Designing a Stabilized Laser System for Laser Cooling

Alexander Clement

Senior Thesis in Physics

Submitted in partial fulfillment of
the requirements for the degree of
Bachelor of Arts

Department of Physics
Middlebury College
Middlebury, Vermont

May 2012

Abstract

The technique of laser cooling radically reduces the temperature of atoms to within one degree of absolute zero. To laser cool atoms the frequency of the cooling laser must be precisely controlled. The temperature, current, and cavity length of a laser all affect its output frequency. Towards this end PID feedback loops are used to hold these parameters constant. The laser frequency can be stabilized near the $F = 3 \rightarrow F' = 4$ transition in ^{85}Rb used for laser cooling. The final, fine frequency adjustment is accomplished using acousto-optic modulators. The methods of frequency control are used on both the "master" laser, which provides the frequency standard for the system, and the "slave" lasers, which provide the extra power needed to cool the atoms.

Committee: ______________________________

Anne Goodsell

____

Stephen Ratcliff

____

Daniel Scharstein

Date Accepted: ________________________

Acknowledgments

I would like to thank the entire physics faculty both as a group and as individuals. You foster a community that is as socially vibrant as it is academically. Stephen Ratcliff, your humor and incredible competence welcomed me into the department. Jeff Dunham, you challenged me to answer every question, especially those that I initially dismissed as trivial or impossible. Susan Watson, you helped me clarify my notation and my thinking. Anne Goodsell, I hope I have helped you enjoy Middlebury as much as I have. Your tireless assistance and friendship made this thesis more valuable than any other experience I have had in my college career.

To my fellow physics majors and tenets of 525, thank you for the company during the long nights and short tea-times. Thanks to my all other friends for teaching me everything I did not have time to learn in the classroom. John, Lani, Amy, thank you for all your support these past four years. If I have ever bored you with too much detail, you have never shown it.

Contents

List of Figures

Chapter 1

Introduction

The goal of laser cooling is to radically reduce the temperature of a gas of atoms to within one degree of absolute zero through light-matter interactions. These interactions only occur when the incident photon has an energy which corresponds to an atomic transition within the atom. To ensure these interactions, we must measure, control and lock our laser to the frequency which corresponds to the atomic transition. This thesis details the methods for the control of the emission frequency of lasers as well as methods to shift the frequency of laser light in the beam path.

Much progress has been made towards characterizing the frequency of our lasers [1]. The frequency of the laser diode is dependent on its temperature, injection current, and cavity length. We use electronic feedback to stabilize all of these parameters and hold the frequency constant. Using Doppler-free spectroscopy, we can measure the frequency of the laser relative to the atomic transitions of ^{85}Rb. We can then use the feedback circuity to lock the laser to any one of these transitions. Both commercial and lab-built stabilization circuits are used and examined in this thesis.

Fine control over the frequency of individual beams produced by the same laser is also a necessary component of the laser cooling apparatus. This thesis makes a study of acousto-optic modulators, devices which shift the frequency of light through interactions with sound. These devices can shift the frequency of the laser to account for the different Doppler shifts associated with counter propagating beams.

These methods are used to stabilize and set the frequency of our "master" laser which is the frequency standard for our entire experiment. Since additional laser power is required to create the optical trap, the master is used to drive "slave" lasers by stimulating the slaves with light from the master. Together, these lasers provide us with enough light of the proper frequency to cool atoms.

Chapter 2

Principles of PID Circuits

2.1 What is a PID Circuit?

The Proportional-Integral-Derivative (PID) circuit is a electronic device that is designed to hold a voltage at a constant value (called the setpoint). This circuit relies on measurements of this variable voltage and responds to the difference between the measured value and its internal setpoint, the error signal. By constantly adjusting its output the PID circuit uses feedback to minimize the error, holding the variable voltage at or near the setpoint. The PID circuit is designed to rapidly correct for changes in the error signal without overshooting the setpoint by too large a margin. When the output is at the setpoint the circuit is considered "locked," this means that small perturbations will not affect the output past a certain tolerance. This is achieved by the use of three different feedback loops within the device that all respond differently to the error signal and then sum their respective outputs. These three loops are the proportional, integral, and derivative loops and they respond to both the magnitude of the error signal as well as its time dependence.

In analog circuits this feedback behavior is the hallmark of operational amplifiers

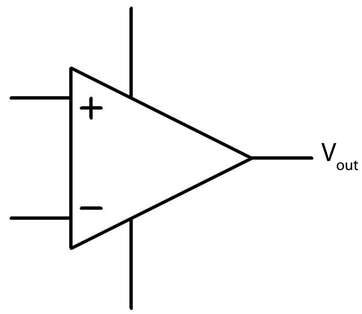

Figure 2.1: This diagram shows the pin-out of a typical op-amp. The inverting and non-inverting inputs are labeled - and +, respectively. The two vertical inputs provide power to the output and determine is range.

(op-amps). These integrated circuits as shown in Fig. 2.1 have inverting (−) and a non-inverting (+) inputs that determine the output of the op-amp according to the simple formula

$$V_{out} = (V_+ - V_-)A \tag{2.1}$$

where A is the open-loop gain of the op-amp, generally on the order of at least 1000. Without any feedback (a connection from the output to the inverting input) the op-amp is almost always "railed." Due to the size of A, any slight difference between the two inputs causes the op-amp to output either the largest positive or negative voltage that it can, limited by its power supply. The op-amp can be used to follow or amplify a signal by employing negative feedback, where there is a connection between the output and the inverting input of the op-amp. The circuit in Fig. 2.2 is an inverting amplifier where the output is given by

$$V_{out} = -\frac{R_2}{R_1} V_{in} \tag{2.2}$$

so that the output is directly proportional to the input. Other negative feedbacks are used to achieve both integral and derivative feedbacks, but this proportional feedback is exactly the sort used in our PID loops.

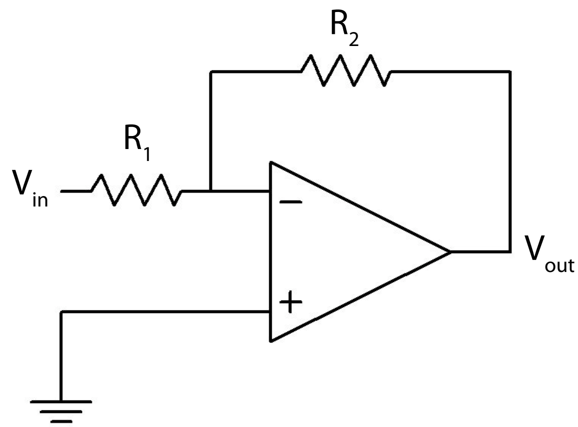

Figure 2.2: An example of a op-amp being used in a negative feedback application, specifically an inverting amplifier.

2.2 Proportional Feedback

Proportional feedback is required output the form

$$V_{out} = K_P V_{err}, \tag{2.3}$$

where K_P is a gain constant that be adjusted to tweak the behavior of the PID loop. This feedback is dominant for small, rapid fluctuations around the setpoint. The proportional gain responds to these small variations by producing a signal of the opposite sign that is proportional to the error. This correction becomes zero at the setpoint so the proportional part of the circuit turns off. Thus the proportional loop is incapable of holding the output constant, but rather corrects for changes that occur. This results in a output that oscillates slightly around the setpoint. Since the proportional feedback loop is the most responsive in this region, the magnitude of the oscillations is directly related to the proportional gain. If the proportional response is very large, either because of a large deviation from the setpoint, a high gain, or a combination of the two, then the PID loop can become unstable and drift too far away from the setpoint to be able to recover. At this point the feedback loop is “unlocked” and the output voltage is unregulated. For very large errors a different type of feedback is needed to bring the circuit back into the realm of proprotional control. This is often done with integral response.

2.3 Integral Feedback

Integral feedback is based on the accumulation of error over time. This can be simply realized as the area under the error curve (where the setpoint is the zero) or expressed as

$$V_{out} = -\frac{K_I}{\tau}\int V_{err}(t)dt \tag{2.4}$$

where K_I is the gain of the integral feedback loop and τ is its time constant. While proportional feedback can struggle to correct for large deviations from the setpoint, the output of the integral feedback grows rapidly with time. This summation of error allows the integral feedback to rapidly bring the system to zero, however since this feedback accounts for the history as well as the present, it is not necessarily zero at the setpoint. In fact, the accumulation of error causes the output to overshoot the setpoint leading to an output that behaves like a damped oscillator. Each correction overshoots the zero by a smaller amount until the proportional loop becomes dominant. Obviously, it would be preferable to minimize this oscillation and overcompensation, which is where the derivative loop comes in.

2.4 Derivative Feedback

Derivative feedback is designed to damp variations in the output, smoothing out large features. As the name implies, this feedback responds to changes in the error signal, rather than its present value, in effect trying to predict what it will become. This can be expressed mathematically as

$$V_{out} = -K_D \tau \frac{dV_{err}(t)}{dt} \tag{2.5}$$

where, again, K_D is the gain of the loop and τ is time over which the slope is measured. As one can easily determine from this relation, this part of the circuit moderates rapid change. This can be very useful once the output has been locked to the setpoint, but can actually inhibit the performance of the other loops. In many applications it is better to achieve a rapid lock with several over-corrections than it is to delay the locking process.

2.5 Putting the PID together

Each section of the PID loop behaves in a way that complements the whole. The proportional feedback addresses small rapid changes to the output. The integral feedback corrects for large changes in the output and helps to rapidly bring it towards the setpoint. Finally the derivative feedback smooths out the over-corrections of the other two feedback and any large jumps that the output takes (e.g. voltage spikes). This produces a very robust control loop shown in Fig. 2.3 that can be applied in many different settings.

However there are limitations on the performance of PID loops. As discussed in the sections on proportional and integral feedback a PID loop cannot hold the output at the setpoint, rather it just corrects for any deviations away from this value. This means that the PID actually causes the output to oscillate around the setpoint. The integral feedback is the primary cause behavior, as it has a "memory" and does not output zero when the output is at the setpoint voltage. Thus, the size of this oscillation is controlled by the gain for both the proportional and integral loops with

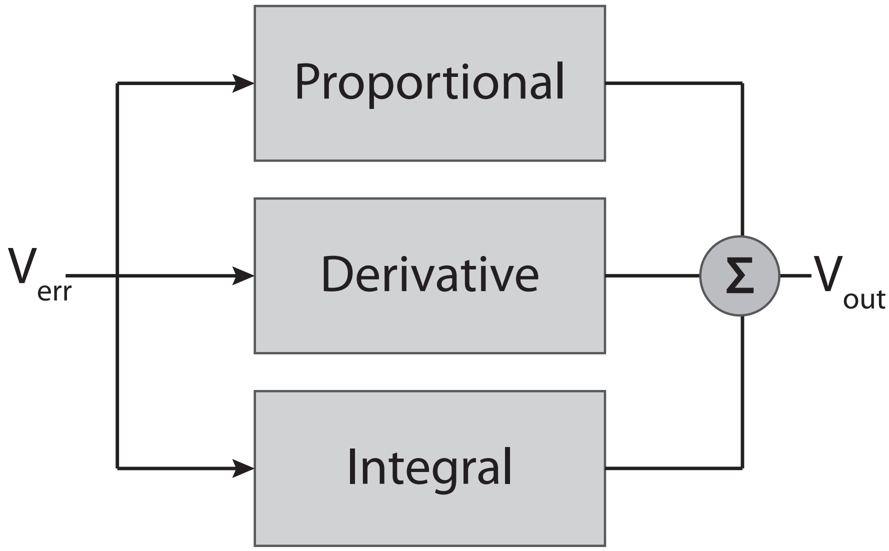

Figure 2.3: A schematic of a PID circuit.

the proportional loop working to minimize the overshoot. With proper gain settings these oscillations can be damped to the point where they no longer interfere with the operation of the device.

Unfortunately, PID loops cannot lock to any arbitrary point on a curve; they are only stable if the setpoint is on a monotonic slope in the device's output. This is most easily explained in the context of the PID's proportional behavior, but applies equally to all three pieces. If the output changes monotonically, then a drop below the setpoint will produce a negative error signal, while a rise will produce a positive one. This means that a change in either direction is opposed by the circuit. However, if we imagine that the setpoint is the peak of a curve, then either a drop or a rise would produce a negative correction. This correction would work on one side of the peak, but would be in the wrong direction on the other side. The circuit has no way of telling the difference between the two changes. Even if they are locked to a monotonic slope, too large a change to the output can shift the realm in which the external system is operating. An example of this potential shift is explained in the next chapter. We address these limitations in both our laser stabilization and temperature control systems, the two major applications of PID circuits to this project.

Chapter 3

PID Circuits in our Laser Cooling Apparatus

3.1 Applications of PID Circuits in Diode Laser Stabilization

We use PID loops in many applications related to stabilizing the frequency of our diode lasers. The frequency of a diode laser is dependent on its temperature, the injection current, and the length of the external cavity. We can use PID loops to control and stabilize all of the parameters, as we do do on our master laser. Each of these parameters has a different effect on the laser frequency. The laser diode expands and contracts with temperature, changing its cavity length and shifting the wavelength of light that is emitted by about 0.11 nm per °C [2]. The injection current changes the frequency of the laser in two ways. Changes in current will cause the laser diode to be warmed or cooled due to Joule heating, thus if the current is to be changed, temperature control must be used. Second, the injection current is used to force the laser into single-mode operation. At lower currents the laser rapidly switches between several different frequencies. In single-mode operation, the laser diode has a single, narrow frequency peak. Once the laser in operating in a single-mode changes to the current can switch the dominant mode of the laser, radically shifting the frequency of the laser. In a Litman configuration [1] as shown in Fig. 3.1 the length of the external cavity is controlled by a diffraction grating mounted on a piezo-electric transducer. This grating can is displaced by the application of a voltage

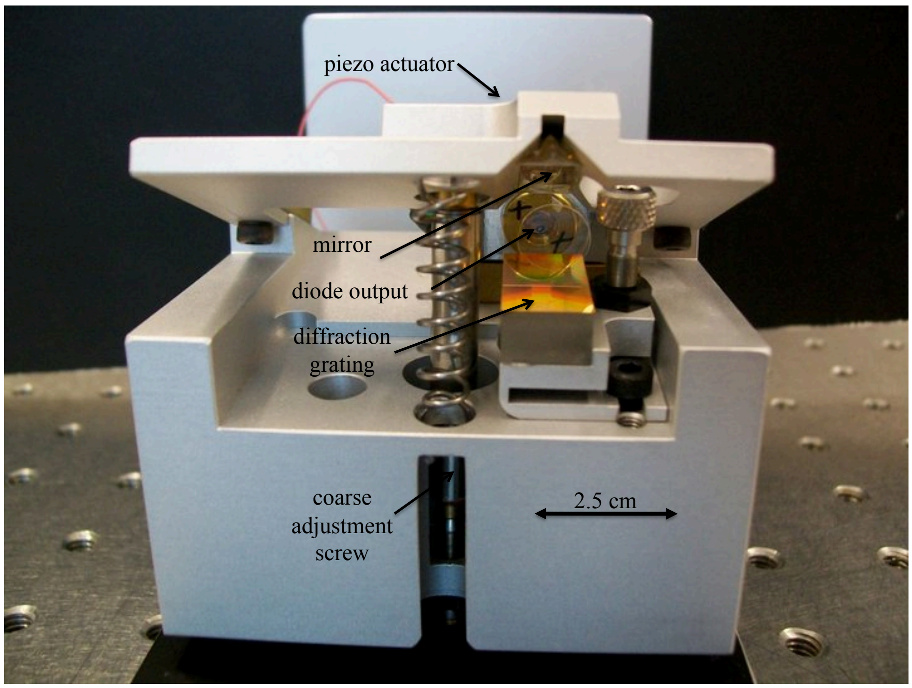

Figure 3.1: A laser in a Litman configuration. The diffraction grating is rotated by applying a voltage to the piezo.

to the piezo. This gives very precise and rapid frequency control and is our method of choice for fine-tuning the laser frequency. To ensure that the laser frequency is only dependent on the piezo voltage, PID loops are used to hold the laser current and temperature at constant values.

3.2 Temperature Stabilization

All our lasers are temperature regulated using thermo-electric coolers (TEC), see Fig. 3.2 for an example. These devices use an electric current to create a temperature

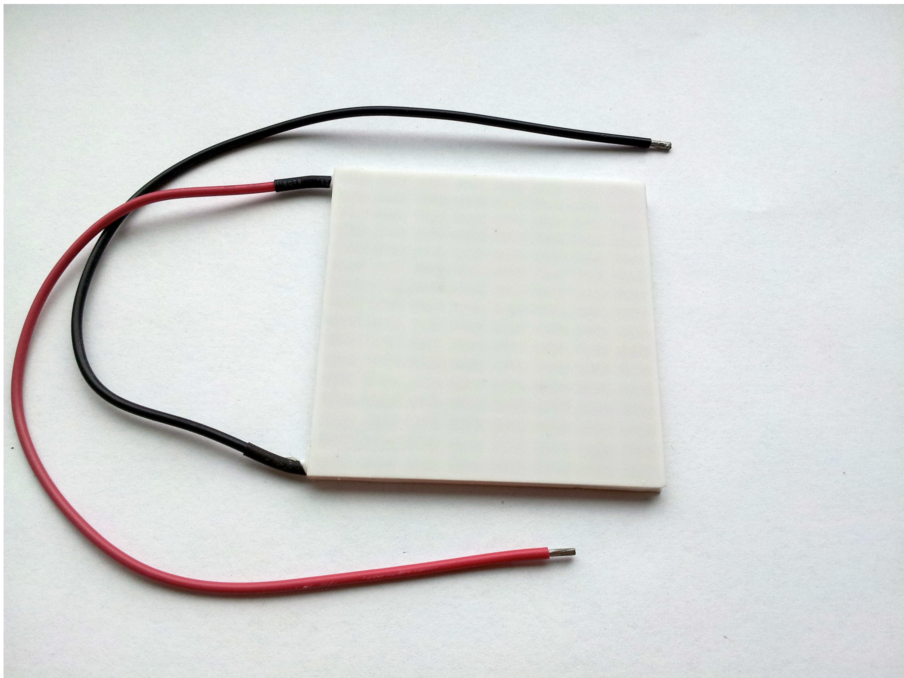

Figure 3.2: A 2x2" themo-electric cooler (TEC). A current can be applied across the leads to cause heat to flow from one plate to the other.

difference between their two faces. Conversely, an externally applied temperature difference will drive a current in the TEC. The cooling system that we I have been assembling for the slave lasers relies on three TEC units. One small (less powerful) one is as close as possible to immediate thermal contact with the laser diode and ensures that it is held at the correct temperature. Two other TECs are used to bring the overall temperature of the laser enclosure much closer to the desired temperature, reducing the strain on the small TEC. The heat produced by the two large TECs is removed by water blocks which lie under the entire unit. This apparatus can be seen in Fig. 3.3, where the platform and aluminum enclosure provide extra thermal mass.

The current supply to these TECs is regulated by a PID loop which uses a thermistor to determine the temperature of the laser diode. This makes a perfect signal for the PID to lock to, as the temperature typically changes smoothly and monotonically. The single small TEC is driven by a different circuit than the two large TECs as they have different current requirements; however the circuits are identical in their

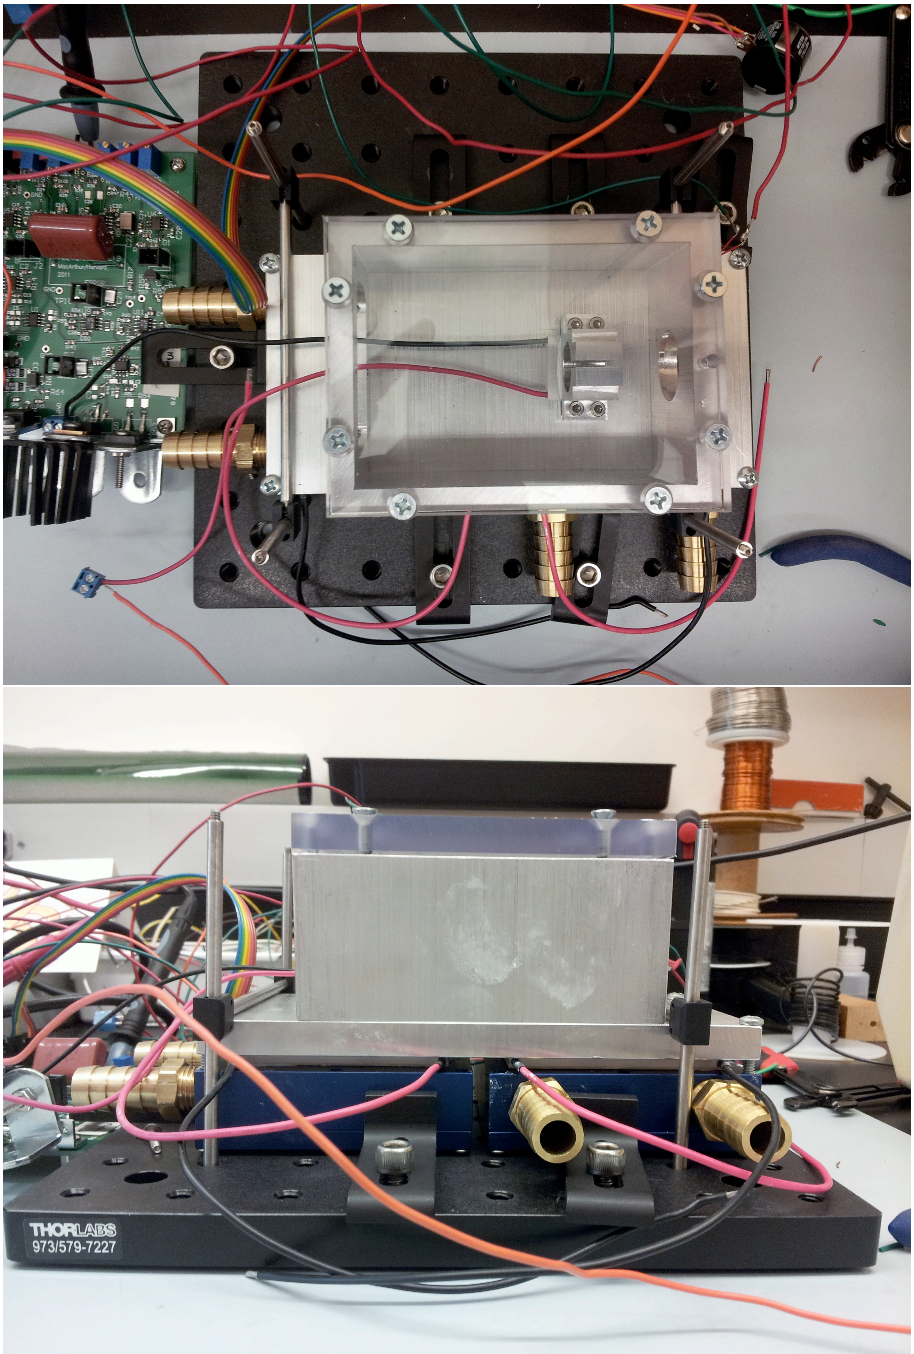

Figure 3.3: The slave laser housing from above (top) and the side (bottom). The blue water plates with their brass connections can be clearly seen, with the two large TECs mounted on top. From above, the red and black leads of the small TEC can be seen mounted just below the laser bracket.

operation. These feedback loops compare the voltage across the thermistor to their external setpoint, and this signal is fed into proportional and integral gain loops. We have chosen not to include a derivative loop on these boards because temperature changes are very slow relative to the response of the circuit and we allow the circuit to stabilize the temperature of the laser diode before we turn it on, so the initial correction can cause large oscillations without causing damage. The corrections due to the proportional and the integral loops are summed and then amplified according to an overall gain.

We have made efforts to characterize the behavior of the PID circuits during different modes of operation. The PID circuit was tested with just the small TEC operational. Its performance was measured with just proportional response enabled, with both proportional and integral response, and finally, as a baseline, with no active response at all. These tests were conducted by allowing the temperature of the diode housing to reach a constant value, then heating the housing. The three trials can be seen in Figures 3.4 and 3.5.

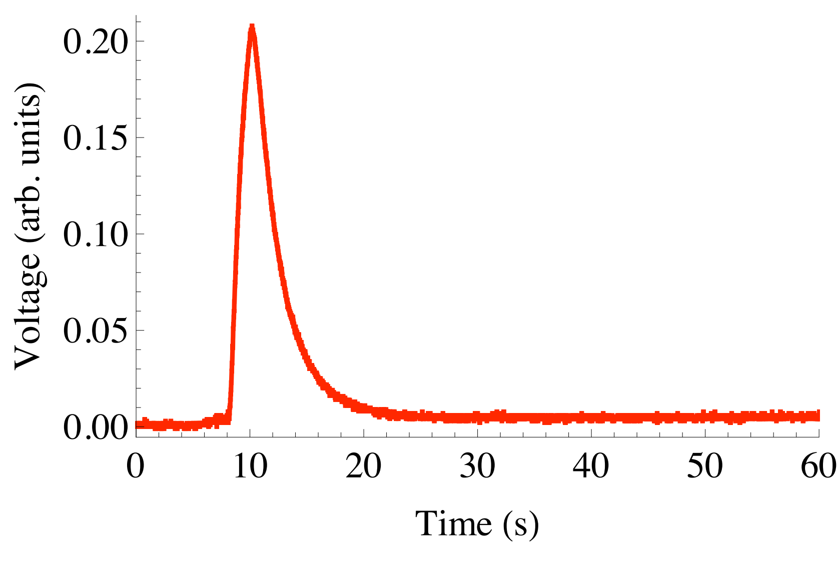

Figure 3.4: The error signal is shown after temporary heating with the PID disabled. Note the slow return towards zero without ever reaching zero.

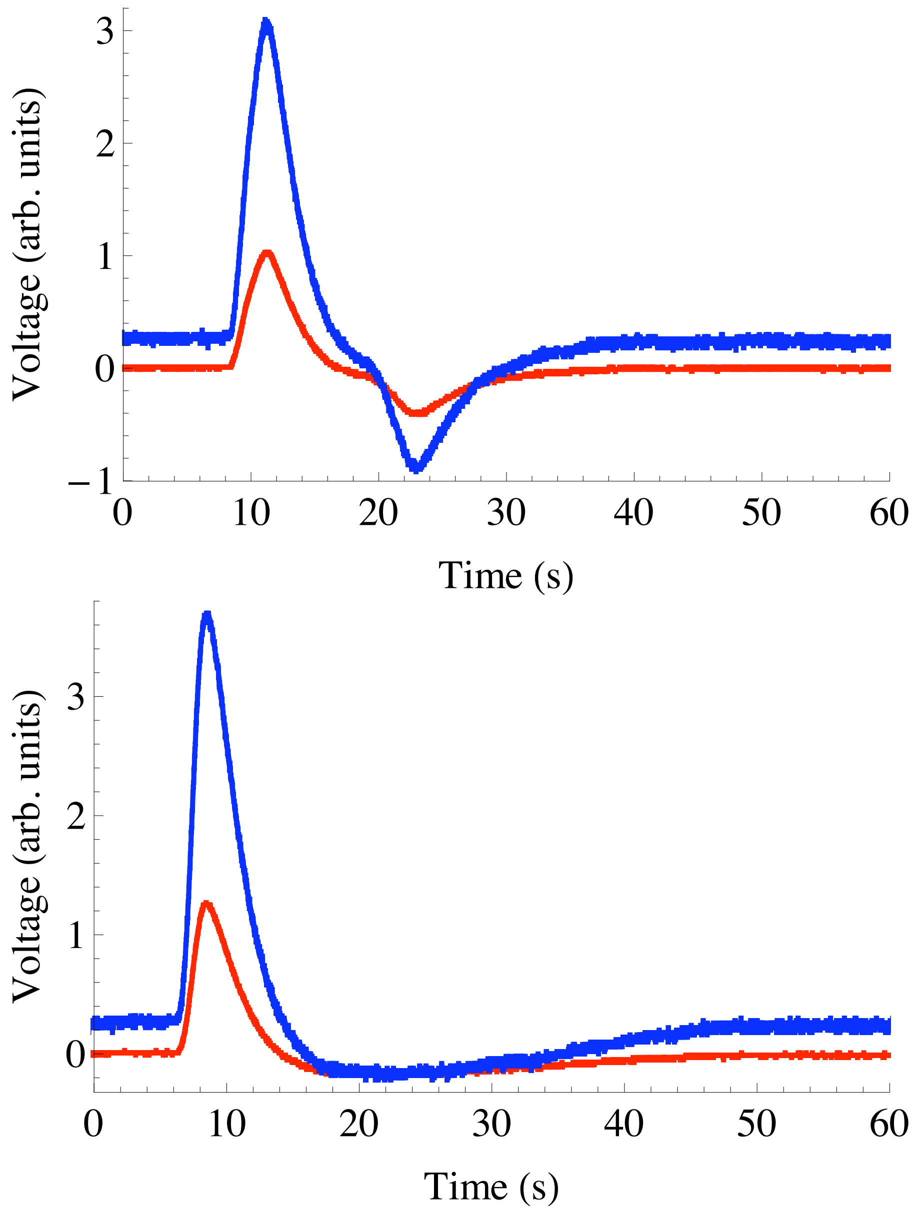

Figure 3.5: The red traces are the error signal, while the blue trace is the output of the PID. (Top) The response of the PID circuit with only the proportional loop enabled. We can see how the proportional response can over correct for small features in the error and cause oscillations. (Bottom) The response of the PID circuit with both the proportional and integral loops enabled. We can see how the integral response forces the output to be negative when the error is zero. The output slowly comes back to zero as the integral loop rids itself of accumulated error.

Despite the fact the the laser's frequency depends on the temperature, we do not adjust the frequency with temperature changes. The reason for this choice is clearly show in the plots above; the thermal mass of the laser housing is too large for the TECs to make rapid adjustments. We can see that it takes around 20 seconds to return the temperature to the setpoint after a large disturbance. Thus, all our lasers are held a constant temperature and the frequency is adjusted using other methods. For the master laser this frequency stabilization is accomplished with another PID device, the Sacher Lasertechnik LB2001.

3.3 The Sacher Lasertechnik LB2001

The Sacher Lasertechnik LB2001 High Frequency Servo Controller is used to lock the frequency of our laser. In this application the error signal is produced from a photodiode which measures the probe beam of a Doppler-free rubidium spectroscopy apparatus [1] [3]. The signal from the photodiode is passed into a lock-in amplifier. The lock-in produces a much larger signal that is easier to work with and allows us to offset the error signal. This offset is critical since this signal serves as the error signal for the LB2001, so it must be zero at our desired frequency. Figure 3.7 is a schematic of these devices and the feedback between them. This signal, shown in Fig. 3.6, is the first we have discussed that is not just a simple monotonic slope. We can tune

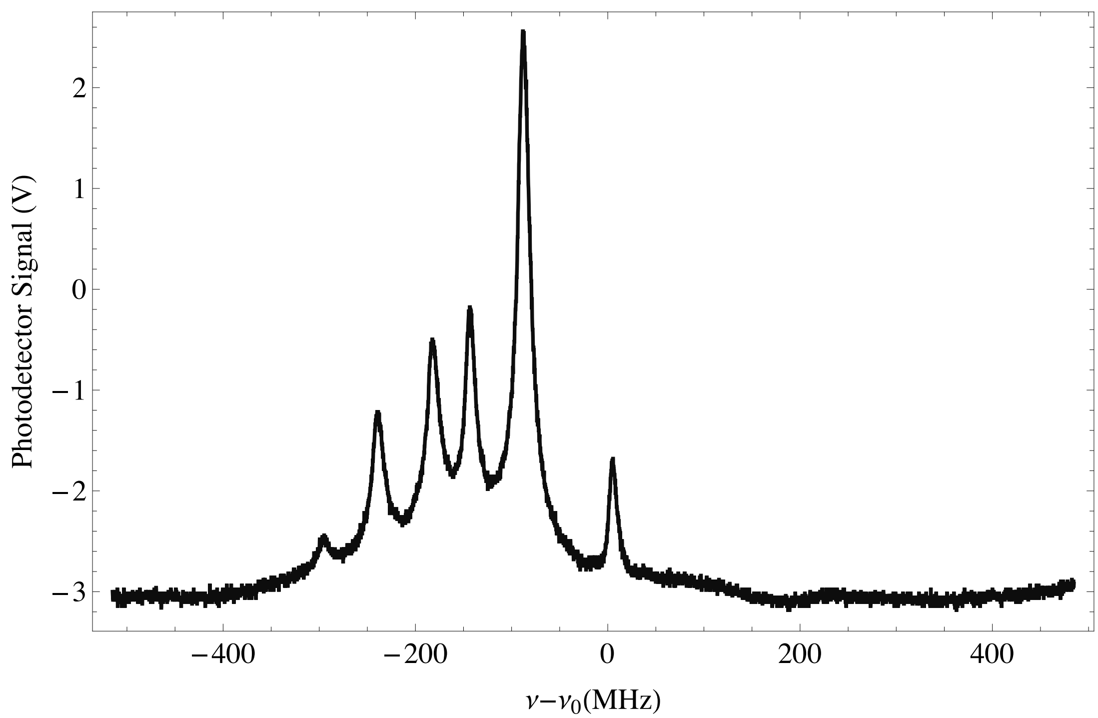

Figure 3.6: This plot shows how the error signal changes over a wide range of frequencies.

the laser over a wide range of frequencies, which allows to identify the proper place to lock the laser. However, there are multiple points where $V_{err} = 0$ and the LB2001 can lock to any of these. Thus, to get the LB2001 to lock to the desired frequency the PID loop must be activated when the laser is very close to that specific zero of the error signal. Not only does this type of signal make locking more difficult, but it makes the lock less stable: a sufficiently large disturbance would shift the error signal to a different region and cause the LB2001 to lock to the wrong frequency.

Despite these challenges the LB2001 can be used to lock the laser by adjusting the voltage to the piezo which rotates the diffraction grating, controlling the length of the external cavity. This voltage, and the piezo attached to it, respond to changes on very short time scales. Locking the frequency of this laser is almost instantaneous, once all the other parameters are stable as is detailed in Chapter 6. The speed at which changes occur requires that painstaking care be taken with the gain settings for feedback loops. Once the laser is locked, it is resilient to most external vibration and noise. Static shocks and obstruction of the spectroscopy beam path will both cause the laser to lose its locking point. Unfortunately, due to the nature of the error signal, the laser must be locked again by hand.

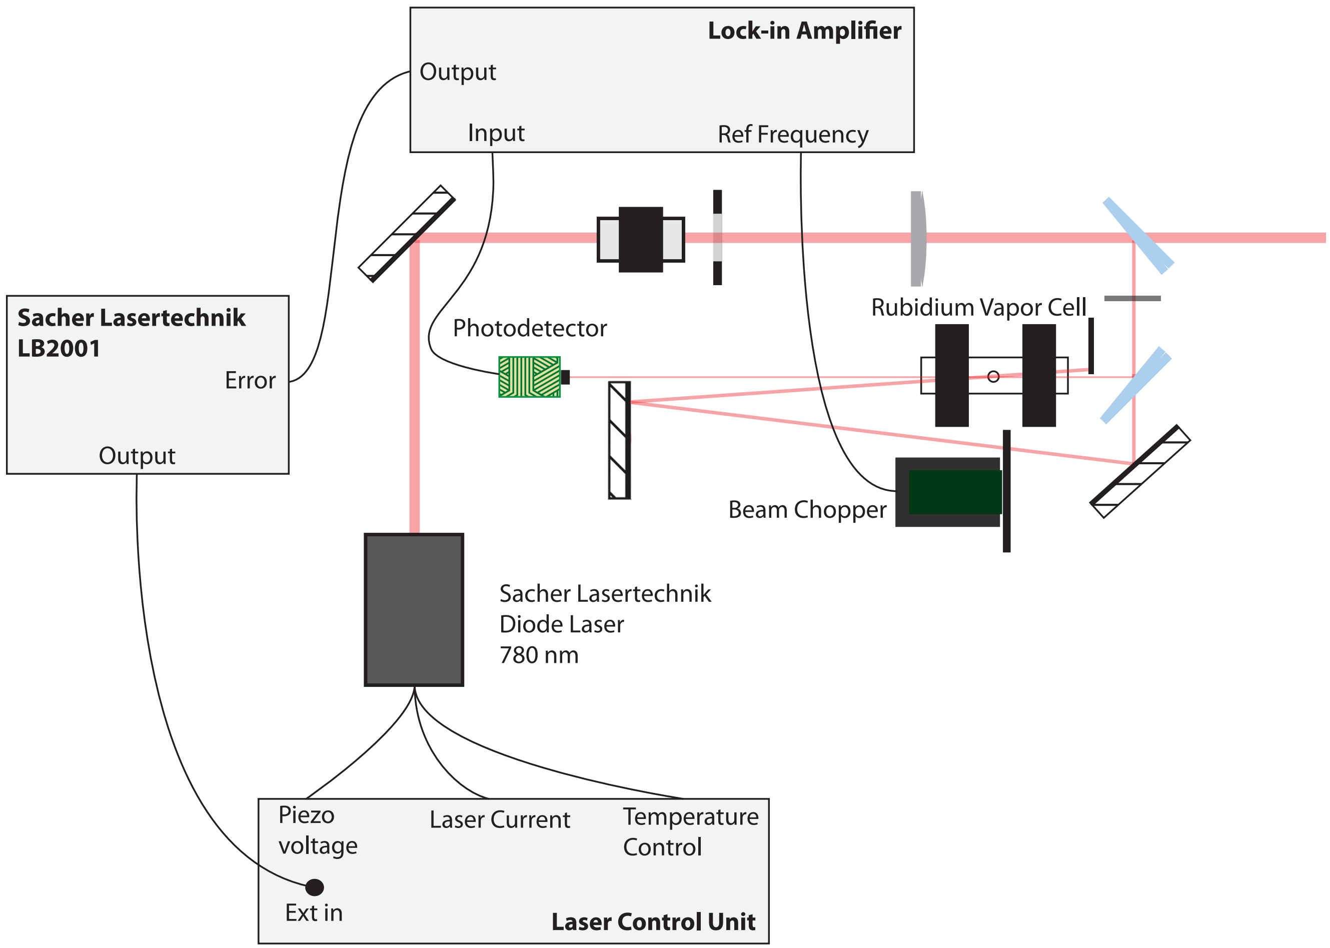

Figure 3.7: This diagram shows how the LB2001 monitors the frequency of the laser and adjusts the laser piezo voltage to compensate for any drift. The frequency of the laser is measured by Doppler-free spectroscopy. This signal is amplified in the lock-in amplifier. The LB2001 then adjust the piezo voltage to compensate for changes in the laser frequency and monitors the effect that they have, completing the feedback loop.

Chapter 4

Acousto-optic Modulators

4.1 Overview

An acousto-optic modulator (AOM) is an optical tool that exploits phonon/photon interactions to alter the frequency, direction and intensity of incident laser beams. Phonons that travel through a crystalline medium create regions of higher and lower density, yielding a grating. Photons that are incident upon this phonon-induced grating are diffracted. While in the medium, photons interact with phonons and experience a frequency shift. For first order diffraction, this interaction can be written as

$$\omega = \omega_0 \pm \Omega \tag{4.1}$$

where Ω is the frequency of the phonon, ω_0 is the frequency of the photon before the shift, and ω is the frequency of the photon after the shift. To account for higher orders of diffraction we introduce m as the diffraction order and rewrite Eq. 4.1 as

$$\omega = \omega_0 \pm m\Omega \tag{4.2}$$

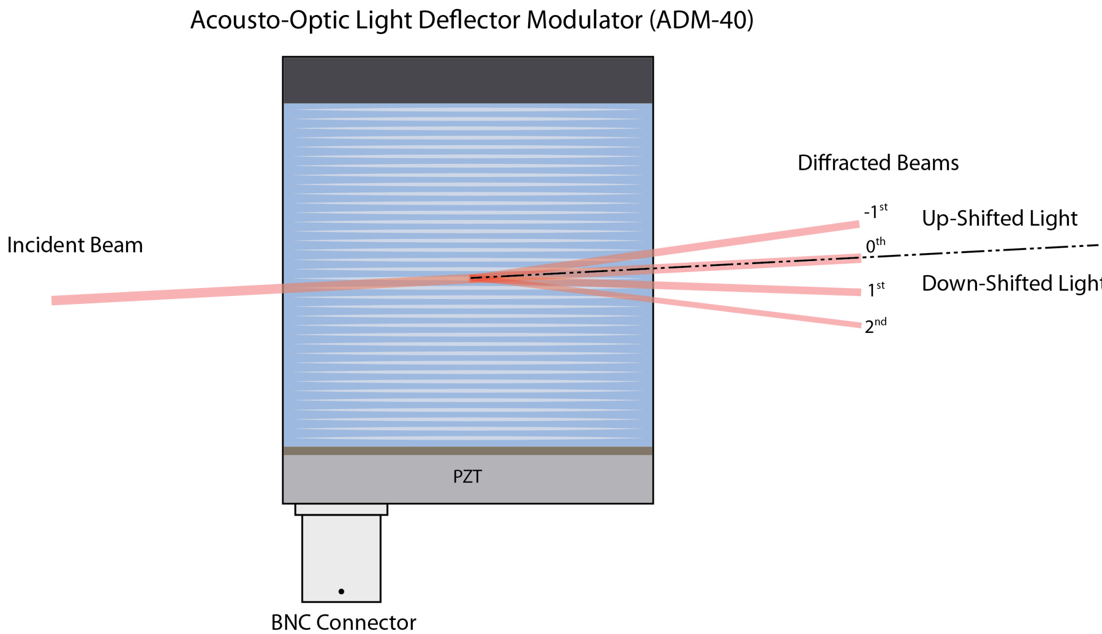

Figure 4.1: This figure shows a schematic of an AOM with the incident beam and the multiple diffracted orders. Up-shift and down-shift refer to whether the frequency of the light is increased or decreased.

. The phonons have much longer wavelengths than the incident light, so the shift in the photon's frequency is small. [4] This frequency shifted light is separated into the different diffraction orders by the AOM's sonic diffraction grating. A single laser beam entering an AOM is diffracted into many different beams with unique frequencies. The beam with the appropriate frequency can be selected by the experimenter.

AOMs for visible light are built by attaching a piezo-electric transducer (PZT) to a transparent crystalline medium. A radio-frequency (RF) driver is connected to the PZT and causes it to oscillate at a given frequency, producing sound waves in the crystal. This driver frequency is the phonon frequency Ω that shifts the photon frequency. With changes to this frequency, the AOM can be tuned to mimic the properties of many different conventional gratings, which rely on their physical structure to diffract a specific wavelength of light. Like a conventional diffraction grating, the incident laser's power is split among the many diffracted orders; however we can subtly rotate the AOM to change the ratio of power between the orders. The AOMs that we use are designed to be driven at frequencies between 30 and 50 MHz. In this range they can be configured so that 85% of the incident power is output to either of the first diffraction orders. This allows us to shift the frequency of our laser beam without losing a large portion of the laser power.

4.2 Use in Laser Cooling

To laser cool atoms we must be able to produce laser light of exactly the right frequency and hold the laser stable at that frequency. We can stabilize the frequency of our laser using the techniques described in Chapter 3. This approach relies on locking the laser to a specific feature on the ^{85}Rb absorption spectrum [1]. However, we cannot lock the laser the peak that represents the appropriate atomic transition because the PID circuit can only lock a system to a monotonic slope. We can lock to the nearby slope shown in Fig. 4.2 and use the AOM to shift to the exact value that we need. This is preferable to the exact resonance frequency of the ^{85}Rb transition since we need the frequency of our light to be adjusted to account for the Doppler shift that the moving atoms will experience. This also allows us to to keep the master laser locked to the correct point on the ^{85}Rb spectrum even if we need to adjust the frequency of our slave lasers. The AOMs allow us to tweak the the frequencies of both the slave lasers independently while leaving the master untouched. However, by glancing at Fig. 4.2 one can tell that there is about a 70 MHz frequency difference between the locking point and the cooling transition, but our AOM is only capable of shifting the frequency of incident light by 50 MHz. We address this issue by employing our AOM in a double-pass configuration.

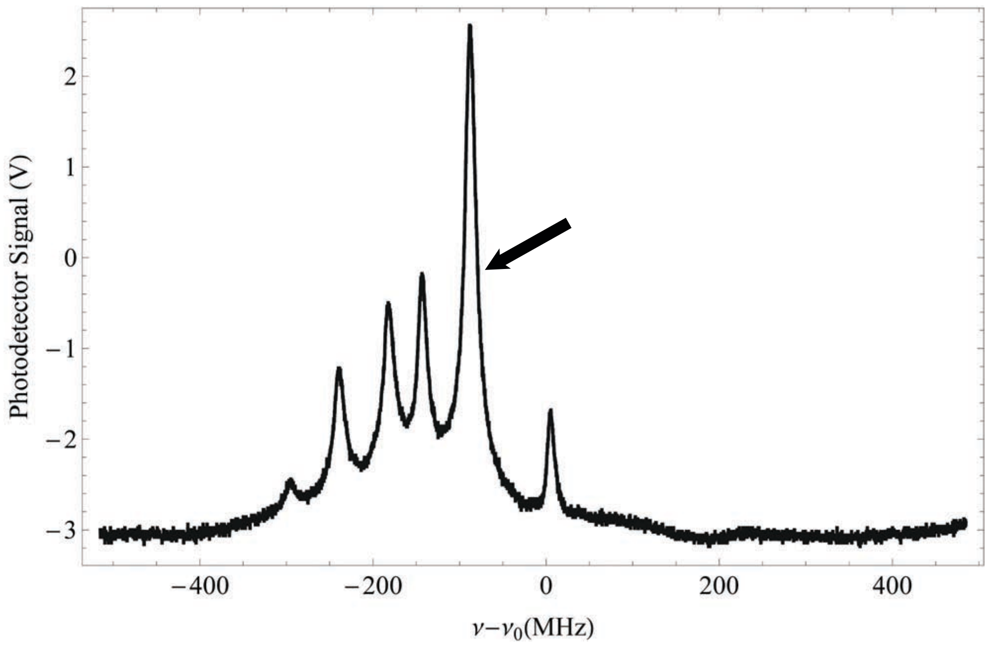

Figure 4.2: Ths figure shows the ^{85}Rb spectrum after the lock-in amplifier, and is marked to show where we wish to lock the laser. The laser cooling transition is the peak seen farthest to the right.

4.3 Double-pass Configuration

Rather than use multiple AOMs to shift our light's frequency by the necessary amount in a series of "single-pass" configurations we can send the diffracted beam back through the AOM from the other direction. This approach, visualized in Fig. 4.3, is known as a "double-pass" configuration, and it is advantageous for several reasons. The frequency shifted beams that emerge from an AOM in the single-pass configuration are also rotated by a small angle

$$\theta = \sin^{-1}\left(\frac{m\omega}{2\Omega}\right), \tag{4.3}$$

recalling that ω is the frequency of the photons and Ω is the wavelength of the phonons. [5] This angular change means that any adjustment in the AOM frequency would lead us to realign all our downstream optics or perform a very finicky lens alignment to collimate all the potential beams. The single-pass configuration is also incapable of introducing an adequate frequency shift to the laser beam, so we opt for the double-pass configuration.

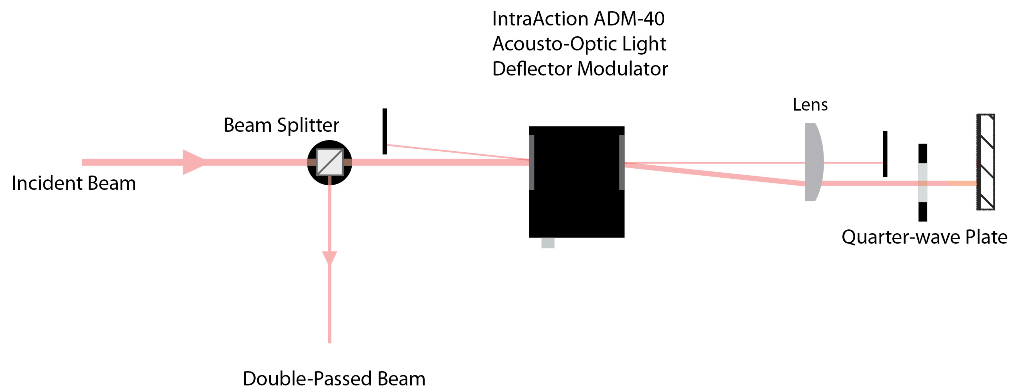

Figure 4.3: An AOM in the double-pass configuration. The incident beam is frequency shifted twice and returned along the same line with the opposite linear polarization.

In a double-pass configuration the incident beam passes though the AOM, which is rotated to favor the first order up-shifted beam. The diffracted beam then passes through a lens. This is placed according to a modified cat's eye configuration, ensuring that the double-passed beam emerges along the same path as the incident beam but with a different polarization. [6] In a normal cat's eye, the spacing between the lens and the mirror would be the lens's focal length, however we would like all the beams leaving the AOM to return along their outgoing path, so the lens is placed one focal length away from the AOM. A quarter-wave plate is placed between the mirror and the lens to rotate the polarization of the double-passed light so that it can be isolated from the incident light. The diffracted beam then passes through the lens and the AOM again where it is frequency-shifted again and redirected back to the path of the incident light. We then can use a polarizing beam splitting cube to separate the frequency-shifted light from the incoming beam. Beam blocks prevent unwanted orders from also getting double-passed and picked off by the beam splitting cube.

Passing our beam through AOMs comes at the cost of dramatic power reductions. The IntraAction ADM-40s used in this experiment are incapable of exceeding 80%

single-pass efficiency. This gives them a theoretical maximum double-pass efficiency of 64%, however their best performance does not exceed 50%. This loss of power is an inevitable side-effect of the double-pass configuration, because the additional optics create back reflections which sap power from the beam. Both of these efficiencies are very dependent on the driver frequency. This can be seen in Fig. 4.5 for single-passed light and in Fig. 4.6 for double-passed light. We do not know why there is a dip in the AOM frequency around 43 MHz; were we to use this RF driver value, we would have to investigate this further. Figure 4.4 shows how the AOM operates while down-shifting the double passed light. The scale of the frequency shifted light (the upper line) has been changed so that the signal in visible, it it much more attenuated than the unshifted light. The efficiency of our AOMs is not of great concern since the resultant frequency-shifted beams drive the slave lasers. They need to have enough power to create feedback with the slave laser, but do not have to meet the higher power requirements of laser-trapping. To shift the frequency from the lock point on the peak to the cooling frequency the AOM will run at 30-35 MHz. In this realm it is 25-40% efficient, so our ADM-40s are employable in our laser cooling apparatus.

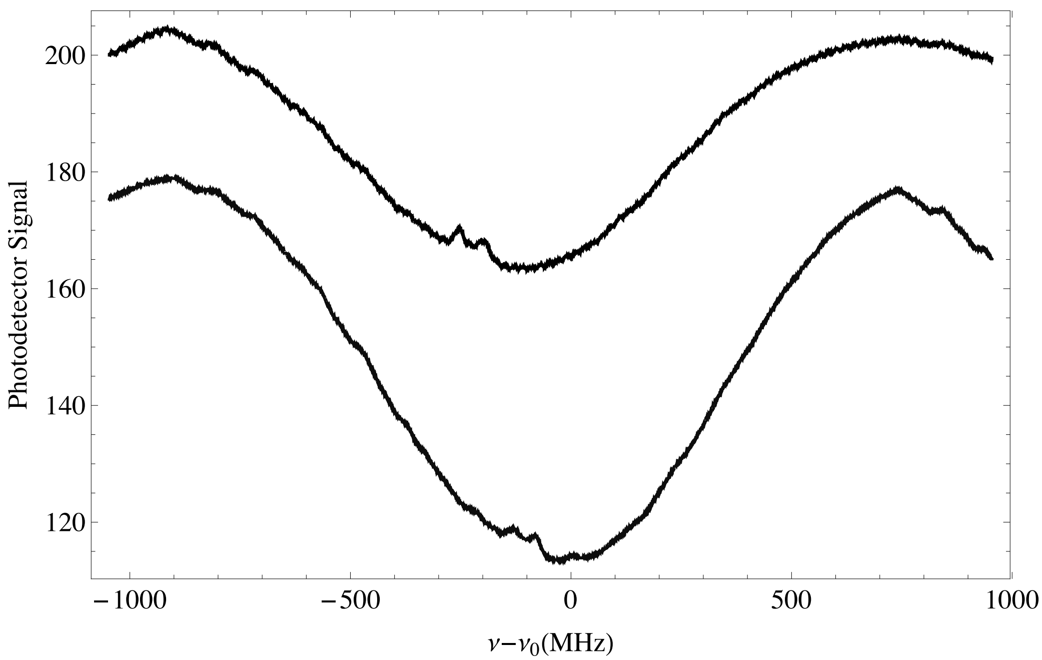

Figure 4.4: The upper trace shows the frequency of light after passing through the AOM while the lower traces shows the undisturbed light. Note that the upper line shows that the light is down shifted by about 100 Mhz, which agrees with our RF driver setting.

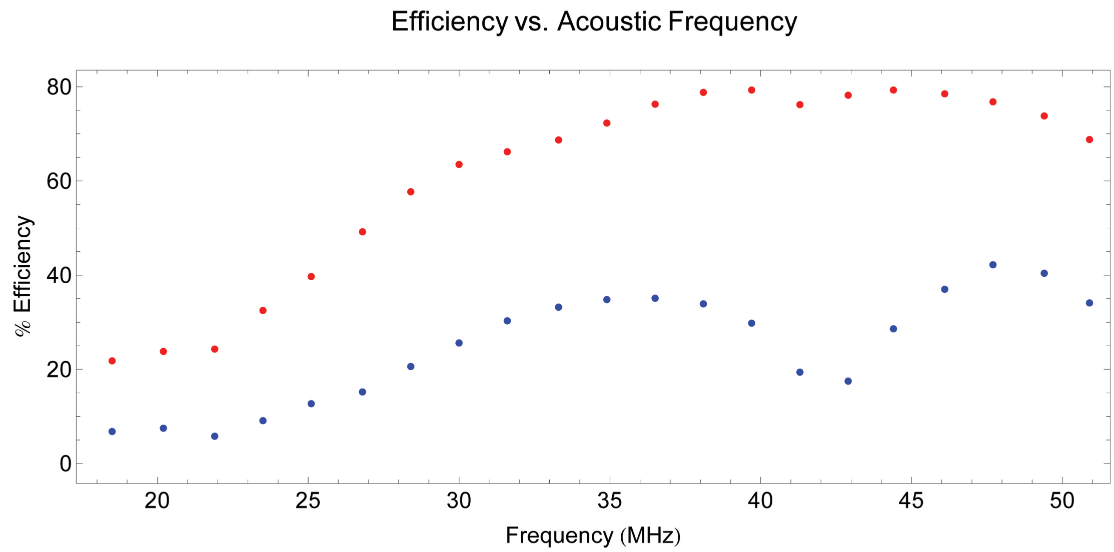

Figure 4.5: Plot of two different ADM-40s' single-pass efficiencies versus the acoustic frequency, one is in red, the other blue.

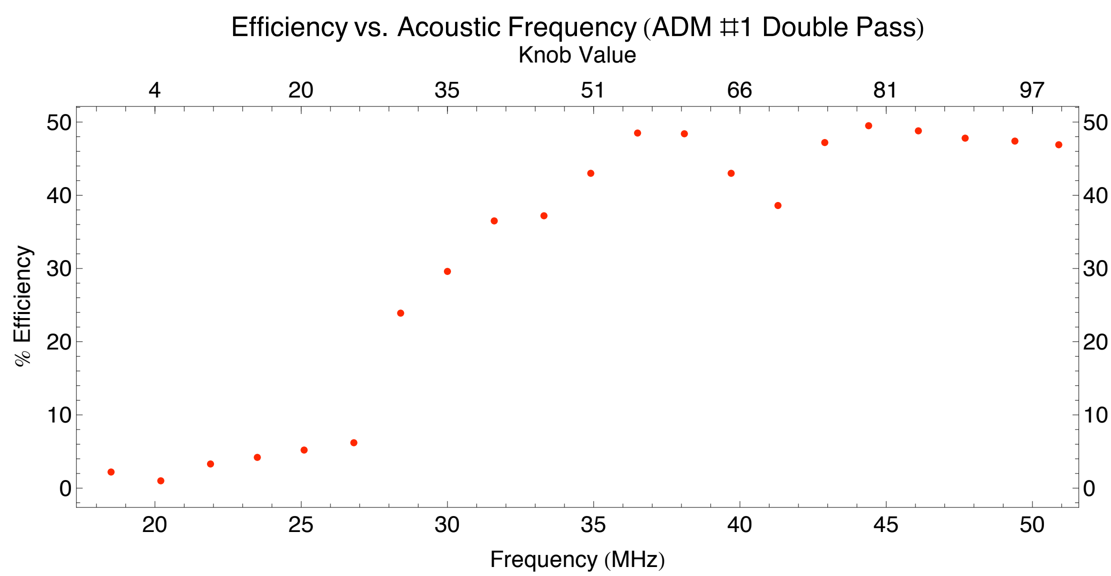

Figure 4.6: Plot of the best ADM-40's double-pass efficiencies versus the acoustic frequency.

Chapter 5

Preparing the Beams for Laser Cooling

Laser cooling is extremely dependent on the frequency of the light used. Several factors must be taken into account in the frequency tuning of the lasers to ensure that the laser light can cool the rubidium vapor. We must select the atomic transition carefully to maximize the number of atoms that our lasers can act on. Then we have to calculate the detuning of our laser to ensure that it slows the moving atoms. Finally we lock our laser to the detuned frequency using PID circuitry, as discussed previously.

5.1 Selecting the Atomic Transition

Our laser can be tuned over a several nanometer range of wavelengths granting us some freedom in selecting the atomic transition to use for laser cooling. The amount by which our laser can decelerate an atom is dependent on how often the atom can absorb and emit a photon [8]. To maximize this deceleration we have to keep our atoms in a two level transition cycle where the atoms are excited to the higher state by the laser and decay directly to the original lower state with a very high probability.

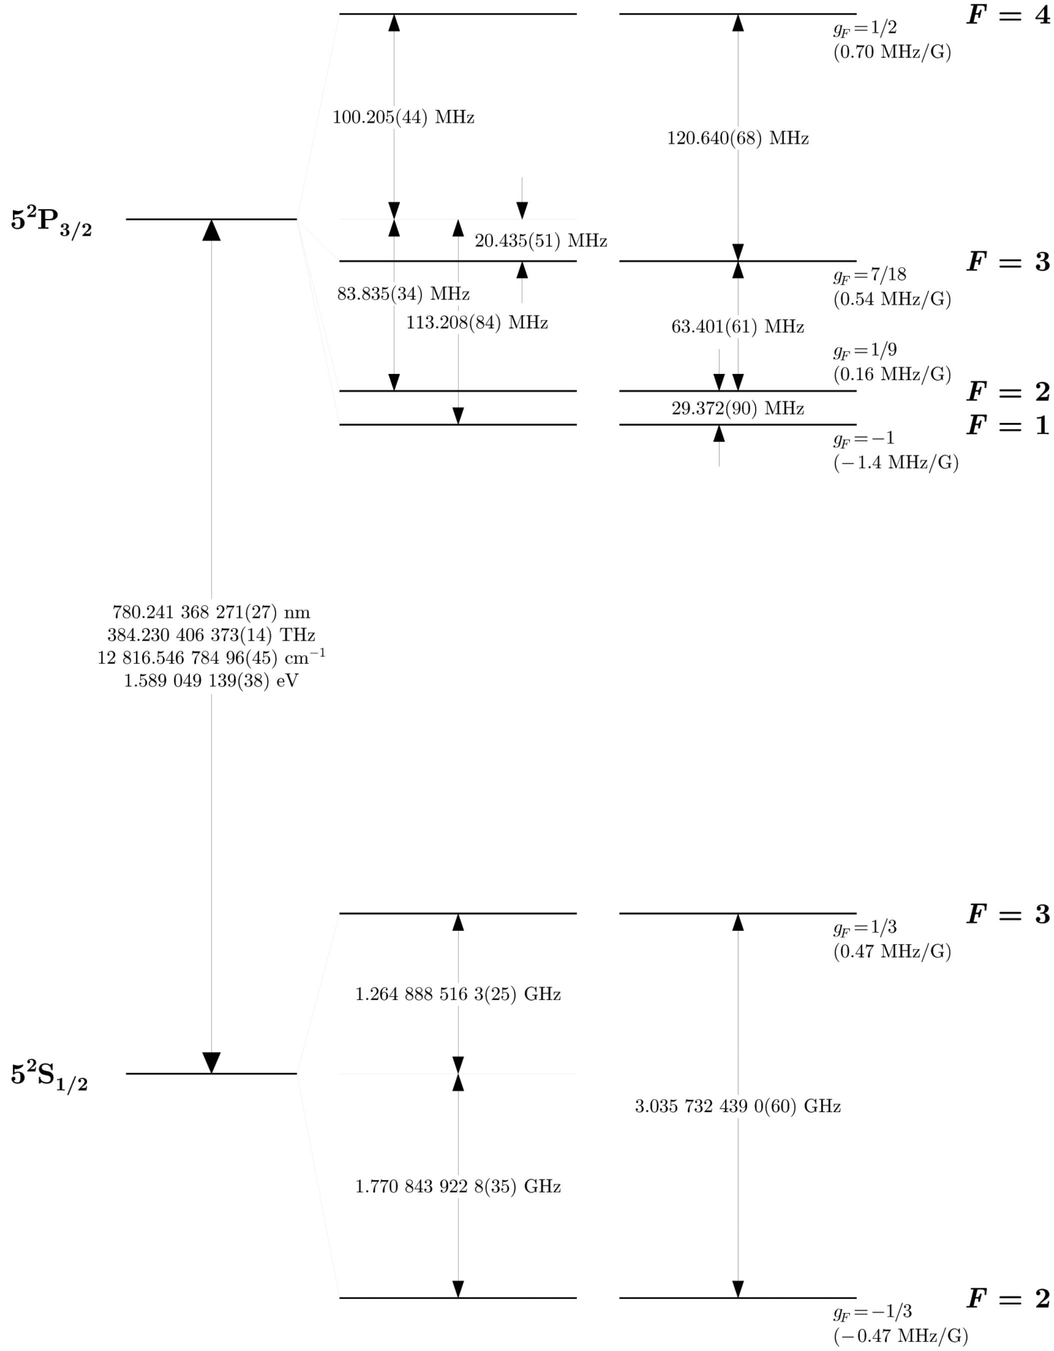

Figure 5.1: The energy levels of ^{85}Rb . After ref. [7]

As we can see in Fig. 5.1 the transition from the $F = 3$ level of the $5^2\mathrm{S}_{1/2}$ manifold to the $F' = 4$ level of the $5^2\mathrm{P}_{3/2}$ manifold fits this bill.

Atoms from the excited state are only allowed to decay back into the $F = 3$ state, so the system is almost closed. However, some atoms can be excited to the $F' = 3$ state and decay to the $F = 2$ state instead of the $F = 3$ state. We need to employ a re-pump laser to excite the atoms from the F= 2 state back into the $F' = 3$ state, guaranteeing that our atoms stay in the cooling transition. [8]

5.2 Finding the Force on the Atom

Now that we know the energy of the ^{85}Rb transition we would like our laser to excite we have to determine what frequency our laser light needs to be. To cool atoms we must create "optical molasses," a region of light that exerts enough force on the atoms traveling though it to slow them dramatically. This force is extremely frequency dependent and falls off rapidly as our laser becomes more and more detuned from the atomic resonance, ω_0. This detuning is often expressed as the difference between the frequency of the laser and this atomic resonance, written as $\Delta_0 = \omega_L - \omega_0$. The force on a single atom from a beam with wavelength $\vec{k}$ is

$$\vec{F} = 2\hbar\vec{k}\delta\omega_0\rho_{ee}, \tag{5.1}$$

where ρ_{ee} is related to the number of atoms in the excited state and $\delta\omega_0$ is a measure of the detuning. This relation is given by

$$\rho_{ee} = \frac{s_0/2}{1 + s_0 + \left(\frac{\Delta}{\delta\omega_0}\right)^2} \tag{5.2}$$

where s_0 is defined by our laser intensity I_L and the saturation intensity I_{sat} as

$$s_0 \equiv \frac{I_L}{I_{sat}} \tag{5.3}$$

and Δ is the effective detuning, corrected for the velocity of the atom and the presence of magnetic fields:

$$\Delta = \Delta_0 - \vec{\mu} \cdot \vec{B}/\hbar - \vec{k} \cdot \vec{\nu}. \tag{5.4}$$

Thus we can write the force on the atom as

$$\vec{F} = \hbar\vec{k}\frac{s_0\delta\omega_0}{1 + s_0 + \left(\frac{\Delta}{\delta\omega_0}\right)^2}.[6][1] \tag{5.5}$$

Now that we have the force that our light exerts on an moving atom, we can relate it to the force required to stop the atom. This will allow us to find the detuning of the laser, Δ, which will cool atoms in our optical trap. This force comes from simple kinematic principles where the distance D traveled by an moving atom that is experiencing an constant acceleration is

$$D = \frac{v_i^2}{2a}. \tag{5.6}$$

We can rewrite this equation in as an expression for a

$$a = \frac{v_i^2}{2D} \tag{5.7}$$

and using $F = ma$ we can express force as

$$\vec{F} = \frac{m|v_i|^2\hat{v}_i}{2D}. \tag{5.8}$$

Combining our two expressions for force we have

$$\frac{m|v_i|^2\hat{v}_i}{2D} = \hbar\vec{k}\frac{s_0\delta\omega_0}{1 + s_0 + \left(\frac{\Delta}{\delta\omega_0}\right)^2} \tag{5.9}$$

which can then be used to find the expression for Δ

$$\Delta = \delta\omega_0\sqrt{\frac{2\hbar\vec{k}s_0\delta\omega_0 D}{m|v_i|^2} - 1 - s_0}. \tag{5.10}$$

This allows us to calculate the amount by which we need to shift our laser frequency to account for the motion of the atoms and ensure that they are slowed before leaving the confines of the trap.

5.3 Locking to the Resonant Frequency

Now that we can find the optimal detuning for our laser we have to lock the output frequency to this value. We first have to find the frequency of our laser. This is done by passing our laser light through a Doppler-free spectroscopy apparatus [1], producing a spectrum that looks like Fig. 5.2. We then place a ND filter in the beam path to reduce power broadening and use a lock-in amplifier to pull the signal out of the noise as shown in Fig.5.3.

PID circuits can only be used to stabilize the laser at certain points on our spectra, which do not correspond to the correct detuning. We cannot use a PID circuit to lock our frequency to a peak of the spectrum as there is no way to electronically distinguish the difference between the voltage drops on either side of the peak. However, PID circuits can lock to monotonic slopes where a change in input voltage is directly correlated with a change in output voltage. Thus we must lock to a slope close to our desired frequency.

We use a Sacher Lasertechnik LB2001 servo controller to lock the master laser to the slope. The locking behavior can be seen in Fig. 5.4. To orient ourselves, channel 2 of the oscilloscope trace is the output of the lock-in amplifier while channel 3 is the error signal of the LB2001 and channel 4 is its output. On the left of Fig. 5.4 we can see the ^{85}Rb absorption spectrum after the lock-in amplifier. This has been offset so

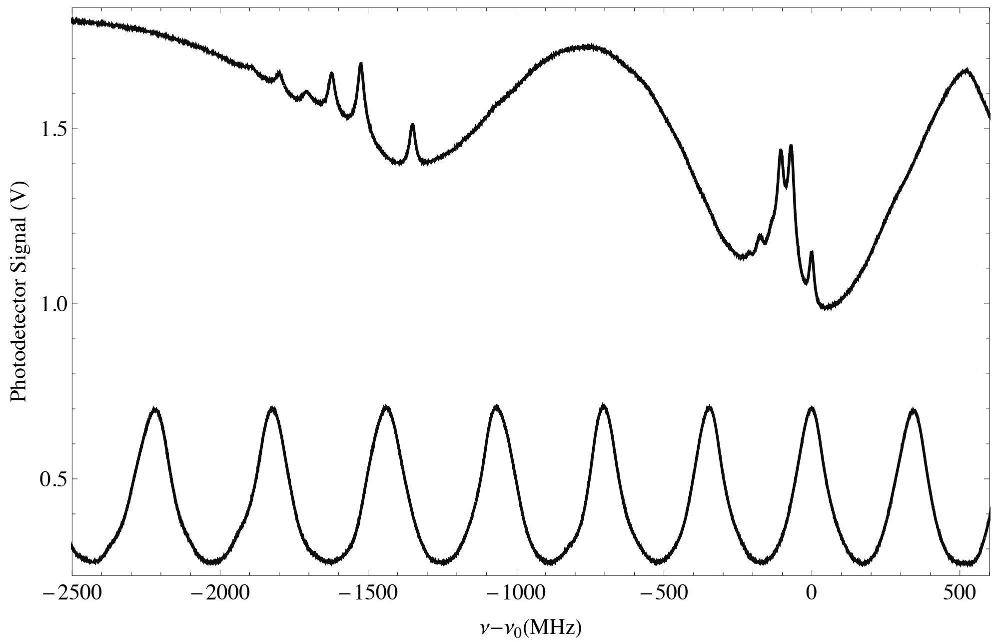

Figure 5.2: The upper curve is a photo detector signal showing the Doppler-free absorption spectrum for Rb while the laser's frequency scans. The lower peaks are from a Fabry-Pérot, where a new peak occurs every time the laser's frequency scans by 300 MHz.

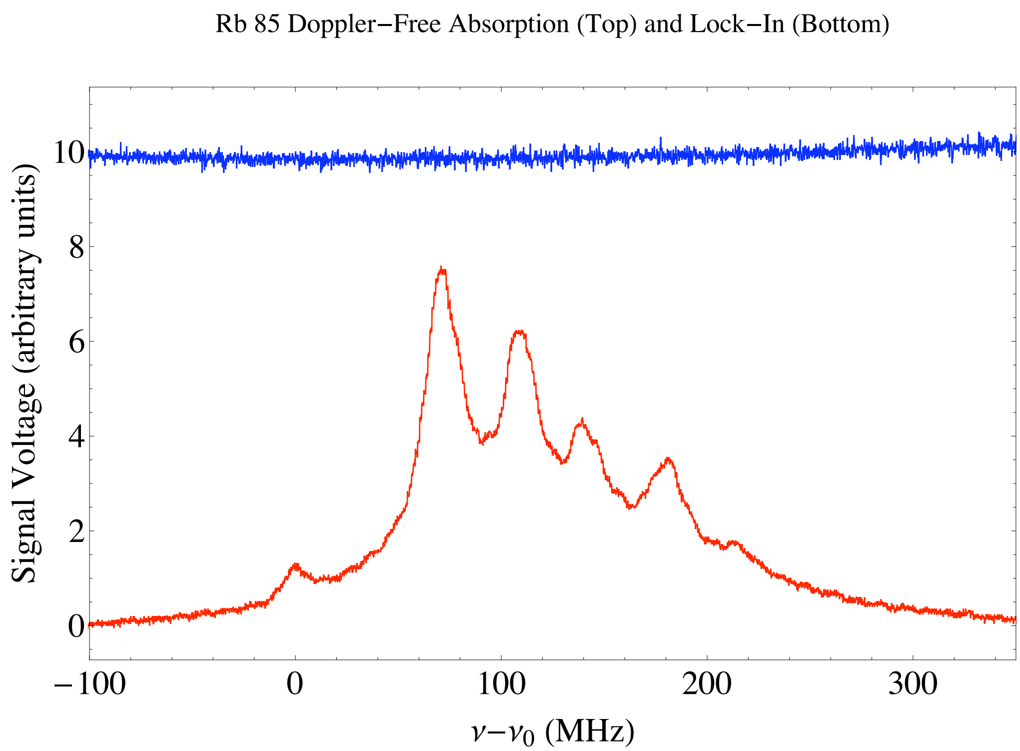

Figure 5.3: The upper curve is the same photo detector signal of the Doppler-free absorption spectrum for Rb after an ND filter was placed in the beam path. The lower curve is the absorption spectrum after a lock-in amplifier.

that the voltage is zero at the frequency where we wish to lock the laser. Channel 4 shows how the piezo voltage is being used to sweep the laser frequency. Once the frequency of the laser is close to the setpoint, the LB2001 is switched into "lock" mode. On the trace, this is when the output drops suddenly. After this change we can see that they LB2001 tries to keep the error signal at or around zero, effectively locking holding the frequency constant. The LB2001 is able to keep the laser locked despite small disturbances to the table or the optical setup. Large voltage changes, like static shocks, and obstructions in the beam path unlock the laser. The locking does not hold the laser to the exact frequency that we desire, but rather fluctuates rapidly around the set-point. These fluctuations are on the scale of 500 mV, or $\Delta\lambda \approx 0.1$ nm. This fluctuation is tolerable and should not reduce the efficiency of our laser cooling apparatus. Once we have the laser locked we can use an ADM to shift they frequency of the light to the optimal value for laser cooling and then feed this light into our slave lasers.

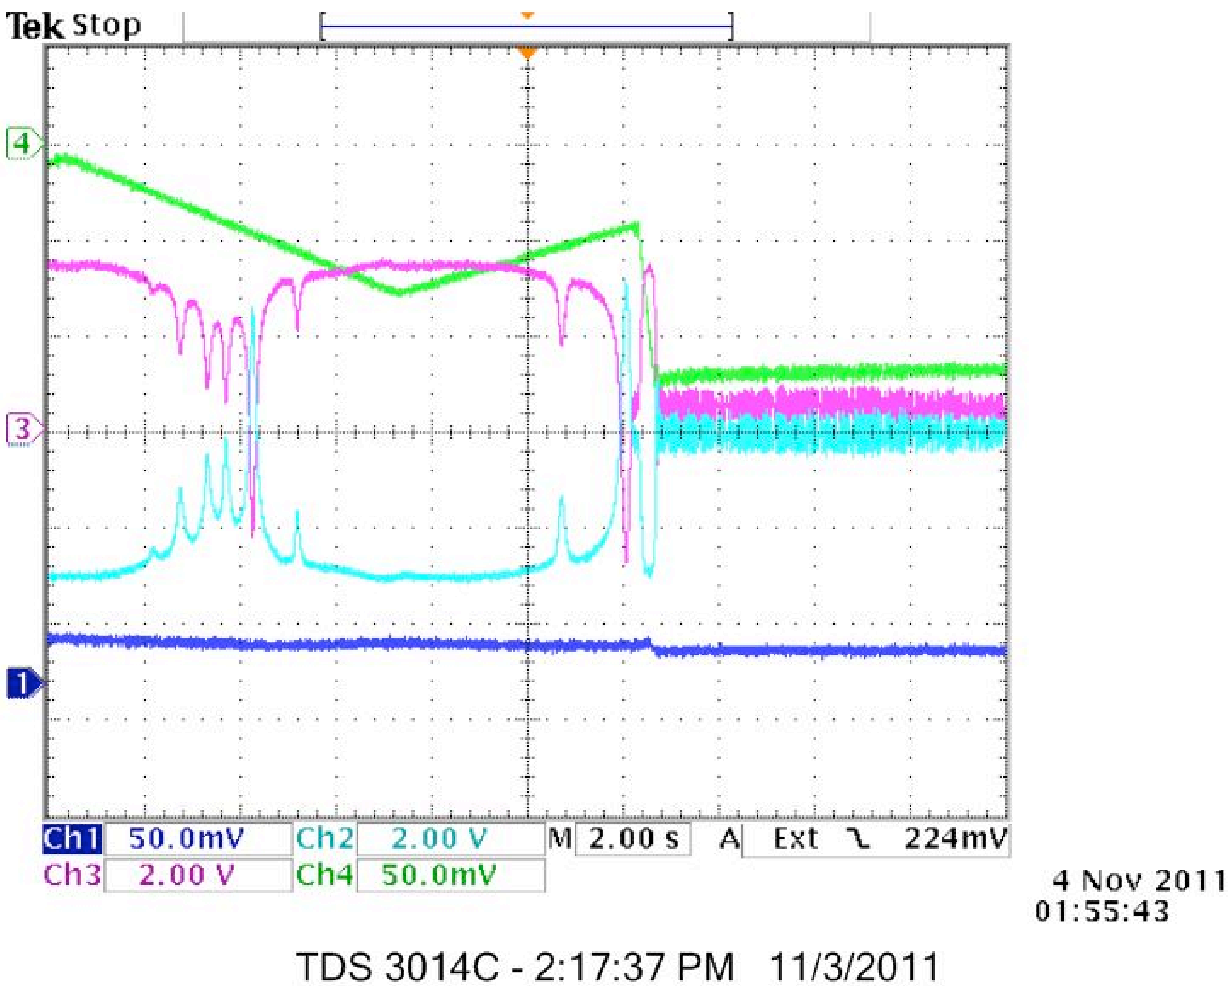

Figure 5.4: This shows how the LB2001 behaves after locking to a slope on the ^{85}Rb spectrum. In this plot the green curve is the output of the LB2001, the pink is the error signal, and the light blue is the signal from the lock-in amplifier. At first the LB2001 scans the laser frequency, then at about 2 seconds, it is turned on and begins to stabilize the laser frequency.

Chapter 6

Conclusion

Laser cooling can only occur when photons from the laser have energies that correspond to an atomic transition within the atom. This task is by no means simple, since the frequency of the laser depends on its temperature, injection current and cavity length. In this thesis methods for stabilizing all of these parameters are proposed using electronic feedback circuits known as PIDs. If all three of these factors are regulated, then small adjustments can be made to hold the laser frequency constant.

Temperature control is accomplished by using a PID feedback loop that drives TECs. This sort of temperature control is used on all our lasers, both master and slave. The response time of the circuit is slowed by the thermal mass of the laser enclosure, so changing the temperature of the laser is not an effective way to rapidly correct for frequency drift. Injection current is set by the operator and could be used to change the frequency of the laser. However, changes in the injection current can lead to “mode-hopping” (rapid jumps in the frequency of the laser) so this is a less desirable option. [1] Finally, the cavity length is controlled by a piezo that can respond rapidly to changes and has a wide range. A commercial servo controller is used to drive this piezo in accordance with a frequency error signal generated using Doppler-free spectroscopy. These techniques allow us to lock the frequency of the laser, however we must detune it from atomic resonance to trap atoms.

Small frequency shifts, like detuning, are the domain of the AOMs. These devices can be placed anywhere in the beam path to make the necessary frequency adjustments wherever they are needed. Using a mix of commercial and homebuilt technology on two different laser housings, this thesis demonstrates how the frequency of a laser is locked to a specific, useful value and how that frequency can be shifted as needed.

Appendix A

Slave Laser Housing

To ensure adequate temperature control of the slave laser, a housing was designed with one small TEC directly below the laser diode and two larger TECs for cooling the entire enclosure. The entire unit was mounted to two water blocks which serve as heat sinks for the large (2x2") TECs. This appendix contains the drawings for the housing. The assembled unit can be seen in Fig. 3.3.

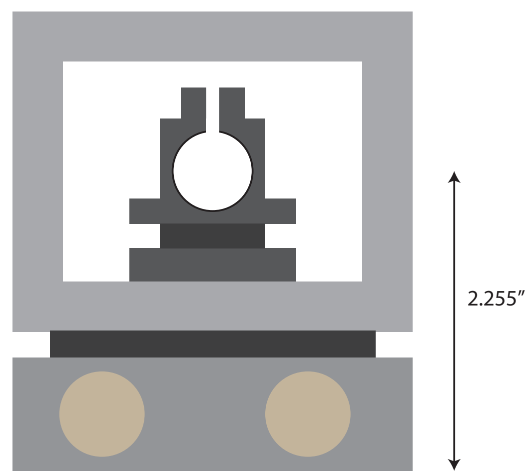
2.255″

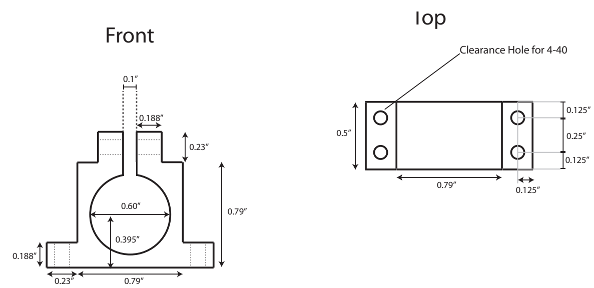
Front
0.1"
0.188"
0.23"
0.79"
0.60"
0.395"
0.188"
0.23"
0.79"
Top
Clearance Hole for 4-40
0.5"
0.125"
0.25"
0.125"
0.79"
0.125"

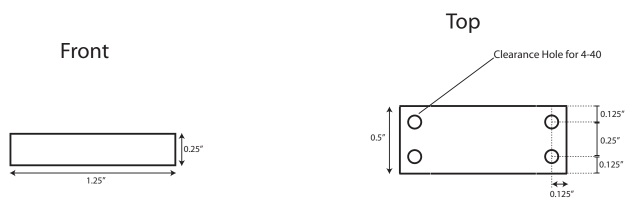
Front
0.25"
1.25"
Top
Clearance Hole for 4-40
0.5"
0.125"
0.25"
0.125"
0.125"

Top

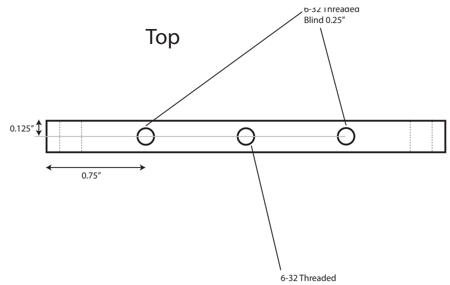

Front

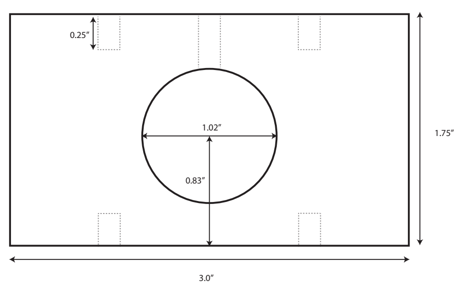

Horizontal Cross Section

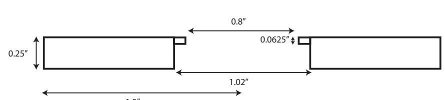

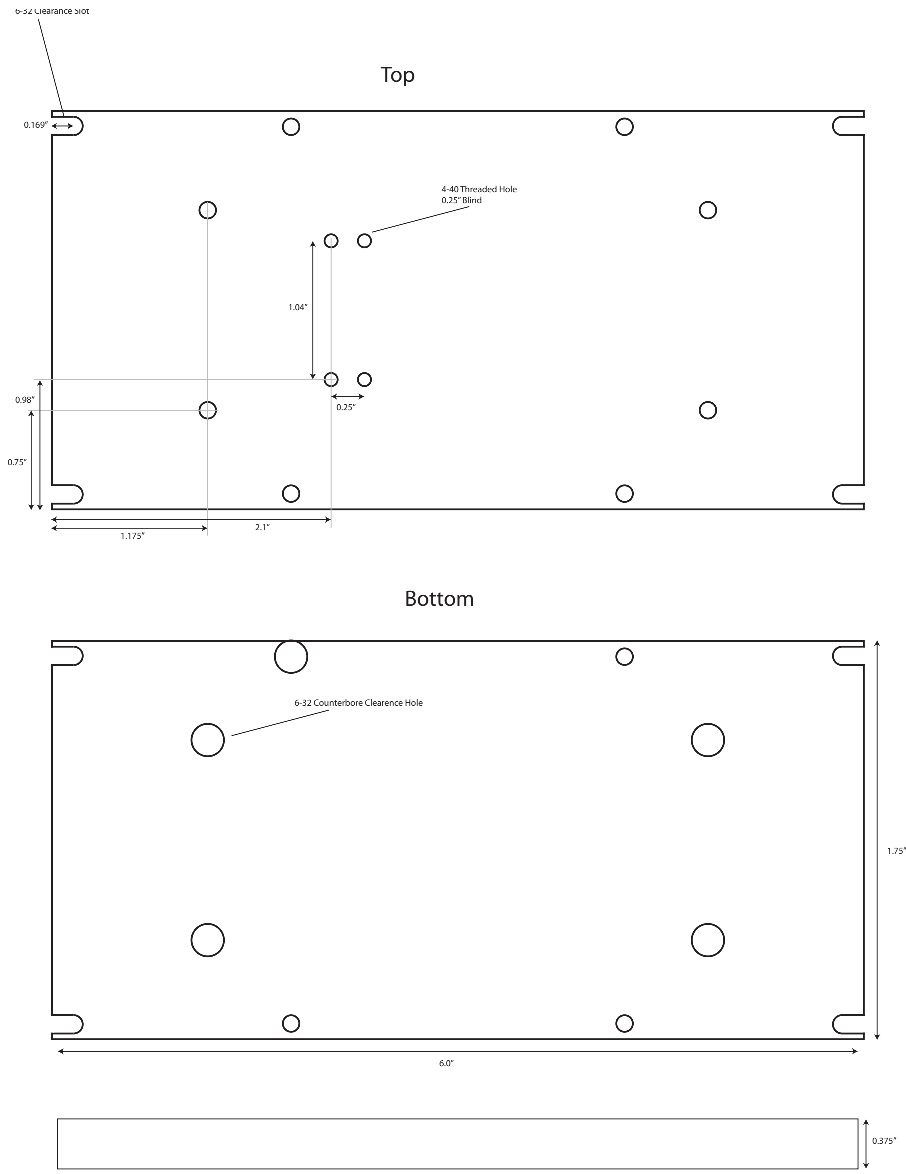
6-32 Clearance Slot
Top
0.169″
4-40 Threaded Hole
0.25″ Blind
1.04″
0.98″
0.25″
0.75″
2.1″
1.175″
Bottom
6-32 Counterbore Clearence Hole
1.75″
6.0″
0.375″

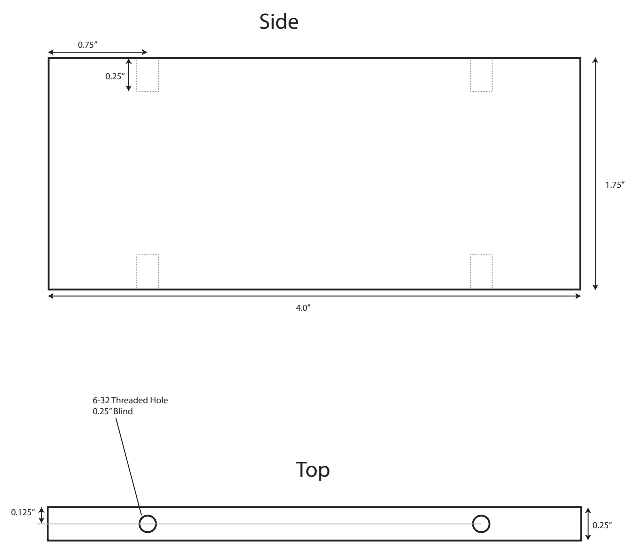
Side
0.75"
0.25"
1.75"
4.0"
6-32 Threaded Hole
0.25" Blind
Top
0.125"
0.25"

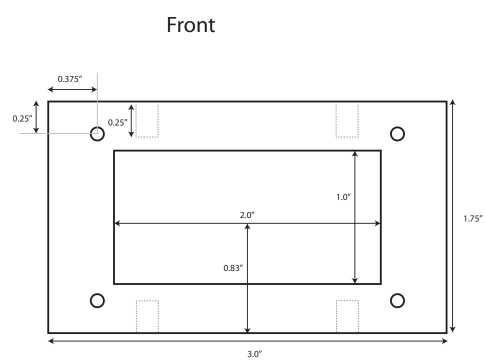
Front
0.375″
0.25″
0.25″
1.0″
2.0″
1.75″
0.83″
3.0″

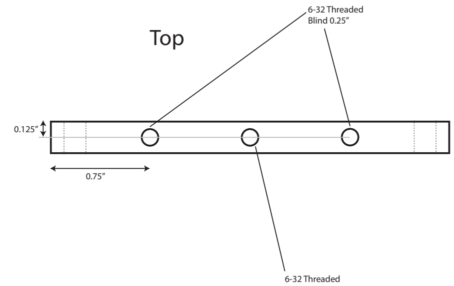
Top
6-32 Threaded
Blind 0.25″
0.125″
0.75″
6-32 Threaded

Bottom

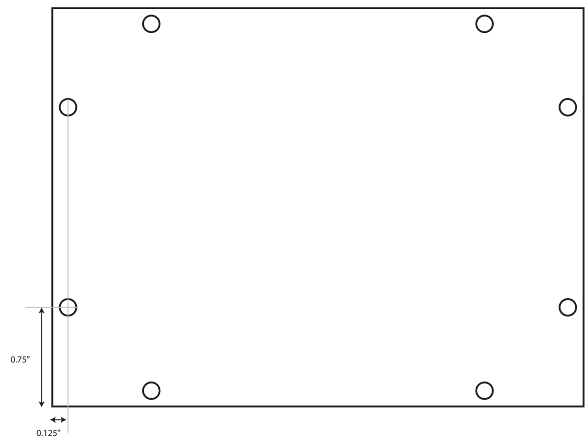

Top

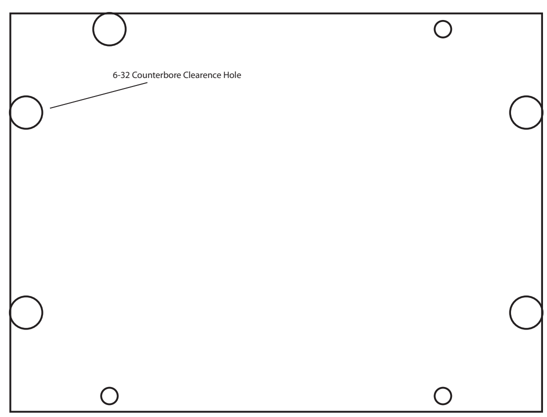

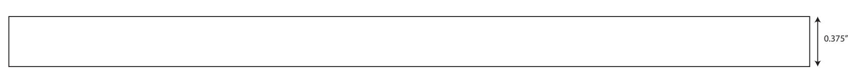

Bibliography

[1] Thomas Stark. The road to laser cooling rubidium vapor, 2011.

[2] Lightwave. Temperature Stability Using the LDT-5948.

[3] C Wieman, G Flowers, and S Gilbert. Inexpensive laser cooling and trapping experiment for undergraduate laboratories. *American Journal of Physics*, 63(4):317–330, April 1995.

[4] S. G. Lipson. *Optical physics*. Cambridge University Press, Cambridge Eng New York, 1981.

[5] Leo Levi. *Applied optics : a guide to optical system design.* Wiley, New York, c1980 1968.

[6] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. Double-pass acousto-optic modulator system, 2005.

[7] Daniel Adam Steck. Rubidium 85 D Line Data, 2010.

[8] Anne L. Goodsell. Capture of laser-cooled atoms with a carbon nanotube, 2010.