
6	Results	46
6.1	Resonance Scans	46
6.2	Linewidth Analysis	48
7	Outlook and Conclusion	54
7.1	Outlook	54
7.2	Conclusion	55

Acknowledgments

I would like to thank my research advisor Prof. Seth Aubin for guiding me through my research, Andrew Rotunno and Shuangli Du for lending their aid whenever I was lost in the lab, as well as everyone else in the Aubin research group for creating a positive atmosphere that gave me inspiration and encouragement throughout the year.

List of Figures

1.1	Four Wave Mixing	2
1.2	Photo-Association Scan Data	3
1.3	Magneto-Optical Trap	7
1.4	Rubidium Transition Levels	8
1.5	ECDL Diagram	9
1.6	Linewidth Plot	10
1.7	Rubidium Absorbtion Spectrum	11
2.1	ECDL	16
2.2	Control Boxes	17
2.3	Housing for the Laser	18
2.4	Testing Setup	19
3.1	Thermistor and TEC Data	21
3.2	Difference Between TEC readout and External Thermistor	22
3.3	Comparison of Temperature Drift with TEC On and TEC Off	23
3.4	Poor PID Feedback Convergence	24
3.5	Decent PID Feedback Convergence	25
3.6	Good PID Feedback Convergence	26
3.7	Unknown Kicks	27
3.8	Long Term Temperature Data With No Insulation	28

3.9	Before and After Insulation	29
3.10	Long Term Temperature Data With Insulation	30
3.11	Long Term Temperature Data With Insulation and Laser Hole Open	31
3.12	Current Setup	32
3.13	Stabilized Temperature	33
3.14	Instability During Scan	35
3.15	Stability During Scan	36
4.1	Unassembled Control Platform	38
4.2	Front View of the Platform	39
4.3	Top View of the Platform	40
5.1	Initial Noise from Scan Box	42
5.2	Noise from Scan Box with Filter	43
5.3	The Floating Power Supply	44
5.4	Noise from Scan Box with New Power Supply	45
6.1	Scan Range from B. Halkowski's Thesis	47
6.2	Initial Small Scan Range	48
6.3	Initial Larger Scan Range	51
6.4	Linewidth Analysis Setup	52
6.5	Linewidth Results	53

Abstract

This thesis summarizes work to stabilize the wavelength of an external cavity diode laser (ECDL) for use in cold and ultracold experiments. This was done through thermal engineering, primarily tuning a proportional-integral-derivative (PID) feedback loop of a thermoelectric cooling (TEC) controller to stabilize the temperature inside the ECDL housing; mechanical engineering, by creating a control platform to make transportation more convenient and to reduce noise from ground loops; electrical engineering, through wiring a new, quieter power supply for a current and piezo feedforward scan system. to reduce noise in scan data from AC power supplies. Temperature variations are the largest contributors to unstable output wavelength, and substantial progress has been made in correcting these variations, as well as identifying and correcting other problems with the ECDL. The linewidth of the ECDL was determined by measuring its beat note with another narrow-linewidth laser.

Chapter 1

Introduction

1.1 Motivation

Ultracold atom experiments require tunable, high stability, narrow linewidth lasers. Cooling atoms to hundreds of microKelvin (μK) above absolute zero and then holding them at that temperature requires very precise laser control. The process, which is described in more detail in subsection 1.1.2, has the best results when using a laser that can output a single frequency for extended periods of time. Various experiments require precise control over lasers to manipulate atoms and collect data.

1.1.1 Experiments with Ultracold Atoms

Four-Wave Mixing

Four wave mixing is a process that requires a stable laser with a narrow linewidth. The phenomenon of four wave mixing occurs when two lasers, the control beam and the probe beam, of different frequencies that are each near to the transition resonance of a cloud of atoms converge at a small angle in the cloud of atoms (Figure 1.1). If the frequency difference between the two input beams matches the hyperfine splitting of the atom, then a third beam, known as the Stokes beam, will be generated from the cloud at an angle to the two incident lasers.




Figure 1.1: Four Wave Mixing: A diagram of the setup to achieve four-wave mixing, taken from [1]. The control and probe beams are directed into a rubidium vapor cell, generating a Stokes beam.

Generating the Stokes beam in practice requires the control beam and the probe beam to be exactly one hyperfine splitting apart, so a narrow linewidth is mandatory (a more in depth look at four-wave mixing can be found in [1]).

Photo-Association

Another ultracold atom process of interest is photo-association (PA), which benefits mostly from a finely tunable laser capable of a smooth scan across a range of frequencies. PA involves the generation of molecules in an ultracold environment that would be impossible to generate otherwise. The process begins when atoms with a slow drift approach each other at a low impact parameter. Under normal conditions, these atoms would undergo a normal collision. PA changes this interaction by increasing the energy of one of the atoms. When this happens, the two atoms combine to form an ultracold molecule.

The calculations to find the correct photon energy to achieve ultracold molecule production are very involved and can be inconvenient. A laser, on the other hand, can scan across a range of frequencies near the suspected PA frequency. Plotting the data, such as in Figure 1.2, and analyzing the results can streamline the search for




Figure 1.2: Example of Photo-Association Scan Data in ^{85}Rb : This is a figure taken from [2]. This is two instances of data taken from scanning the frequency of a laser, which is incident on a cloud of atoms, and measuring the number of atoms, with the frequency on the horizontal axis and the number of atoms on the vertical axis. When the PA frequency is reached, there is a drop in the number of atoms measured.

the exact PA frequency, and can be used to compare the expected and measured PA frequency (a more in depth look at photo-association can be found in [2]).

1.1.2 Laser Cooling Production of Ultracold Atoms

For any ultracold atom experiment to work, there needs to be a supply of ultracold atoms. Stable, narrow linewidth lasers are used to cool atoms and keep them at a low temperature. In this thesis, the primary atom used is ^{87}Rb . Laser cooling makes use of the transitions of the electron in the outer shell of ^{87}Rb , specifically between the $5S_{\frac{1}{2}}$ and the $5P_{\frac{3}{2}}$ states. Using lasers operating on this transition, we can manipulate the behavior of the absorption and emission of atoms to lower the energy to hundreds of μK (i.e. a walk speed of 10cm/s in Rb).

Magneto-Optical Trap

The laser cooling takes place in a magneto-optical trap, or the MOT (Figure 1.3). The MOT consists of the area where the atoms are collected, six lasers oriented in counter-propagating pairs along three orthogonal axes, and two coils of wire to generate a magnetic field. This magnetic field uses the Zeeman effect to split the atomic states into more energy levels.

Trap Laser: The trap laser is responsible for cooling and holding atoms in a process known as Doppler cooling. The primary transition from the $5S_{\frac{1}{2}}$ state to the $5P_{\frac{3}{2}}$ state shown in Figure 1.4 is the transition that an electron is most likely to make. The trap laser is detuned to have a slightly lower frequency than this transition.

When an atom is moving toward the laser, the laser is Doppler-shifted to a slightly higher frequency in the frame of the atom. Since the laser is detuned to have a frequency below the frequency of the transition level, the higher frequency means the atom is more likely to absorb a photon. Once the photon is absorbed, the electron eventually drops down to a lower energy state, the atom emits a photon, and by the law of conservation of momentum, the atom loses some speed in the process. When an atom is moving away from the laser, the laser is Doppler-shifted even farther from resonance, so the atom is less likely to absorb a photon.

Over time, the atoms lose more and more speed, and therefore thermal energy, until they reach an equilibrium energy in the range of $100 \mu\text{K}$ or so. Since this is an equilibrium state, the atoms will remain at the ultracold temperature until the laser is turned off or they interact thermally with an external object.

Probe Laser: The probe laser is used to measure the number and temperature of the atoms in the atom cloud. The process involves turning off the lasers holding the

cloud in place and turning on the probe laser. The probe laser is incident into the cloud and is tuned to a frequency that causes the cloud to fluoresce, which is picked up by a CCD camera. The number of atoms is calculated by the number of photons (the intensity of the fluorescence) while the temperature is calculated by finding the area of the cloud.

Since the trap laser and probe laser are not on at the same time, and they are already tuned closely to each other, it is not uncommon for the trap laser to also serve as the probe laser.

Repumper Laser: The rule for transitions between states is that the hyperfine F number is only allowed to change by $+1$, 0 , or -1 . As seen in Figure 1.4, these splittings only allow an electron in the $F=3$ level of the excited state to fall to the $F=2$ level of the ground state. However this splittings allow for a scenario where an electron transitioning to the excited state does not go all the way to the $F=3$ level of the excited state, but is instead excited to $F=2$ level of the excited state. When it is in the $F=2$ excited state, decaying to the $F=1$ level of the ground state is an allowed transition. When in the $F=1$ ground state, the trap laser no longer holds the atom in the cloud, and the atom will eventually drift out of the cloud.

A repumper laser is tuned to operate so that the electron can transition to an energy level where it can decay back into the $F=2$ level of the ground state and enter the cycle again. Since the transition to the $F=2$ excited state has a probability of occurring every time the electron is excited, and decaying into the $F=1$ ground state has a probability each time an electron enters the $F=2$ excited state, it is expected that over time the cloud will lose the majority of its atoms and disperse. The repumper laser is critical to maintaining the atoms in an ultracold state. The most probable use of our ECDL will be as a replacement or enhancement of the lab's current repumper

laser.

1.2 External Cavity Diode Laser

The external cavity diode laser (ECDL) is a laser that has the traits outlined in section 1.1. It is specifically designed to have a consistently narrow linewidth and be tunable while maintaining the linewidth. An ECDL is constructed with two primary sections, the diode laser and an external cavity, as depicted in Figure 1.5. Normally, a cheap diode laser would fall short of the requirements to be used in laser cooling and ultracold experiments. The addition of an external cavity greatly improves the effectiveness of the laser to the point where it can be used for laser cooling or as a laser for ultracold atom experiments.

1.2.1 Narrow Linewidth

A primary function of the ECDL design is the narrow linewidth, which is defined as the frequency/wavelength spread of the output light of a laser. An ECDL incorporates an additional external cavity that is many times longer than the internal cavity of the diode laser. The secondary cavity acts as a filter; extraneous wavelengths that are permitted by the primary cavity are eliminated if they are not permitted by the secondary external cavity. This smaller linewidth is shown in Figure 1.6, where both light sources are from a diode laser, but the ECDL has the additional external cavity. The linewidth can become small enough that measuring it in wavelength (nm) becomes impractical, so it is common to measure the linewidth in frequency (Hz).

In section 6.2 of the Results chapter we assess the effectiveness of our ECDL by measuring its linewidth using the beat note created when combining the light from two lasers.




Figure 1.3: Magneto-Optical Trap: A diagram of a magneto-optical trap (MOT) adapted from [2]. The path of the pump and repump lasers are marked by the arrows, and the blue sections are where the blue coiled wire generates a magnetic field. The bright spot in the middle is the trap, where atoms are collected as they are cooled.




Figure 1.4: Rubidium-87 Transition Levels: This figure, adapted from [3], shows the energy levels of the $5S_{1/2}$ and the $5P_{3/2}$ states. Electrons are only allowed to transition if the F state changes by 0, 1, or -1. Therefore, any electron in the $5P_{3/2}$ $F=3$ level can only transition down to the $5S_{1/2}$ $F=2$ level (dashed red line). This is the cycle the trap laser operates on (solid red line). If an electron is only excited to the $5P_{3/2}$ $F=2$ level (solid green line), then it can transition down to either the $5S_{1/2}$ $F=2$ level or the $5S_{1/2}$ $F=1$ level. A transition to $5S_{1/2}$ $F=1$ (dashed green line) would leave the electron in a state where the photons from the trap laser do not have enough energy to cause it to transition up; these atoms can then drift out of the trap. To counteract this, a repumper laser is tuned to operate on the $5S_{1/2}$ $F=1$ to $5P_{3/2}$ $F=2$ transition (solid blue line), where there is the probability the electron will fall back to the $5S_{1/2}$ $F=2$ level and return to the trap cycle.

1.2.2 High Tunability

The ECDL is also designed so that its output wavelengths can be tuned over several nanometers. The ECDL in this thesis uses the Littrow design, where the diffraction grating can be rotated to change the length of the cavity and therefore




Figure 1.5: ECDL Diagram: A diagram of an ECDL setup taken from [5]. The external cavity is created by the diffraction grating. The length of the external cavity is controlled by changing voltage to the piezoceramic mount. This allows for fine tuning of output wavelengths. The collimation lens aligns diverging light into parallel waves. θ is the angle of the diffraction grating, which splits the incident light so some is reflected back into the cavity and some is reflected as the output beam.

change the peak output wavelength. In many cases, the rotation is controlled by an electrical component called a piezoceramic, which changes size depending on the voltage applied across it. By changing the voltage across the piezo, the diffraction grating is rotated and the output wavelength is changed. Using this, it is possible to scan across a range of frequencies, and maintain the small linewidth while doing so. The drawback with the Littrow design is that changing the angle of the grating changes the angle of the output beam. The effects of this are most notable when coarsely adjusting the screws using a screwdriver (usually on a larger scale than microns). For the piezo scans used in applications such as PA, the change in output direction will usually be small enough to not have a noticeable effect on the ending position of laser at short (on the order of meters) distances.

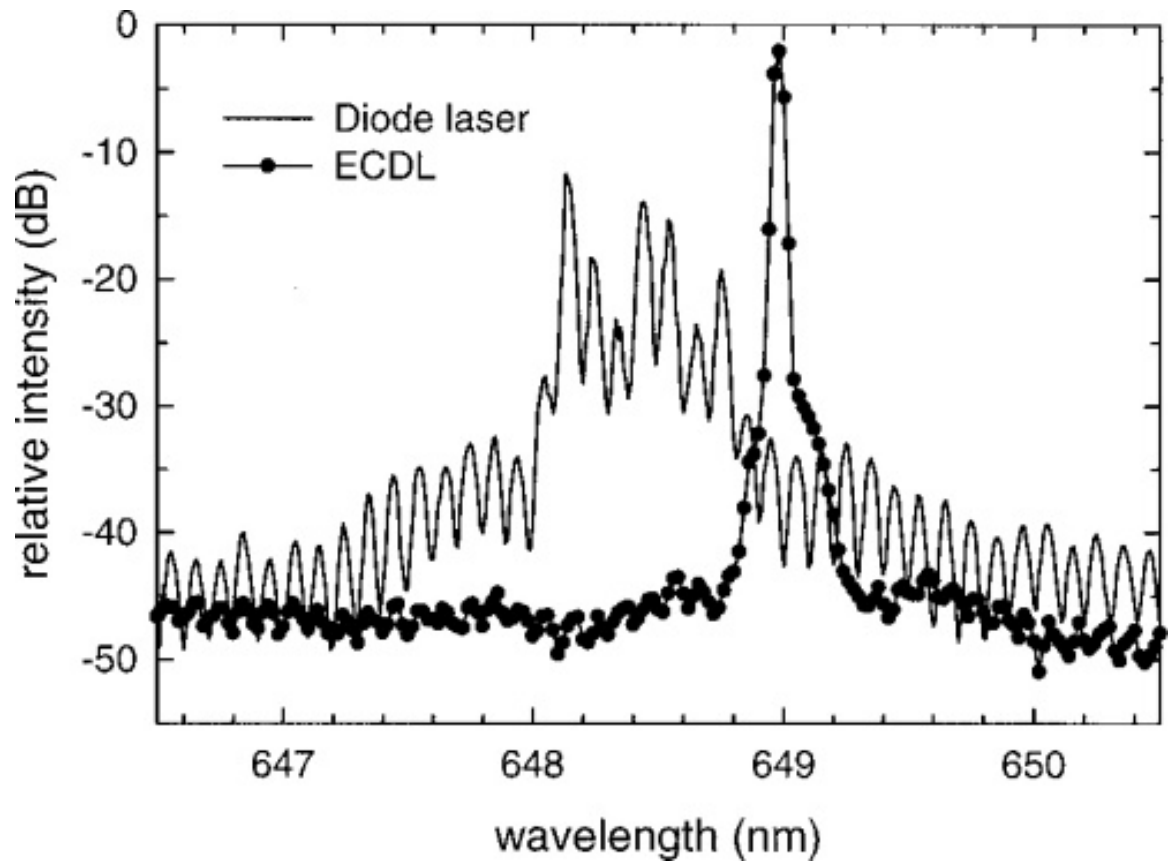


Figure 1.6: Linewidth plot (From [4]). By adding an external cavity, the linewidth of the output light from a laser diode significantly shrinks. The primary output wavelength for this ECDL is 649 nm, while the one we are constructing has a primary output at 780 nm.

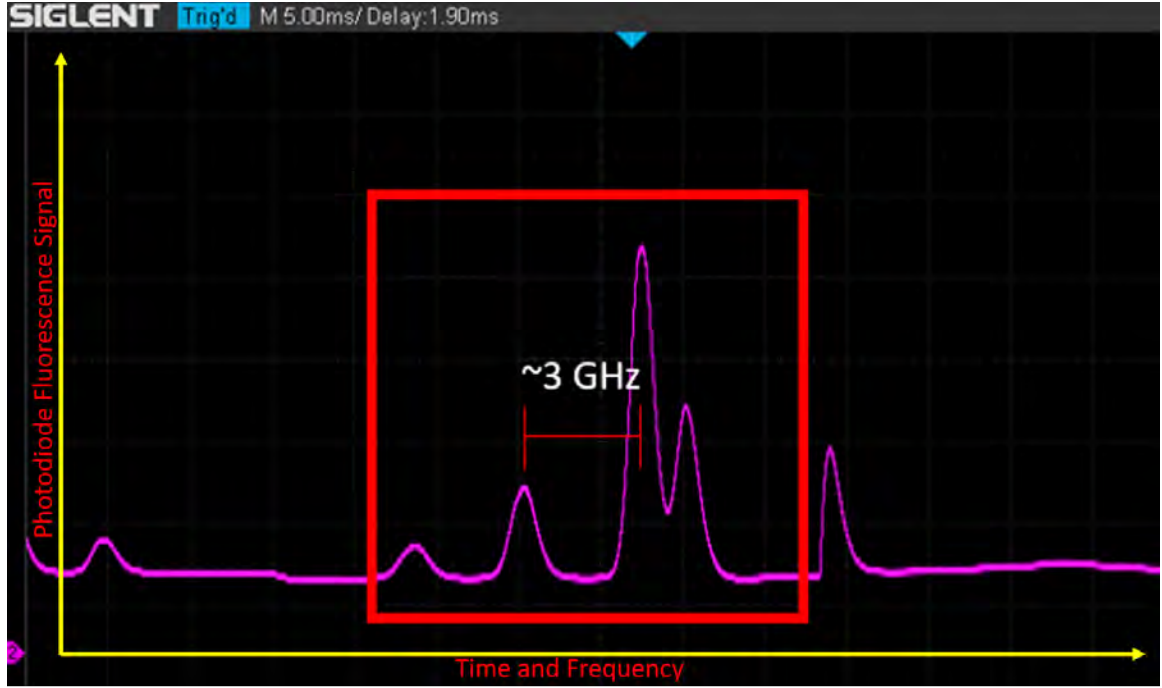


Figure 1.7: Rubidium Absorption Spectrum: A diagram of the absorption spectrum of Rubidium, taken from a scan using one of the lab’s commercial lasers (Toptica). The total length of the spectrum, and the goal for the scan range of the ECDL is on the order of 10 GHz.

In Section 6.1 of the Results chapter we assess the effectiveness of our ECDL by testing how far it can scan across the Rubidium 87 (^{87}Rb) spectrum, shown in Figure 1.7. ^{87}Rb is one of the primary elements used in ultracold physics, and was used by B. Halkowski in [5] to measure the scan range.

1.2.3 Low Cost

The last major benefit of an ECDL is the cost. Other lasers have been designed to counteract the problems that arise with non-ideal construction and operating conditions. This is usually achieved with higher precision and caution when creating and assembling the laser and/or using better materials to build the laser, which inevitably leads to higher prices. An ECDL is constructed with a laser diode and a diffraction

grating, as well as some control surfaces, that can cost significantly less than other lasers with similar functionality.

1.3 Thesis Structure

The remainder of this thesis is structured in the following manner. Chapter 2 gives an overview of the state of the ECDL when I began my work in September 2018. Chapter 3 explains the diagnosis of thermal fluctuations and the engineering steps taken to fix the causes of the fluctuations. Chapter 4 describes the design and construction of a laser controller platform. Chapter 5 gives details on steps taken to minimize electrical noise. Chapter 6 describes basic performance results, including scan tests and linewidth measurement. Finally, chapter 7 concludes the thesis and provides an outlook for future work.

Chapter 2

Existing ECDL

The laser we are working to improve is the one built by B. Halkowski in [5], based on earlier work by S. Ryan and E. Urbach. It comes with the ECDL already constructed, control boxes, and a housing for the laser. The ECDL had not yet been optimized for stability or scan range, and the linewidth was still unknown. The goal of this thesis is to make the necessary adjustments to the pre-existing components to optimize the performance of this ECDL.

2.1 Laser Cavity

The laser cavity, shown in Figure 2.1, consists of a diode laser, which sits on a thermoelectric cooler (TEC), a diffraction grating, which sits on a piezo mount, and a collimation lens between the grating and the laser diode. Previously, B. Halkowski aligned these components so that the peak transmitted wavelength is 780 nm. This adjustment was done coarsely through turning screws on the mount until the output wavelength was within fractions of a nanometer of the desired wavelength. The collimation lens has been placed so that the majority of the light is aligned after coming out of the laser diode. The notes left by B. Halkowski also describe early troubles with vibrations, so the aluminum breadboard sits on top of a rubber sheet (sorbothane). Adjacent to the side with the laser and the grating is a metal panel

with connections for the controllers and scan box. Additionally, there is a protection circuit for the laser diode against the panel.

2.2 Control Boxes

There are 3 control boxes and a feedforward scan box (Figure 2.2). The control boxes were stacked above the surface where the laser is located, while the scan box was not in use. The controller for the piezo voltage is on the bottom, and set to not exceed 100V. Its job is to adjust the voltage across the piezo, which turns the diffraction grating to change the output frequency. The controller for the laser current is in the middle and has a 105 mA limit. It controls the current going to the laser diode, which affects the output power as well as the output wavelength of the laser. The TEC controller is on the top, and has a 0.399 A limit. The TEC controller is located at the top because the control knob is very sensitive; an accidental touch can change the equilibrium significantly and cause the system to destabilize.

2.3 Laser Housing

B. Halkowski built a laser housing to help with temperature control (Figure 2.3). It sits on the plate below the rubber sheet and is held in place by components installed on the breadboard. It has a hole in the front for the laser output, and space on the back for the connection panel. From the lab notes, the functionality greatly improved after installing the housing, but it is nonetheless not airtight, so there are improvements that can be made.

2.4 Method for Testing Effectiveness

The method that B. Halkowski used to determine scan range and the time spent on resonance was to use a photodiode to measure the intensity of the light in a

rubidium vapor cell, shown in Figure 2.4. If the laser is on resonance at 780 nm, then the rubidium atoms are excited and fluoresce in the infrared spectrum. This light is picked up by both a photodiode, whose signal is sent to an oscilloscope for observation, and to a CCTV camera, which is sensitive to the emission wavelength, and displayed on a CRT television (Figure 2.4). Using this television, it is easy to tell if the laser is on resonance or not. In general, the television is used to tune the laser, and the oscilloscope can record long term data and scan data.




Figure 2.1: ECDL: The constructed ECDL consists of the laser diode, which sits on a thermoelectric cooling (TEC) device and is aimed at the piezo mount. All of the parts are fastened to a breadboard, which sits on a layer of rubber and aluminum to prevent vibrations from causing the laser to mode hop, a phenomenon where the peak output wavelength suddenly changes.



Figure 2.2: Control Boxes: The laser uses three controllers to operate, and also requires a feedforward scanbox, which was not in use at the beginning of the year. The top device is the the TEC controller, which is responsible for keeping the laser at a constant temperature. The middle device is the laser current controller, which changes the input current to the diode and can also be used to make small adjustments the output wavelength of the laser. The bottom device is the piezo voltage controller, which is used to control the input voltage to the piezo mount.




Figure 2.3: Housing for the Laser: The housing helps to insulate the laser from thermal fluctuations and air currents, both of which can cause changes to the cavity length on the order of nanometers to microns, which have a significant effect on the peak emitted wavelength. There are some gaps where the sides meet, and in the back next to the connector panel where the controllers connect with the components.




Figure 2.4: Testing Setup: This is the setup used to test the effectiveness of our laser. The laser comes from an optic fiber cable and shines through a tube containing rubidium vapor. The strength of the fluorescence is directly related to how close the laser frequency is to one of the resonance peaks. Data taking is done by sending the photodiode signal to an oscilloscope, where short term and long term behaviour can be observed. The camera is connected to a CRT television that displays the fluorescence in real time. When tuning the laser for resonance, we can use the TV screen as a way to check our progress in real time.

Chapter 3

Thermal Engineering

3.1 Diagnosis of Initial Problems

The strongest influence on varying wavelength and mode hops (i.e. sudden changes in output wavelength), appeared to be changes in temperature. Mode hops are easily identified when the fluorescence observed through the photodiode and on the TV screen disappears seemingly instantaneously rather than dimming to nothing. Through observing the controllers over a period of approximately 10 minutes, the piezo voltage and current controllers were both holding constant within the accuracy of the displays, but the TEC controller showed a fair amount of oscillation. Therefore, we began to investigate why the TEC showed temperature oscillations.

Our initial action was to make sure that the TEC controller's display was not showing phantom fluctuations due to faulty electronics or other undesirable sources.

To verify that the oscillations were due to real world variations, we compared the readout of the controller to the output of a thermistor placed on the TEC itself with the results shown in Figure 3.1. The two temperatures generally trend together, but there appears to be some difference between the two. In order to see if this was noteworthy or just random noise we analyzed the difference between the two readings in Figure 3.2. A linear fit gives a near zero slope indicating that the two




Figure 3.1: Thermistor and TEC Data: Data was taken concurrently with a thermistor and the readout on the TEC controller. This was necessary to determine if the oscillations that the controller was displaying were actually happening, or if the controller was picking up on phantom changes in temperature due to faulty circuitry. The data shows that both measurements basically vary together.

signals are correlated over the long term, so we concluded that the short term noise was random, and that the controller could be considered trustworthy. Since we were confident that the fluctuations in temperature were not due to faulty electronics, we could proceed to figure out what was causing the fluctuations. The two most plausible causes for temperature variations were errant air currents changing the temperature inside the housing, or a poorly tuned feedback loop on the controller preventing the temperature from stabilizing. Since we were not confident that the feedback loop had been carefully tuned before the year started, we decided that it was likely a factor in the temperature instability. To test this, we covered the ECDL with the housing and compared the oscillations in temperature with the TEC controller off vs on (Figure 3.3). There are noticeably more fluctuations with greater amplitude when




Figure 3.2: Difference Between TEC readout and External Thermistor (placed on TEC): The difference between the thermistor and TEC data was taken to see if there was any significant deviation in the rate of change between the two. The plot shows the difference remaining around 200 Ohms consistently with no real pattern in the variations, so we could confidently conclude that the readout from the TEC was accurate.

the TEC controller is on, so we decided that the next step to solving the temperature issues would be to more carefully tune the feedback loop.

3.2 Implementing Solutions

The TEC controller uses Proportional-Integral-Derivative (PID) feedback for temperature stabilization. The process of tuning this feedback involves manually changing the weight on three input dials on the control box that correspond to weights given to three parameters in the PID equation. Temperature control loops have a slow response time compared to electrical control loops, so the process of fine tuning the parameters to their ideal values is on a timescale of tens of minutes per adjustment. In order to streamline the process, we would identify similar patterns to tell if a change was for




Figure 3.3: Comparison of Temperature Drift with TEC On and TEC Off: To determine if the temperature was varying because of the controller or because of an outside source such as air currents, two readings of the same length were taken at 2 times. The first was taken with the controller off, and the second after the controller had been turned on and given time to settle. With the controller on, the temperature had a much larger amplitude and frequency of oscillation, leading to the conclusion that the feedback loop was the cause of the variations.

the better or for the worse. The quality of the tuning involves the time needed to settle at the desired value, the number of times the feedback overshoots the value, and the amplitude of the overshoots. Figure 3.4 shows a worst case scenario tuning. In this case, it is likely that the proportional or derivative or both proportional and derivative parameters were tuned too high. The initial settings were close to, if not slightly worse than Figure 3.5. With these settings, the laser would occasionally stay on resonance for about 5 minutes, which is passable, but not nearly enough to be reliably used in experiments. The final configuration can be seen in Figure 3.6



Figure 3.4: Poor PID Feedback Convergence: This is an example of a poorly tuned PID loop. Temperature is on the left axis, time is across the bottom. Over time, the system is not stabilizing, so operating the laser with the intent of having minimal or no mode hops is an excersize in futility.
(Temperature is on the left axis, time is across the bottom)

3.3 Further Improvements

When we were tuning the PID feedback loop, we noticed the stability of the system was disturbed on longer timescales by kicks from an unknown source (Figure 3.7). The disturbances were significant enough to cause the laser to mode hop and prevent long term data recording, so it was necessary to find how to prevent the kicks from happening.

First we needed to determine the cause of the kicks. One cause was potentially the feedback system causing itself to go haywire, so we decided to run tests with the TEC on but the feedback loop off. To achieve this, we found an approximate range of currents that the TEC controller supplied (.255A-.265A) and replaced the controller




Figure 3.5: Decent PID Feedback Convergence: This is an example of a decently tuned PID loop. Here the system is slowly settling in to a stable temperature, but the oscillations at the end are still too large. The peak to peak at the right part of the screen is somewhat greater than .04 Celcius. This is approximately what the system looked like before adjustments, and mode hopping is still a problem.

with a power supply with a constant current output. The leads to the TEC's internal thermistor were put into a multimeter with the capability to internally store data. To determine the constant current that would keep the TEC at 15.480 kOhms, we set the current source to .260 A, waited to see where the temperature settled, and repeated until we had stability at .229 A. From there, we recorded data from the multimeter connected to the TEC and a second multimeter connected to a thermistor that was suspended inside the housing but not in contact with the TEC. The initial data, shown in Figure 3.8, shows significant variations in temperature of both the suspended thermistor and the TEC. We decided the best course of action would be to better insulate the housing by taping over smaller holes and using padding to cover the back, which had both holes in the backplate and significant gaps between the




Figure 3.6: Good PID Feedback Convergence: This is an example of good tuning on a PID loop. While it is not optimal, it still decays in a reasonable time, and the final oscillations are within .04 Celcius, which was the target stability. A perfectly tuned system would have no oscillations whatsoever, but time is an issue, as each tune-and-check iteration takes around 10 minutes.

backplate and the housing (Figure 3.9).

After insulating the housing, we repeated the measurements using the same multimeters on the same timescale to see what improvements had been made (Figure 3.10). The results clearly showed that air currents had been the source of the kicks. We further tested the temperature on the TEC with constant current with the laser hole open (Figure 3.11). After some iterations, we found that using a pen tube for the opening (Figure 3.12) allowed the laser to pass through while minimizing the air currents getting into the housing.

After these improvements, the temperature instability is greatly reduced from before. Figure 3.13 shows the temperature of the TEC with the controller connected.




Figure 3.7: Unknown Kicks: The system is disturbed by an unknown source (Data taken over a long period of time, approximately 23 minutes). The resultant temperature variations are of a large enough magnitude to unsettle the laser to the point of mode hopping. Where the oscillations in the feedback loop are on the order of 10 s period, the kicks are closer to 5 minutes apart. The large discrepancy in the timescale means that there may possibly be another source of destabilization other than the feedback loop.

Currently the laser can maintain resonance for upwards of one hour.

3.4 New Problems

After reaching a promising point with stabilizing the temperature while the laser is at a constant output, we decided to test how the temperature behaves during a scan. We noted that sometimes the system would get a temperature from anywhere between 3 and 5 Ohms when turning on the function generator or from changing the settings on the current controller or the potentiometer on the front of the scan box. Figure 3.14 shows the temperature returning to stability after a kick started by one of the aforementioned cases. The problem is that when the kick occurs, the laser tends




Figure 3.8: Long Term Temperature Data With No Insulation: We took temperature data over the course of 70 minutes using the multimeters connected to both the TEC and to the suspended thermistor. The temperature from the suspended thermistor in particular has variations at a high frequency and amplitude. In these variations, there appear to be peaks, likely caused by air currents, that could potentially be the source of the kicks. The temperature from the TEC also has peaks at approximately the same times, but they are not as pronounced as the ones in the other plot. The two otherwise appear to be independent of each other to a degree, but it is difficult to tell if that is expected with only one set of data to analyze.

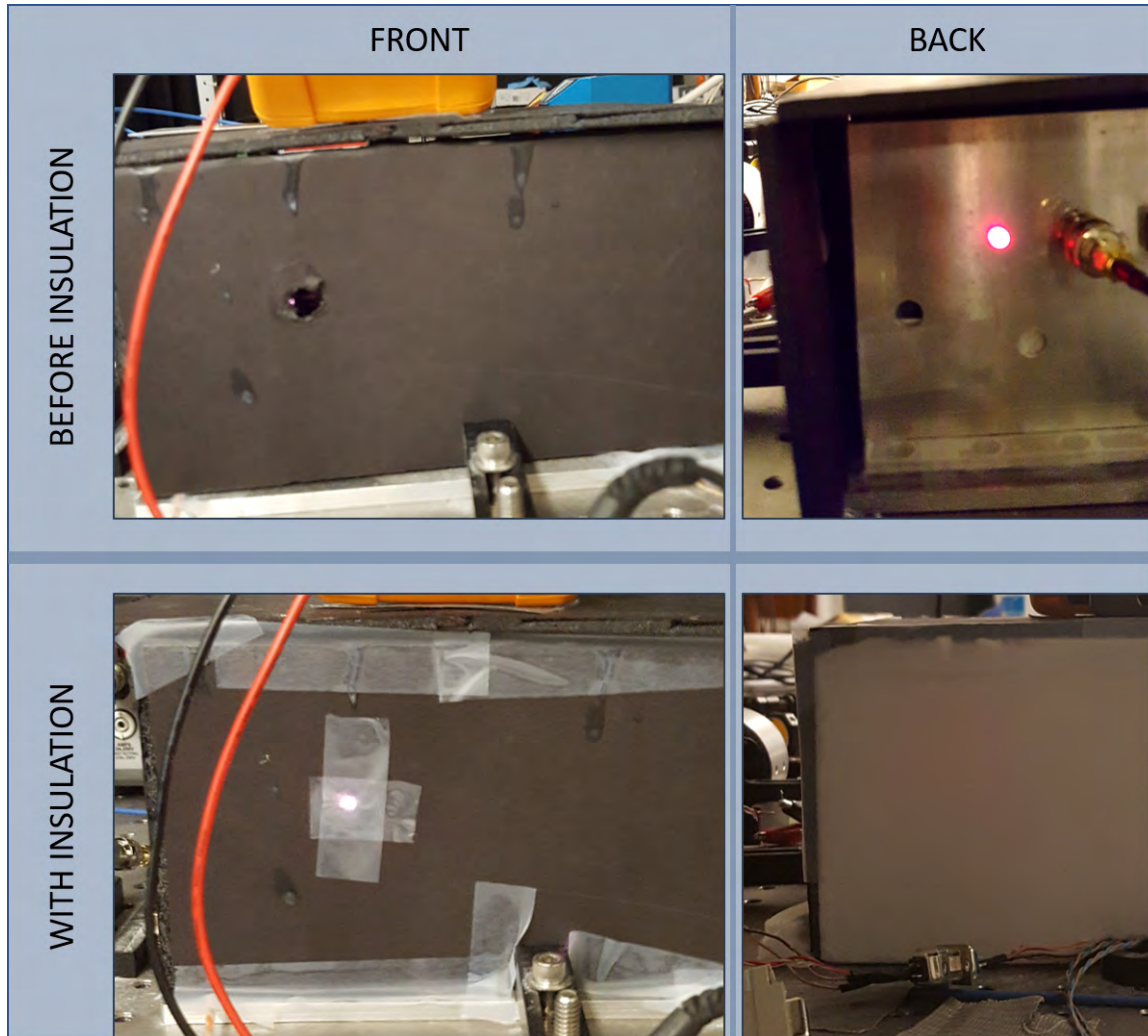


Figure 3.9: Before and After Insulation: The housing lid does not sit flush with the walls, the corners of the walls have some gaps, and the backplate does not sit close to the walls and lid. The most straightforward fix was to simply put tape over all the areas with holes, and since there was so much to cover in the back, to fit some padding in the back. We also covered the laser hole, to get a benchmark for a completely sealed housing. Since we believed the issue came from air currents, we were not primarily concerned with the optical output; insulating the housing so air currents could not get in was the main priority.

to mode hop, and achieving resonance while scanning is much more difficult than when scanning is off. These kicks are inconsistent, as shown in Figure 3.15 sometimes when turning the scan on there is no kick. The theory we have is that the changing




Figure 3.10: Long Term Temperature Data With Insulation: The data taken with the housing sealed from air currents. The time taken and the scale on the vertical axes are the same as the counterparts in Figure 3.8. In this case, there is almost no short term variation, a stark contrast to the other data taken. This data was enough for us to decide that air currents had been the primary source of the kicks.




Figure 3.11: Long Term Temperature Data With Insulation and Laser Hole Open: It would be impossible for us to use the laser with semi-opaque tape blocking the laser hole. We decided to test how the system behaves with the hole open, and found that while it is an improvement over the original housing, it was on the border of acceptable variation. To further insulate the cavity from air currents, we added a tube from a scrapped pen, as seen in Figure 3.12.




Figure 3.12: Current Setup: The current iteration of the housing setup. The insulating padding and pen tube prevent most of the air currents from reaching inside the box, while still allowing the laser to be transmitted and used in the lab. The length of the pen tube restricts the directions air currents can enter into the box, which reduces the volume of incoming air, so there is more overall insulation than if a hole was open on the side of the box.

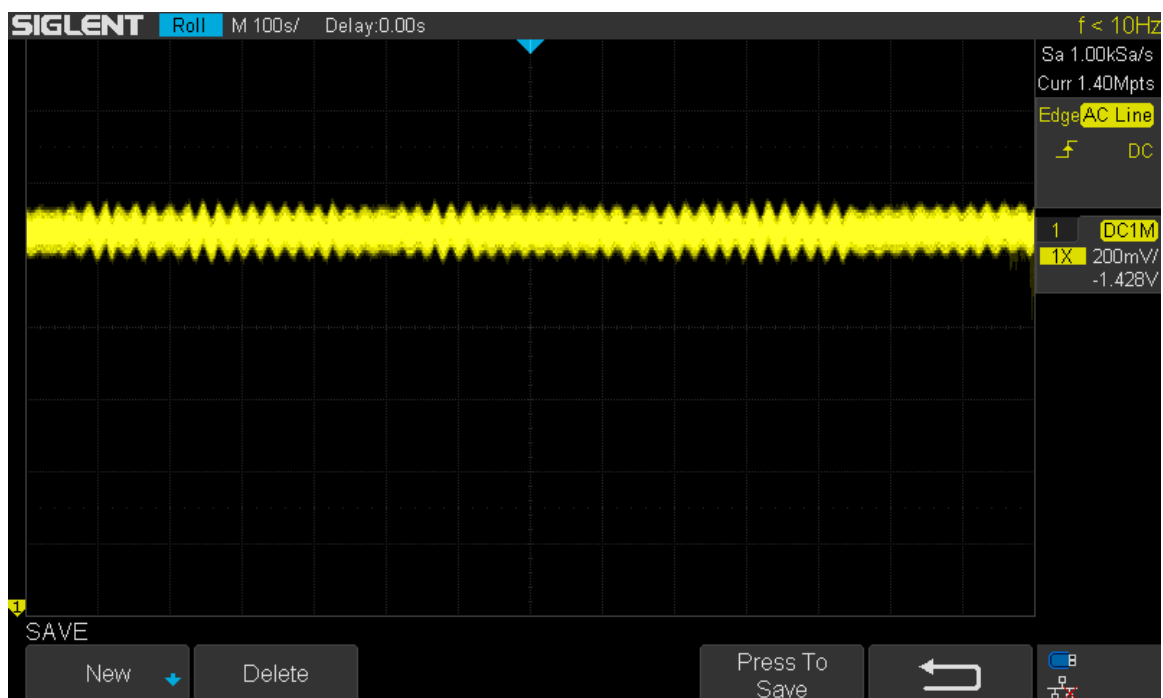


Figure 3.13: Stabilized Temperature: The data for this screenshot was taken for the same length of time with the same vertical scale as Figure 3.7. In this case, the laser varies by less than 1 Ohm, which is within the target range of variation.

current changes the thermal power released by the laser, which in turn changes the temperature of the laser, which causes the feedback loop to engage and the system receives a kick. The proposed solution is to possibly affix some thermal mass to the laser diode itself, so when the power changes, most of the thermal energy is given to/taken from the thermal mass and not from the TEC.

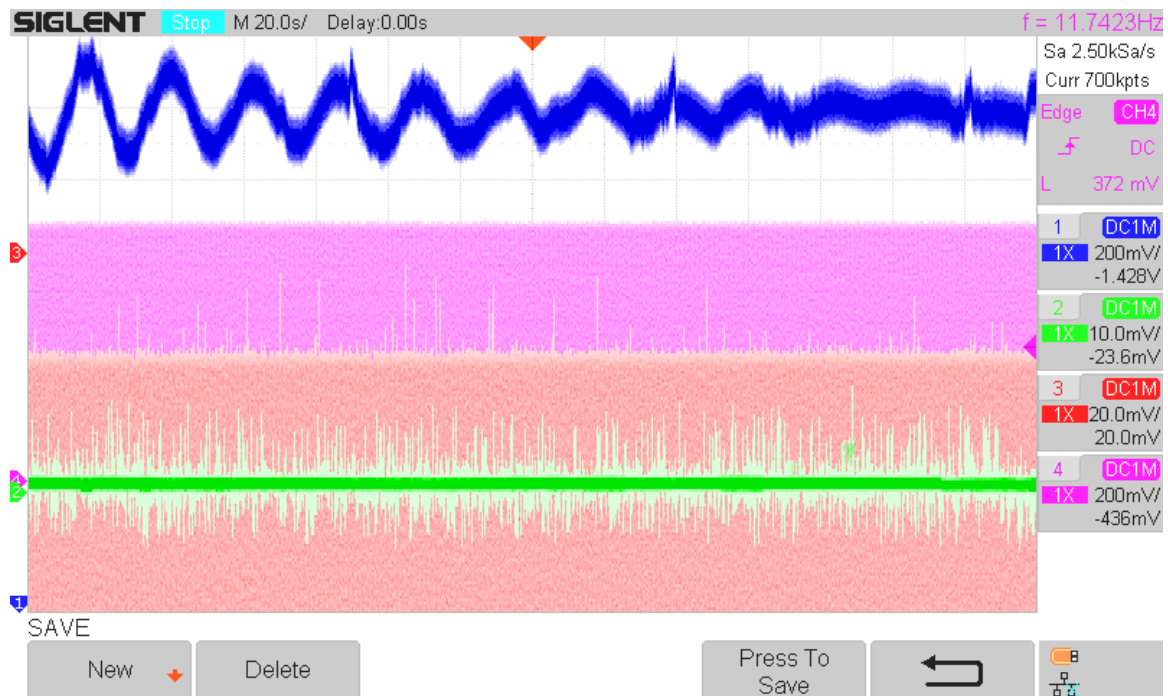


Figure 3.14: Instability During Scan: When the scan was started, there was a kick to the system. The amplitude of the kick is comparable to the amplitude of kicks before the insulation, but the system damps down to a more stable state. Kicks like this cause problems because the laser will mode hop during the large amplitude oscillations. There are 20 seconds per division, so it takes around 220 seconds to return to equilibrium. When this happens the current solution is to stop scanning and wait for the oscillations to die down before trying again.

Signal Sources:

Blue: Temperature

Purple: Piezo

Red: Current

Green: Photodiode

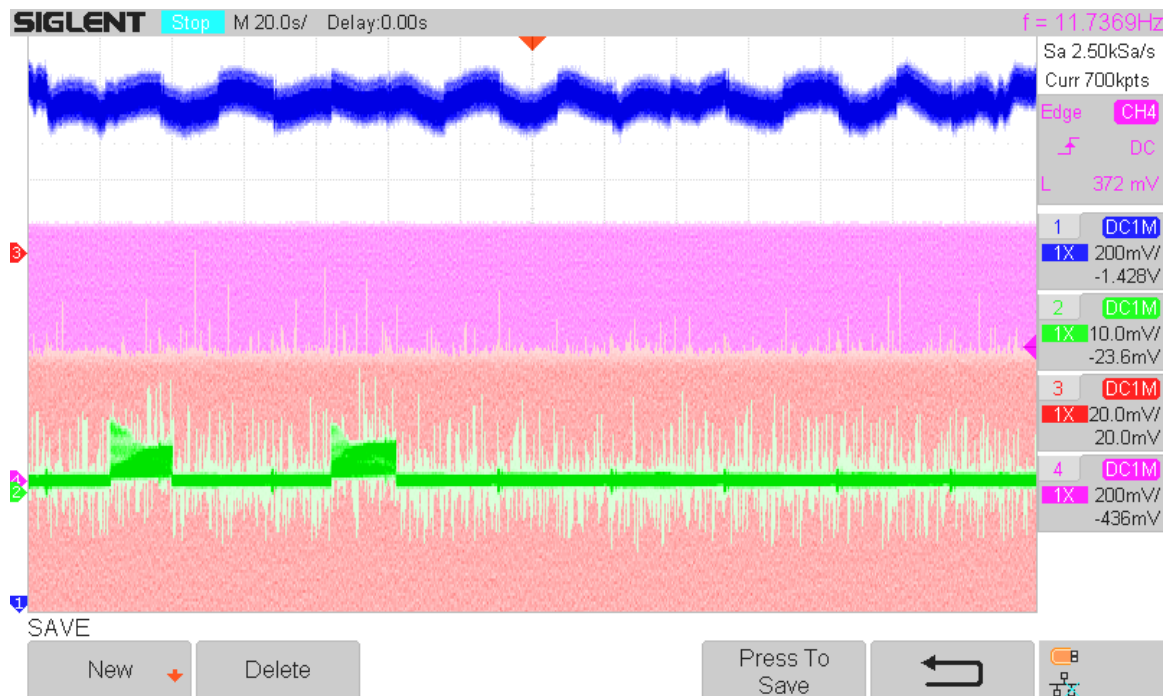


Figure 3.15: Stability During Scan: There are times where there is not a kick associated with beginning a scan. This screen shows that during the scan, the temperature is more unstable than when the scan is off. However, the variations are small enough that there is a low chance of mode hopping occurring.

Signal Sources:

Blue: Temperature

Purple: Piezo

Red: Current

Green: Photodiode

Chapter 4

Mechanical Engineering

4.1 Basic Problem and Solution

The temperature, current, and piezo controllers are all separate units, as is the lab-built feedforward scan box. There is no existing physical connection between the units, so transportation is difficult. The units being separated can also potentially contribute to additional noise in the form of ground loops, which are alluded to in the lab notes of B. Halkowski from 2017-18. Furthermore, the feedforward scan box does not have a power supply (B. Halkowski used 9 V batteries to suppress ground loop noise).

4.2 Implementing Solutions

The solution to all of these problems is to integrate and power all of the control units into a single device. The batteries used to provide power to the feedforward scan box can be difficult to change when they run out of power, so a small floating linear power supply is added. The basic parameters and approach for the open frame integrated system of controllers is the following:

- Edges must be secure so control boxes do not fall off.
- There must be enough room for fresh air to access the controllers to prevent



Figure 4.1: Unassembled Control Platform: These are the basic pieces of the control platform. The sides and back use the same sized L-brackets, while the front uses a smaller L-bracket to allow access to the controls on the piezo voltage controller and the laser current controller on the bottom of the stack.

overheating.

- There must be enough room for cables.
- There must be enough room for the power supply and its housing.

We began with a sheet of aluminum measuring 1'x2'x.125". We cut this down to 12" x 18" to form the baseplate. Next, we cut L-brackets to fasten to the four sides (Figure 4.1).

In order to assemble the platform, we used 10-32 screws to fasten the parts together through the use of tapped holes, untapped holes, and nuts. In general, the screws were chosen so the profile would minimally interact with the components on the platform. In order to consolidate space, a power strip was affixed to the back of the platform. All of the control boxes and the power supply for the scan box are plugged into the




Figure 4.2: Front View of the Platform: A photo of the front of the control platform. This is the view that the researcher would see when adjusting the settings on the controllers and while changing the scan box settings. The devices are (left to right, top to bottom) the feedforward scan box, the TEC controller, the piezo voltage controller, and the laser current controller.

power strip. Additionally, two vertical plates were attached on the left and right side to aid with the stability of the stacked elements. Figures 4.2 and 4.3 show the final platform with the housing removed from the power supply.

This platform increases the ease of transportation of the control systems, and the improvements in the ground loop problem are further described in Chapter 5.



Figure 4.3: Top View of the Platform: A photo of the control platform from above. This shows the layout of the control boxes as well as the location of the power supply.

Chapter 5

Electrical Engineering

5.1 Ground Loop Supression

Ben Halkowski had noted during his scan attempts that there was a considerable amount of electrical noise that was not of thermal or acoustic origin (Figure 5.1). One step he took to fix this noise was to add a low pass RC filter (Figure 5.2). This took care of the effects, but did not address the root of the problem. In his notes he mentioned that through changing the lab bench power supply to batteries and moving the feedforward scan box closer to the controllers, the noise was greatly reduced. This indicated that a strong candidate for the source of the noise in Figure 5.1 was the long cables carrying the signal from the scan box to the controllers created a ground loop. A ground loop can serve as an anteanna that picks up on noise from stray signals and transmits them to the system.

While source of the noise had been dealt with, batteries are innefficient and inconvenient, and constantly changing them would be a problem, so we decided to search for a better solution. Therefore, as discussed in the previous chapter, we added a floating power supply to the integrated control system for convenience and to help suppress possible ground loops.

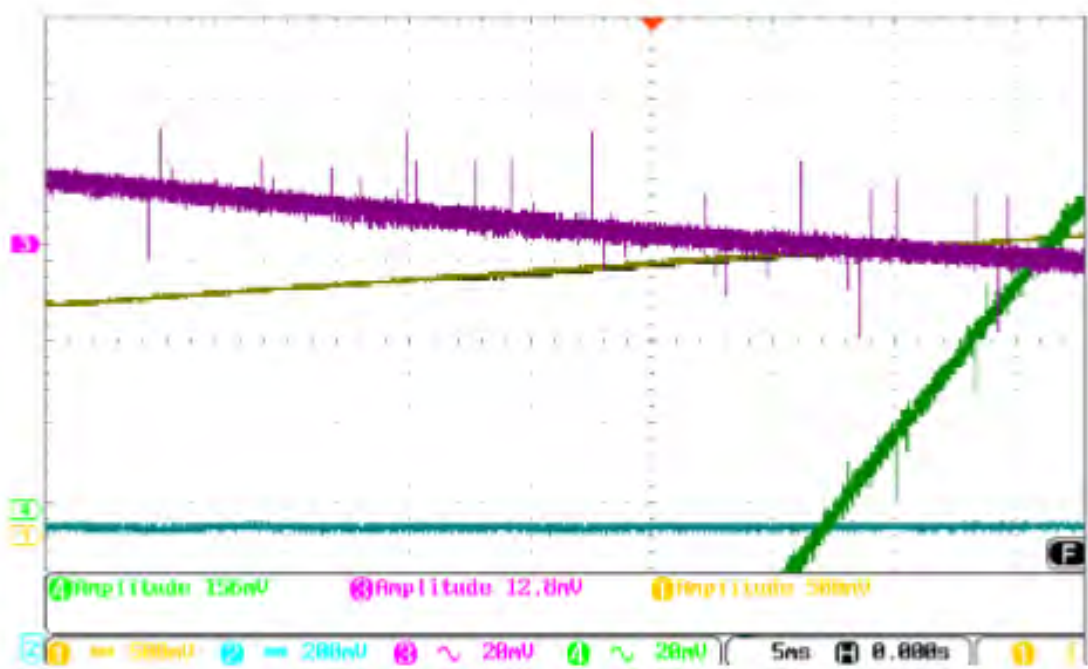


Figure 5.1: Initial Noise from Scan Box: This is a snapshot from [5] of an oscilloscope screen showing the noise on the current (green) and piezo (purple) signals before attempts were made to dampen the noise. He theorized that the spikes on both signals may be the cause of mode hopping, so he added a low pass RC filter to the signal.

5.2 Integrating the Floating Power Supply

The intent behind using an integrated floating power supply is that it should have reduced noise compared to a non-integrated lab bench power supply while being more convenient and efficient than batteries. The power supply consists of a transformer and circuitry to step the incoming AC voltage from the wall (120V 60Hz) and convert it to a constant 12 volt output. While the internals do most of the work, it was necessary to:

- Add an on/off switch.
- Connect the correct terminals on the transformer together.

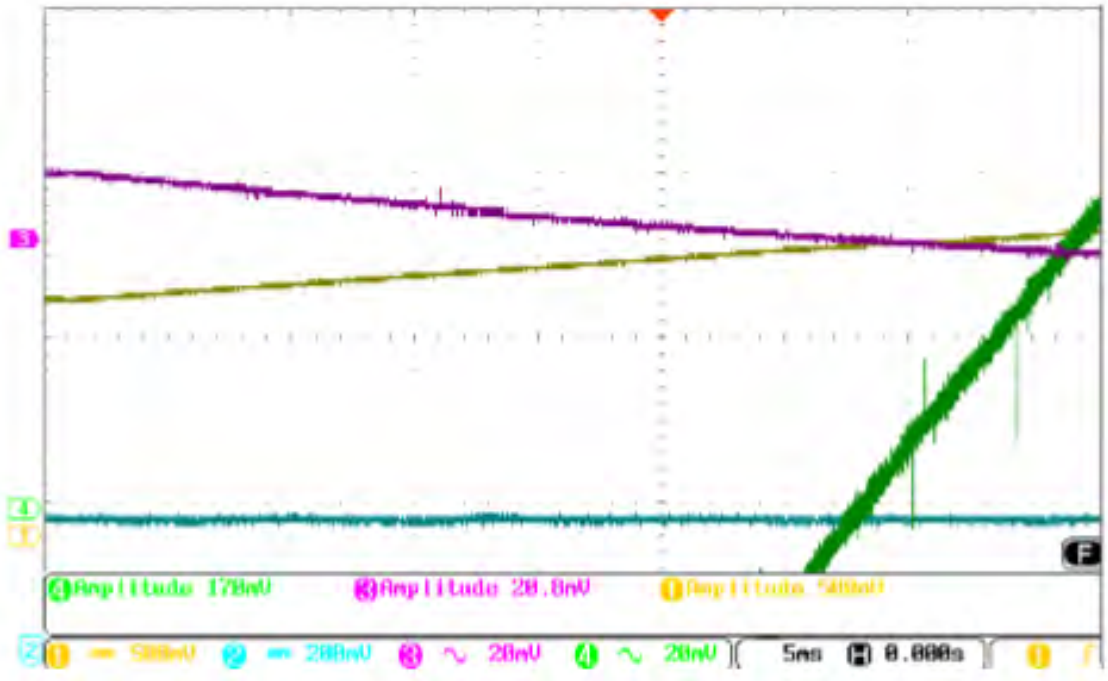


Figure 5.2: Noise from Scan Box with Filter: This is a snapshot from [5] of an oscilloscope screen showing the noise on the current (green) and piezo (purple) signals after a low pass RC filter was placed in the circuit.

- Solder output wires to connect to the scan box.

Due to the high voltages, exposed elements are taped over. Additionally, the power supply is also very sensitive to short circuits, so a fuse is connected in series with the on/off switch. Figure 5.3 shows the finished power supply, with most of the added wiring and some of the insulating tape visible, and is protected further by being encased in a sheet metal housing.

Once the power supply was attached, we began testing to determine its effectiveness at preventing ground loops. Using the same settings on the oscilloscope as B. Halkowski in [5], we generated Figure 5.4. The number of random spikes is virtually nonexistent, which provides strong evidence supporting the claim that adding the




Figure 5.3: The Floating Power Supply: A zoomed in view of Figure 4.3 featuring the floating power supply. The switch is soldered to the fuse and to the input power from the wall. The transformer and the terminals are hidden behind the control box on the left. In the center, the positive, negative, and ground outputs are visible. Tape is used both to insulate where the outputs meet the power supply and to prevent other exposed parts on the back of the power supply from coming into contact with other objects.

integrated floating power supply has greatly suppressed noise from ground loops.

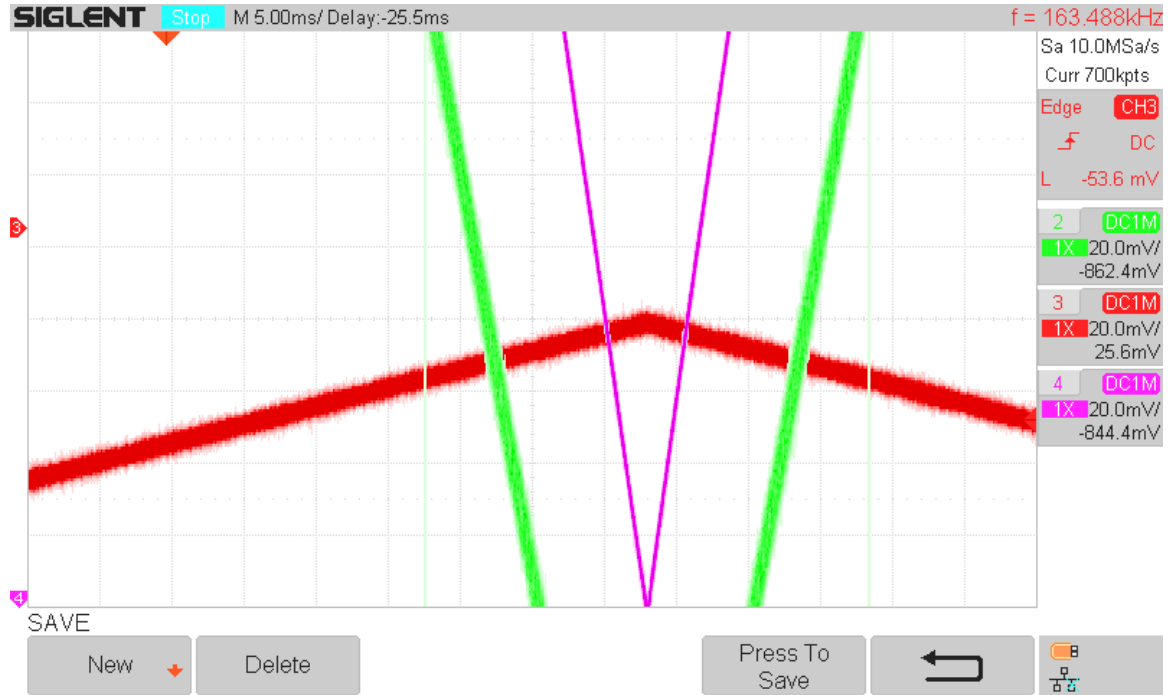


Figure 5.4: Noise from Scan Box with New Power Supply: This is a snapshot from an oscilloscope screen showing the signals from the piezo (red), current (green), and the function generator input (purple) without the RC filter and with the new power supply. The scale of the signals is comparable to the scale on the figures from [5]. The effectiveness of the floating power supply can be determined by comparing the noise on the current signal from this scan with the noise on the current signals on the scans from Figures 5.1 and 5.2. Evidently, the spikes are no longer present and the magnitude of the other noise of the peizo signal is similar to Figure 5.1, so it is likely that the spikes were caused by a ground loop.

Chapter 6

Results

In order to determine if progress had been made, we measured scan range without mode hops to compare to with previous measurements, such as the ones found in [5]. We also made attempts to measure the linewidth of the laser, which had not yet been done precisely, due to past thermal drift and mode hops of the laser.

6.1 Resonance Scans

To find the scan range, we used the feedforward scan box (more details of operation can be found in [5]) with a triangle wave input from a function generator. The steps taken are:

- Adjust piezo and current settings so the laser is on resonance.
- Turn on function generator.
- Adjust the size of the scan range by turning the ‘Master’ potentiometer on the front.
- Observe laser output on scope.

In order to maximize the range, the ratio between the scan slopes of the current and piezo can be changed by adjusting the ‘Piezo’ potentiometer on the front. When

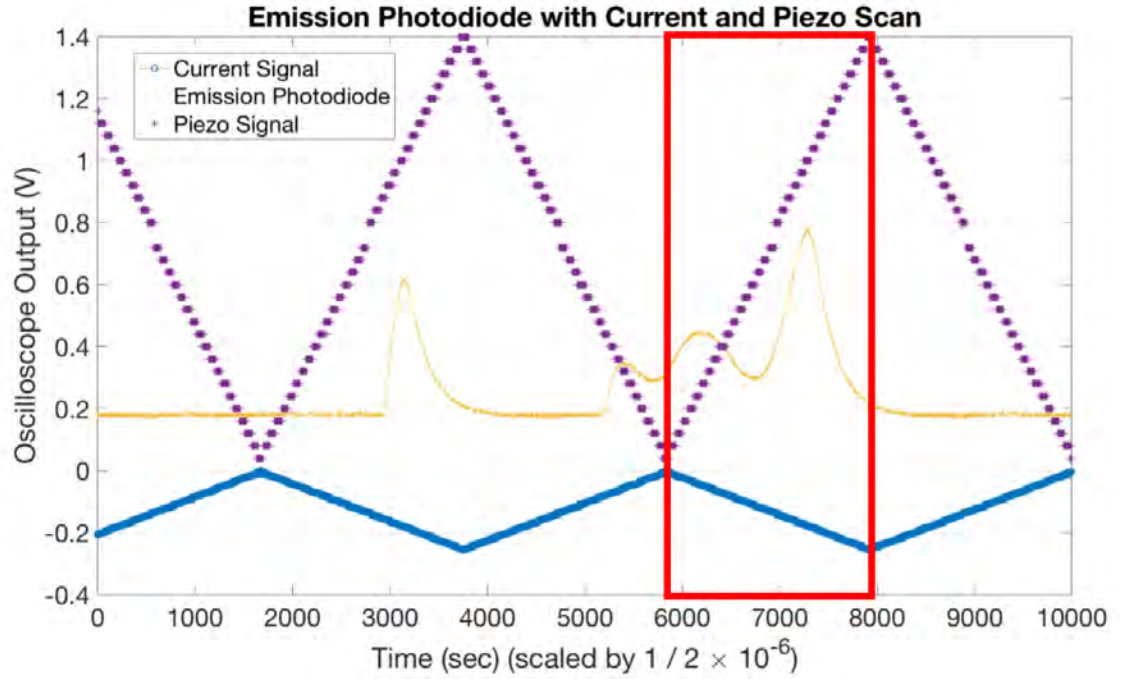


Figure 6.1: Scan Range from B. Halkowski's Thesis: Adapted from [5]. The red rectangle highlights the full range in one direction of the scan. In this case the laser was not consistently on resonance, and the scan image was taken the brief moment the signal was visible.

going through the steps, it is common to lose resonance, in which case we turn the generator off and begin from the first step again. The first notable signal we found is shown in Figure 6.2. While it is not a large range, it is confirmation that the process can show results. After more attempts, we found the signal shown in Figure 6.3. This scan shows at least a half of one peak; the range is smaller but similar to the example in Figure 6.1. The most promising result of these tests is the length of time the laser would maintain resonance during a scan. Where B. Halkowski had a split second chance to snapshot Figure 6.1, the resonance in Figure 6.3 was maintained for tens of seconds. This implies that the changes we made to improve stability had a positive effect on the laser's scanning capabilities.




Figure 6.2: Initial Small Scan Range: The red rectangle highlights the full range in one direction of the scan. The purple signal shows the intensity of the laser light; as the wavelength is scanned across the emission peak, the intensity changes. With this method we can observe the amount of the spectrum the laser can scan across. In this case, the covered area is only a portion of one peak, so it is hard to determine how large the scan is.

Source Signals:

Yellow: Temperature

Purple: Photodiode

Green: Piezo

Blue: Current

6.2 Linewidth Analysis

One of the main strengths of an ECDL is its narrow linewidth. As stated in the discussion of the planned applications in chapter 1, it is critical that our ECDL have a

sufficiently small linewidth. The general target is to have a sub-MHz linewidth, with an optimal case being less than 100 kHz and the threshold for success of approximately 0.5 MHz. Due to the instability of the laser in the past, it was difficult to accurately measure the linewidth, as it requires a constant output wavelength for tens of minutes of time at the minimum. Now, with the improved stability, we were able to record data and make initial estimates for the linewidth.

When two lasers at very similar frequencies are combined, the resulting combined laser has a unique spectrum of frequencies (beat note) created by the interference between the spectrum of output frequencies of the two incident lasers, the setup of which is detailed in Figure 6.4. As is evident from the beat note equation 6.1, if one of the lasers has a much smaller linewidth than the other, then the resulting beat note is almost exactly the optical spectrum of the laser with the larger linewidth. In this way, we can measure the linewidth of our laser using a control laser from the lab.

$$\omega_{beatnote} = | \omega_{laser1} - \omega_{laser2} | \quad (6.1)$$

For our control laser, we used a Toptica Photonics DL Pro, that we knew had a narrow linewidth (on the order of 100 kHz) and an output wavelength that could be tuned without worry of mode hopping. The first step was to combine the beams using a optical fiber combiner (ThorLabs PN780R5A2) specifically designed to combine two fiber inputs with output to two fibers. We sent one of the outputs to a rubidium vapor cell, and the other to an optical detector connected to a spectrum analyzer. Using the rubidium vapor cell, we found the voltages for the Toptica laser that corresponded with the peaks in the rubidium spectrum. Next, after we found resonance with our ECDL, and then we matched the two wavelengths by scanning the Toptica laser across the four resonance peaks (shown in Figure 1.7) and watching the spectrum analyzer for a signal. Once we had the signal, we narrowed the scan range until we could

make a decent estimate of the linewidth. There was some trouble at this point, as the signal was subject to significant frequency jitters, so we took several screenshots (Figure 6.5) and found the mean full width half max, or FWHM, linewidth to be approximately 0.5 MHz.

This is a very promising result, as the linewidth of our laser is likely close enough to the linewidth of the Toptica laser that the small linewidth approximation no longer holds. In this case, the beat note on the scope would actually be larger than the linewidth of each of the incoming lasers, so this measured linewidth is likely an upper bound for the real linewidth. Given our initial conditions, this is considered a success.



Figure 6.3: Initial Larger Scan Range: Taken on the same day as Figure 6.2. The red rectangle highlights the full scan range. Here, the range is clearly a half of one peak and the start of another, which is almost as large as the max range found in Figure 6.1. Unlike in 6.2, the scan is undergoing a one direction mode hop. There are a few variables that can cause this, and we did not take the time to isolate what the primary cause is.

Yellow: Temperature

Purple: Photodiode

Green: Piezo

Blue: Current

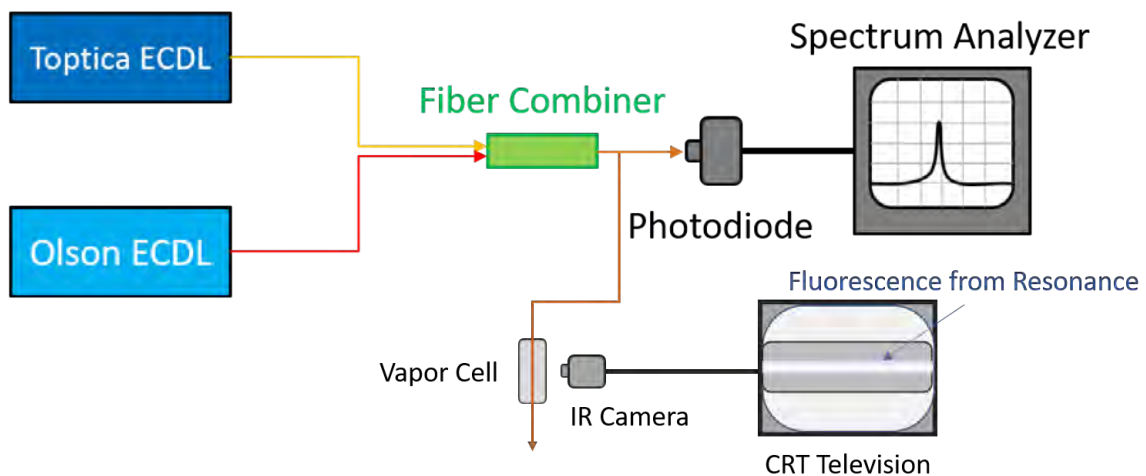


Figure 6.4: Linewidth Analysis Setup: A diagram of the setup to use beat note analysis to determine linewidth. In general, it is difficult to align two lasers to the same wavelength, which is a necessity when measuring the beatnote. Using the resonance peaks of rubidium and observing fluorescence on the CCTV, we can get the lasers very close to the same wavelength without using a wavemeter. This process can take from 10 to 30 minutes depending on the stability of the lasers. Previously the laser would mode hop much too frequently for analysis to be done, but the improved stability has made completing this process possible.




Figure 6.5: Linewidth Results: By analyzing the frequency spectrum of the beatnote created by the control laser (with a narrow linewidth) and our laser (with an unknown linewidth) we are able to estimate our laser's linewidth. To do this, we used the number of boxes as a measurement tool to find the width of the signal at half of the maximum height and convert that into a linewidth.

Chapter 7

Outlook and Conclusion

7.1 Outlook

In the future, attempts to further improve the laser should be through environmental insulation. The strongest improvements in the performance of the laser have consistently been through improvements in thermal insulation. The largest thermal improvement came through eliminating air currents by patching the holes in the existing housing. The most straightforward way to further optimize this would be to create a more permanent housing. One idea that we currently have is to create something with an inner and outer shell. This way, there would be more insulation from the long term temperature changes in the room and it would provide more insulation from air currents.

A further improvement involves isolating the laser base from acoustic vibrations. While measuring the linewidth, we found that walking near the table and tapping on the table significantly increased the jitters for the duration of the noise. We concluded that with the other variables managed to a good degree, the issue of vibrations in the environment transferring to the ECDL setup is now an obstacle to optimal laser performance. One theory we have for the cause of the transfer of vibrations is the current housing may be in contact with the table and with the base of the

setup. Currently, the setup is isolated from the table with sorbithane rubber, but the housing may potentially be bridging the environment and the setup and transferring vibrations. The solution to this is, again, to create a better housing for the laser. In this case having two shells, with one in contact with the table and the other not in contact, could provide a simple solution to the problem.

7.2 Conclusion

In this thesis, we have outlined the process of further refining and tuning an external cavity diode laser in order to use it in cold and ultracold experiments in the Aubin research lab. To do this, we assessed problems preventing optimal operation of the laser, and applied solutions to them from the perspectives of thermal, mechanical, and electrical engineering. Increased thermal stability resulted in longer periods of time on resonance, which is necessary for long term data taking and further assessment of the lasers capabilities. Improvements in the physical layout of the control boxes and changing the scan box to be powered by a floating power supply do not seem to have negatively affected the scan range of the laser.

The final tests of the effectiveness of the laser were done more to get a benchmark for the current state than to fully explore the capabilities. The results of these tests seem promising enough that we can confidently say that the modifications to the laser and the control system have positively impacted the performance of the laser overall and provide a strong foundation for further improvements to the laser to make it ready for use in the lab.

Bibliography

- [1] Francesca Fornasini, The Application of Four-Wave Mixing to Cold and Ultra-Cold Atom Imaging, Honors Thesis (College of William and Mary) (2010).
- [2] Brian Chase, Developing a Dipole Trap Laser for the Purpose of Cold and Ultra-Cold Molecule Production, Senior Thesis (College of William and Mary) (2012).
- [3] N. Vujicic, T. Ban, G Kregar, D. Aumiler and G. Pichler, Velocity-selective double resonance in Doppler-broadened rubidium vapor, Physical Review A. 87, (2013).
- [4] Toni Laurila, Timo Joutsenoja, Rolf Hernberg, and Markku Kuittinen, “Tunable external-cavity diode laser at 650 nm based on a transmission diffraction grating,” Appl. Opt. 41, 5632-5637 (2002).
- [5] Ben Halkowski, External Cavity Diode Laser for Ultra-cold Atom Experiments, Honors Thesis (College of William and Mary) (2018).