Frequency Stabilization of Diode Lasers for Ion Interferometry

Jarom S. Jackson

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

Bachelor of Science

Dallin S. Durfee, Advisor

Department of Physics and Astronomy

Brigham Young University

April 2013

ABSTRACT

Frequency Stabilization of Diode Lasers for Ion Interferometry

Jarom S. Jackson
Department of Physics and Astronomy
Bachelor of Science

This thesis describes the work done to set up two external cavity diode lasers to be used in an ion interferometer. These lasers will be used to trap and cool atoms and to manipulate and probe their energy states. The main cooling laser is generated using a doubled IR laser, but is very unstable due to the various stages—diode, external cavity, doubling crystal and cavity—that all need feedback to lock the frequency. The probe laser is much simpler because the wavelength needed matches that of a readily available laser diode. The stability and scanning ranges of both lasers are presented.

Keywords: lasers, atoms, cooling, trapping, locking, matterwave, interferometry, MOT

ACKNOWLEDGMENTS

Thank you to those who have made it possible for me to be involved in this research. In particular, I am grateful for my parents, who provided me with the opportunites that have led me to be where I am, and Dr. Durfee, who has invested a significant amount of time in helping me to learn how to work in a physics laboratory.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 BYU Matterwave Interferometry Lab

The atom is such an incomprehensibly small and foreign object, that observing or manipulating it is a seemingly impossible task. The invention of the laser, however, has made it possible to probe and interact with the internal energy structures of atoms, thus opening up a field of study and application dealing with the manipulation of atoms and their wave functions.

At the BYU matterwave interferometry lab, we are in the process of building an experiment that will take advantage of the wave nature of atoms in order to perform matterwave interferometry. Such an experiment requires a host of electronic and optical subsystems, each of which can be a lengthy project in and of itself. The purpose of this thesis is to detail the work done on two of the key laser systems that will be used in the ion interferometer. However, before we get to the details, some introduction to the field is in order.

1.2 Atomic Physics and Precision Measurement

Energy levels, atomic transitions, and light interactions are the realm of atomic physics. Energy levels in an atom are quantized, and because of this an atom can only absorb or emit light at certain frequencies. These energy levels are dependent on properties internal to the atom, as well as external fields. Experiments in atomic physics make use of these facts to manipulate and probe atoms.

Precision measurement is an area in which atomic physics has excelled. Since the energy levels in an atom are very sensitive to external fields, the presence and magnitude of those fields can be found experimentally by probing atoms in various manners. Examples of things measured using similar ideas include the gravitational constant and the value of the earth's gravitational field [1].

Making use of atomic physics for experiment and measurement requires some controllable manner in which to interact with and control atoms. This requires a great deal of precision and control, which are not easy to achieve. The best tool that we have available for such things is the laser, a key instrument in atomic physics.

1.3 Laser Use in Atomic Physics

Most of the current work in atomic physics would be difficult, if not impossible, without lasers [2]. It is hard to imagine how we could accomplish much without them. What makes lasers so useful for probing the atom is the high level of control that is possible over various characteristics of the light, most important of which, for our purposes, are the wavelength and linewidth. By controlling these variables it is possible to manipulate atoms in various ways.

Control over the wavelength, linewidth and stability of a laser is key to its usefulness in experimental physics. Linewidth is a measurement of the spread of frequencies that make up the spectrum of the laser (see Appendix B). The light output by some lasers may have a linewidth that is much larger than the natural linewidth of atomic transitions (i.e. the range of frequencies that the atom will absorb for a given transition). Attempting to use a laser with such a wide linewidth in order to manipulate and probe atoms would be like trying to use a blow torch to solder tiny components on an electronics board. It is just not the right scale for the job.

Even if a laser does have the right linewidth, however, its instability could still make it useless. Instability is inherent in the way that many lasers work. At the heart of most lasers, including diode lasers, is an optical cavity in which the light resonates. Depending on the light and form of the cavity, there are usually many different ways in which the light can resonate. These are called modes, and each one tends to have slightly different characteristics—most importantly, for our work, their wavelength. The light coming from a simple laser is typically a mixture of various competing modes, and they are constantly changing. The wide linewidth and the instability of many lasers come from the competing modes, as well as the linewidth and instability of individual modes.

A laser is called multimode if it has more than one competing mode lasing at a time. Each mode may have a relatively narrow linewidth, but since they are all slightly different in frequency, the overall linewidth of the laser is widened. Figure 1.1 gives a qualitative comparison between the frequency spectra of multimode lasers and the frequency spectra of single mode lasers. Even if the diode has just one mode, however, that mode will typically be short lived, and the laser will constantly 'hop' from one mode to another.

If the linewidth and instability of a laser can be controlled, then the laser becomes a powerful tool for studying the atom. Atoms absorb and emit only photons that correspond to certain discrete wavelengths. This property is useful for both manipulation and measurement. Since the atom only absorbs certain wavelengths in any given energy state, controlling the wavelength of the laser allows for measurement of the energy levels of the atom. Conversely, by manipulating the energy levels in the atom (e.g. with magnetic fields), it is possible to control the conditions in which an

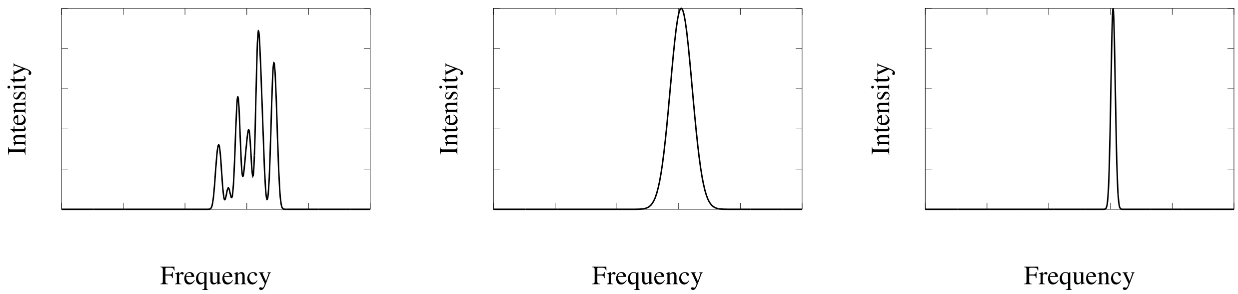

Figure 1.1 Comparison of laser spectra. The frequency spectrum of a multimode laser is spread out and non uniform (as is illustrated in the first graphic). Even if the laser is single mode, it may still be too wide (like in the second). Compared to the first two we need a single narrow mode like in last figure.

atom will emit or absorb light. Absorption of a photon affects the kinetic properties of an atom, so this also allows for a controllable way to move, speed up, slow down, or even trap an atom.

One of the key developments of modern experimental atomic physics is the use of lasers to cool and trap atoms [3]. It is possible to generate a position and velocity dependent force on an atom by clever manipulation of the atomic energy levels with a position dependent magnetic field, and by detuning a laser from the natural resonant frequency of an atomic transition. This trapping force is used to cool the atoms and confine them to a small volume in space. These ultracold atoms are then available to perform any number of experiments on. This type of an atom trap is referred to as a magneto optical trap (MOT) .

The MOT is often a starting point for experiments in atomic physics. It provides a stable group of atoms in a well defined state. From that starting point there are a lot of interesting experiments that can be, and are being, done. Atoms from a MOT are used to perform precision measurements by probing their states before and after performing various manipulations. In our lab, we plan to generate a beam of ions from a MOT, which will go on to form the two arms of an ion interferometer.

1.4 Diode Lasers

Despite the utility of lasers in probing, manipulating and trapping atoms, there are many challenges to using them. To name a few, they can be very expensive to purchase and to upkeep, they often require a lot of time and attention to both setup and maintain, and they can be very fragile. There are various types of lasers in use today, each one of which has its own strengths and weakness among those mentioned above. The type that we use in our lab, which will be the focus of this thesis, are diode lasers.

Diode lasers have a number of advantages. They are inexpensive, relatively small, and don't require anything to run other than a steady current and a controlled temperature.

The key problems with diode lasers are their availability, tunability, linewidth, and instability. Availability and tunability are a problem because diodes are manufactured to emit a very specific wavelength of light and cannot be tuned very far from it. As such, they can only be used in an experiment if a diode close enough to the wavelength needed is commercially available, which is not always the case. Linewidth and instability are also key problems, because a bare laser diode is typically multimode. Even if it is at roughly the right wavelength, the competing modes and resulting instability and wide linewidth make the laser useless for controlled interactions with atomic transitions.

In order for diodes to be useful in our experiments, we need some way to overcome these deficiencies. Luckily, there are some ways to do so, which include temperature and current control, as well as setting the diode into a external cavity arrangement. These techniques will be described in more detail in Chapter 2.

When these technical difficulties are overcome, the diode laser can be a powerful and affordable tool for manipulating and probing atoms. We make use of them in various ways, the end goal of which is to perform experiments that utilize the wave nature of matter.

1.5 Matterwave Interferometry

Matterwave interferometry is an exciting area of atomic physics in which the wave nature of matter is used to make measurements [4]. Matterwaves have extremely small wavelengths (typically much smaller than visible light), which makes them very sensitive and useful for precision measurement. Interferometry is one method of making use of the nature of waves to measure movement—such as displacement and rotation—and fields—such as electric, magnetic and gravitational fields. While light is usually used in interferometry, matterwaves have some unique properties that light does not. They have mass, a much smaller wavelength than visible light, and are affected by electric and magnetic fields.

Interferometry with matterwaves works, at least conceptually, the same as with light. A light interferometer splits an electromagnetic wave into two parts, which travel some distance and are recombined. When the two halves of the wave meet up again, the phase difference in the two waves determines to what degree they interfere constructively or destructively. Since the wavelength of light is so small compared to most things we work with on a macroscopic scale, light interferometers are really good at measuring small displacements on the order of a wavelength of light (half a micrometer). A tiny movement (down to fraction of a micrometer) in one of the interferometer arms causes a large change in relative phase, and that change can easily be observed in the output of the interferometer.

Because interferometers measure phase differences, they can be sensitive to anything that causes a change in phase between the two arms of the interferometer. Displacement is just one example. Light accelerometers can also measure acceleration, for example, by using the phase difference in the interferometer arms caused by the acceleration of the device. Similarly, matterwaves can be used to measure anything that affects their phase.

The phase evolution of a quantum wave is proportional to $e^{-i\frac{E}{\hbar}t}$, where E is energy. Because the energy of a particle includes potential energy from electric, magnetic or gravitational fields, its phase evolution does as well. That means that a matterwave interferometer can be sensitive to differences in field potential along the two different arms of the interferometer. In many cases this sensitivity is quite extreme, allowing for measurements that are orders of magnitude more precise than are possible otherwise [1].

1.6 Ion Interferometry

If atoms can be used for precision measurement, could it be possible to use ions as well? Ions provide a significant advantage over neutral atoms, for certain measurements, in that they are much more susceptible to electric fields. They also present some additional challenges, since the net charge causes them to repel each other. Despite these difficulties, if interferometry could be performed with ions, then it could be used to make very fine measurements of electric fields. We are building an ion interferometer for this purpose.

The potential sensitivity of ion interferometry is several orders of magnitude better when making certain precision measurements involving electric fields than what has been accomplished to date. For example, laboratory experiments attempting to measure the photon rest mass (a theoretical value that is smaller than our capability to measure, if it exists at all) have established an upper limit of about 10^{-20} times the mass of an electron. An ion interferometer with the dimensions of what we are building has a theoretical sensitivity that is two orders of magnitude better. [5]

1.7 Lasers in an Ion Interferometer

To implement an ion interferometer, we need to generate a slow, stable beam of ions. This requires a series of lasers to trap, cool, and ionize the atoms (see Fig.1.2). We plan to generate this beam

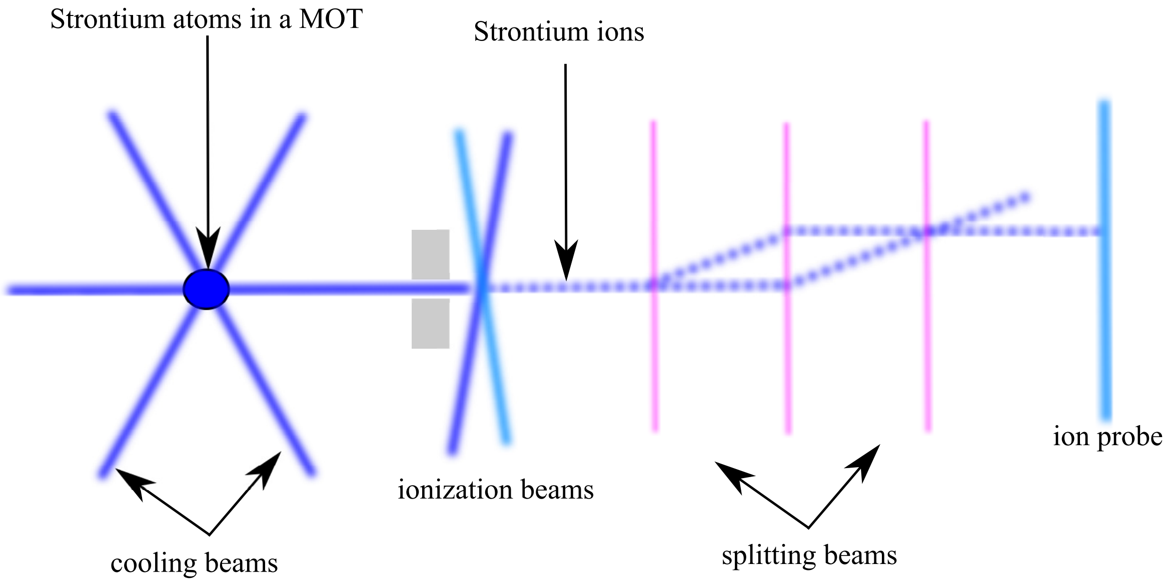

Figure 1.2 Schematic of the interferometer layout. The atom beam travels from the MOT (left) through a hole in a mirror, then through a pair of lasers where it is ionized, another set of lasers where it is split, recombined and then finally probed.

from a MOT of Sr atoms.[1] Creating the MOT requires a laser that is frequency locked to the resonant transition of strontium (i.e. the the lowest allowable transition from the ground state). We will use this laser to cool and trap the atoms .

From the MOT we will then generate an atom beam by leaving a small hole in the trapping potential. The hole in the trapping potential was created simply by drilling a hole in one of the mirrors used to retro-reflect a beam back on itself. Atoms escaping from this hole in the light field will be relatively slow and homogeneous in speed and direction.

The atoms in the beam coming from the MOT will then be ionized to form the slow ion beam. The ionizing process we will use is a two photon transition that uses the same excited state we are using for laser cooling. Because of that, we will reuse the cooling laser as part of the ionizing

[1] Sr was chosen because its cooling and trapping properties are understood, it can be easily ionized, and the ion has energy levels corresponding to wavelengths that are convenient for us to generate.

process. In addition, we need one more laser at ~405nm to ionize the Sr. Since the linewidth of this transition is relatively wide, the 405nm laser will not need to be actively locked.

As the ion beam travels through the interferometer chamber, it will be split and recombined. A pair of lasers locked at a precise offset from each other will be used to accomplish this (Offset Reference, Appendix C).

Finally, we will need a laser to probe the ions as they leave the interferometer. This will be a 408nm laser—right on the borderline between visible and UV light. This laser also needs to be precisely tuned, and will be frequency locked to the resonant transition in the Sr ions. Ions exiting from the interferometer will be in one of two hyperfine states. Interference of the matterwaves as they are recombined defines the probability of the ion exiting in one state or the other. By probing the ions to find which state they are in, on average, we can determine the difference in phase of the two paths as the ions leave the interferometer.

Implementing these laser systems is one of the most challenging and time consuming aspect of the construction of the ion interferometer. This thesis will focus on that task, specifically on the cooling (461nm) and probing (408nm) lasers. Many of the techniques will be discussed in general, as they apply equally to both. Other aspects of the interferometer, including the vacuum systems, electronics, the theoretical groundwork, and the lasers used to split and recombine the ion beam will not be discussed herein.

Chapter 2

Methods

2.1 What Makes a Good Laser?

Before we discuss in detail the methods used to make diode lasers useful for experimental atomic physics, we need understand and quantify what we are looking for. What exactly is it that we require in a laser?

The answer consists of three main factors. The laser needs to have a narrow linewidth, be tunable, and be temporally stable. Those are the general specifications of what were looking for, but how do we measure the laser, and what are the tolerances?

Lets start with the first requirement, linewidth. The linewidth of a diode laser is very closely related to the number of modes. We need the laser to be single mode, and we need that mode to be very narrow. To be more precise, the constraints depend on the atomic transitions that we wish to drive. That is, the linewidth of the laser must be at least as narrow as that of the natural linewidth of the transition it is to drive, which is inversely related to the lifetime of the upper state in the transition. The two transitions that we are concerned with are the resonant transitions of the Sr atom, and of a singly ionized Sr ion. The life time of the first excited state of Sr is 5×10^{-9}s, corresponding to a natural linewidth of about 200Mhz, and the natural linewidth of the transition we will probe in the Sr ion is 140MHz [6]. So the first requirement of our lasers is that their linewidth be on the order of 100MHz or less.

The second thing mentioned was tunability, which refers to our ability to alter the wavelength of light the laser emits. We have two requirements for tunability: rough adjustment and continous scanning ability. Ideally we only need the laser at a single very specific wavelength, or frequency. Our instrumentation, however, is limited in its ability to measure wavelength to the precision we require. At best we can measure the wavenumber of the laser to about $.1\text{cm}^{-1}$, which corresponds to about 3GHz (Appendix A). So our requirement for rough adjustment is that we can get the wavelength to within about 3GHz of resonance with the atomic transition.

Once the laser is roughly tuned (within about 3GHz) we then need to be able to scan (i.e. continuosly tune) the frequency over a range of about 3GHz without the laser mode hopping, or the amplitude varying drastically. We need this flexibility to find the exact frequency at which the laser resonates with the atoms.

The final requirement is temporal stability. Instability in the system means that we must frequently spend time retuning and stabilizing it. The more stable things are the less time is wasted in that way. That being said, the absolute minimum required is that the laser remain stable over the time period we need for an experiment. This will vary depending on what we are doing, but for the purposes here we will assume a minimum of at least 10 min will be required to realistically get anything done. Since the laser will not be stable over this length of time by itself, we will need to use active feedback mechanisms to achieve long term stability. For the active feedback to work, however, we also need enough passive control over stability (e.g. by minimization of vibrations, temperature fluctuations and air movement, and any environmental variables that could affect the laser) to keep the laser from mode hopping, or changing faster than the active control can compensate for.

The requirements for our lasers are that they have a linewidth of less than 100MHz, be tunable to within 3GHz of the target wavelength and capable of continuous scanning over a range of 3GHz or more, be stable enough with passive control for active feedback to work, and be stable enough with active feedback to maintain a frequency lock for more than 10 minutes.

2.2 Temperature and Current Control

Transforming an instable multimode diode laser into a narrow linewidth stable and tunable laser requires a few different techniques. The first and simplest thing to do in order to stabilize the laser is to stabilize environmental and control variables. Specifically, we want the temperature and the current of the diode to be very consistent.

We have developed electronics in our lab based on this goal and, using very stable circuits [7] and active feedback for the temperature (see Section 2.5), are able to control these variables to a satisfactory level of stability. These electronics are the focus of others work and will not be discussed further here.

2.3 External Cavity Diode Laser

Temperature and current control help with stability. As a result the linewidth is narrowed as well, but not enough to bring it to an acceptable level. The best tool we have for that purpose is the external cavity diode laser (ECDL).

An ECDL is a laser diode with a secondary optical cavity feeding light back into it. This cavity is made by placing a diffraction grating directly in front of the diode. The angle of the grating is set so that the first order diffraction fringe from the grating feeds directly back into the diode (see Fig. 2.1). The angle of diffraction is dependent on frequency, so only light at a certain frequency is actually coupled back into the diode. Feedback of this nature causes one of the modes (whichever

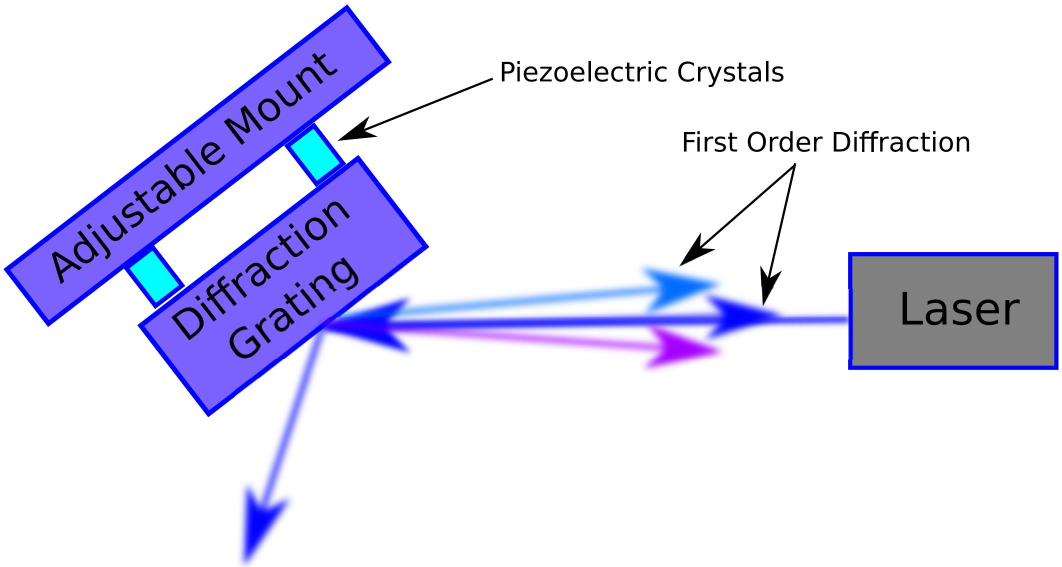

Figure 2.1 Light from the laser is reflected off a diffraction grating. Most of the light reflects as it would from a mirror, but some of it reflects back in the first order diffraction beam. The grating splits light according to it's wavelength, so only a very narrow range is fed directly back into the diode. This feedback causes one particular mode to overcome the rest, and the majority of the light comes out with a very narrow frequency linewidth. The angle of the grating is adjustable by external knobs and piezoelectric crystals for fine tuning).

is coupled back into the diode best) to be amplified more than the others. The result is that the laser settles into a single mode. In addition to narrowing the linewidth significantly, this also provides a way to control the frequency of the laser—with the grating angle.

The characteristics of an ECDL are ideal for our use. They have a very narrow linewidth, are affordable and—with some work—can be tuned.

The variables we can control that influence which frequency mode an ECDL operates in are the temperature, current, grating angle and cavity length. By manipulating these, we can tune a diode laser to the desired frequency. Unfortunately, the ideal parameters can be difficult to find since we have an overdetermined system with non-linear responses and interdependency among the input variables.

Each of the basic variables that determine laser gain can be influenced by various factors. External cavity length changes with grating angle adjustments and cavity length adjustments. The grating feedback is affected primarily by the grating angle. Temperature has an impact both on the internal cavity length and on the gain function of the lasing medium.

To gain some conceptual understanding of how the laser changes with respect to any one variable, we can think of each of the input variables as affecting one or more individual gain masks that dictate how much light of a given frequency will be attenuated. The overall gain of the system response is found by multiplying each of these masks together, as demonstrated in Fig. 2.2. The diode will tend to lase where the peaks of this net gain function are. That occurs wherever peaks from all the individual gain masks come closest to lining up.

Variables, such as temperature, tend to drift on a small scale. As it does so, the peaks of the net gain function also drift, and with them output wavelength of the laser. To make matters worse, the overall gain sometimes changes in such a way that the frequency the laser is currently at becomes attenuated, causing the laser frequency to jump to a higher peak of the combined gain function (often referred to as 'mode hopping').

Preventing the laser from mode hopping requires careful design and passive control of environmental variables. This includes use of optics tables on 'floating' mounts that reduce mechanical vibrations from the lab, as well as external enclosures to minimize air flow and temperature fluctuations.

Even with passive control methods, stabilized current and temperature, the frequency of an ECDL will still drift more than we can tolerate. Preventing the laser from doing so requires some form of active control. For any type of correctional control, however, we need some way to measure the error of whatever value we want to correct.

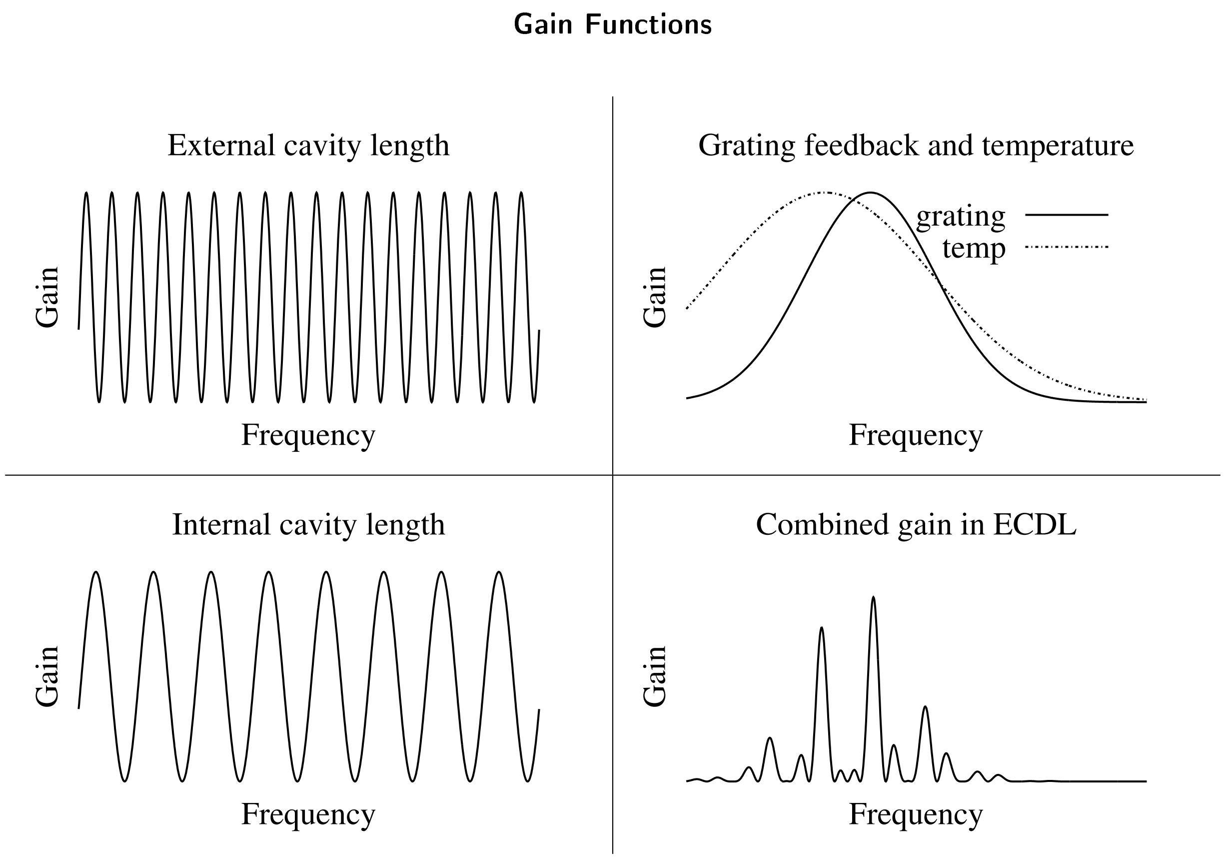

Figure 2.2 Gain functions. There are many variables that influence the overall gain of the diode-external cavity system. In this figure, the qualitative form a few gain masks (internal and external cavity length, grating feedback and temperature) are shown, along with their product—the effective net gain for the laser system. The frequency modes of the laser will be centered around the peaks in the combined gain function.

2.4 Error Signals

To keep lasers at the desired frequency, the first thing we need is some way to measure how far from the desired value the laser's frequency is. This turns out to be a non-trivial problem that has been solved by others in some very ingenious ways. Following are the methods that we use and the adaptations we have made.

The easiest way, conceptually, to measure error is to make an absolute measurement of the quantity to be controlled, and then take the difference with the desired value. However, differences tend to introduce a lot of error when the two magnitudes being compared are very close together. A more precise, and often simpler, method is to find some way to directly measure the difference itself. This is possible wherever there is a measurable effect that is dependent on the difference in question. For example, the absorption (A) of light by atom in a cell takes on the following form:

$$A \propto \frac{1}{1+4\delta_l^2/\Gamma^2},$$

where $\delta_l = \omega_0 - \omega_l =$ laser detuning, $\omega_0 =$ angular frequency of atomic transition, $\omega_l =$ angular frequency of laser, and $\Gamma =$ linewidth of the transition.

This is useful since transmission $(1-A)$, is readily measurable by shining a laser through a cell with an atom vapor and measuring the output on a photodiode. If the transmission is plotted vs. frequency, then there is a dip centered on v_0.

2.4.1 Deriving Error Signals from Absorption

The transmission function is symmetric about the point of interest (i.e. the center of the dip, where $\delta_l = 0$). This means that with just an absorption value, we can potentially tell how far off the frequency is but *not in which direction*. In addition, the light intensity may drift with time, so we can't depend on the magnitude of transmission to remain constant even when the frequency is not drifting.

All of the information we need can be gleaned from the absorption signal, but what we need to do is extract the slope of the signal. The way we do that electronically is to modulate the frequency of the laser (cause to oscillate slightly around a fixed value), with a modulation depth of only a fraction of the width of the absorption peak in question. The signal out will then be oscillatory with an amplitude proportional to the slope of the absorption vs. frequency curve. This amplitude is found by mixing the output signal with a copy of the modulating signal. Since the modulating signal and the output signal are 180 degrees out of phase when the slope is negative, and in phase when it is not, the output from mixing is negative when the slope is negative and positive when the slope is positive. The result, once the higher frequency components are thrown out (done via filters in the electronics), is the derivative of the absorption signal. This provides a signal that is approximately linear in the vicinity of the resonant frequency, and passes through zero at exactly the point in question. That signal is perfect for our needs.

To generate such a signal, we need some way to modulate the laser frequency. This is done by sending the laser through an acousto-optic modulator (AOM). The AOM adds a frequency offset to the optical light by bouncing it off acoustic wavefronts in a crystal, which is equal to the frequency of the acoustic waves. The frequency offset can then be modulated to add a modulation to laser frequency . See Appendix D for more details on the driving circuit for the AOM.

2.4.2 Saturated Absorption Spectroscopy

There are various problems with directly measuring the absorption in a vapor cell. The primary one that concerns us is that of Doppler broadening. Doppler broadening is an effect caused by the Doppler shift experienced by atoms moving towards or away from light. This motion causes each atom to sense a slightly different frequency of light. This, coupled with the random distribution of velocity, causes an overall broadening of the transmission dip with frequency that makes it unsuitable for obtaining a precise error signal. To overcome this limitation, we use a technique common in spectroscopy known as saturated absorption.

Saturated absorption spectroscopy works by aligning two counter propagating beams through a vapor cell so that the Doppler shift is opposite for the two. Because the beams are traveling in opposite directions, the Doppler shift due to movement of any particular atom is opposite for the two beams. For example, if the laser is red detuned (frequency is below resonance) then the light from one laser will be shifted into resonance for an atom traveling towards it, whereas that same atom will see the other laser shifted even farther from resonance. The result is the atom will only absorb light from the laser directed against its direction of motion. Because of this the two lasers will interact with different sets of atoms. One laser beam will only be absorbed by atoms moving in the +x direction and the other beam will only be absorbed by atoms moving in the -x direction, where x is the beam axis.

When the light is on resonance, however, that means that the atoms that it interacts with are those that have approximately zero velocity along the axis of the beams. In this case, both beams *are* interacting with the same set of atoms and, if the intensity is set right, the transition becomes saturated (meaning no more light can be absorbed). As a result, there is actually a spike in transmitted light in the center of the transmission dip. The width of this spike is much closer to the natural linewidth of the transition than the Doppler broadened transmission dip, and it is good enough to get an accurate error signal from.

2.4.3 Spectroscopy on Sr ions

Using saturated absorption spectroscopy to lock the ionization lasers will add in an additional complication when compared to locking the cooling laser. To look at absorption by Sr ions, we first need to produce ions in the cell. This will done by overlapping two beams (the cooling laser and a 405nm laser). These two lasers together are tuned to drive a two photon transition into an auto-ionizing energy state, producing ions along the trail where they overlap. The ion probe laser can then interact with the ions by overlapping it with the other two. The resulting setup is more complicated optically, due to the need for all three lasers to overlap along a single path through the cell, but this can be accomplished with dichroic mirrors (which reflect one wavelength while transmitting another) and polarizing beam splitters (which reflect light at one polarization and transmit light at another).

2.5 Proportional/Integral/Differential Control

The method that we use to control the ECDL parameters based on the error signal is proportional/integral/differential (PID) control. This method calculates a feedback signal based on three terms, each proportional to the error signal (ε), or the integral or derivative of the error signal:

$$\text{Feedback} = \alpha\varepsilon + \beta \int \varepsilon dt + \gamma \frac{d\varepsilon}{dt}.$$

The constants α, β and γ affect the response rate of the feedback, the precision to which it locks in the steady state, and the tendency towards oscillation. In practice we have a separate control knob (variable resistor) so we can tune each individually. We optimize the settings by tuning the proportional gain (α) first, with the others disconnected. This is done by turning it up until the system just begins to oscillate (which it typically will when the proportional gain is too large), then turning back down to about 80% of that value. Then we tune the integral and differential gains in a similar fashion, while leaving the proportional gain on. For a more information on this method of PID control see reference [8].

We use PID controllers for a lasers temperature, current and grating angle. The temperature is controlled separately from the others and held at a constant value. It's controller is slightly different because it does not include the differential gain (i.e. $\gamma = 0$). The control of current and grating angles is made more complicated because feedback to both is generated from a single controller.

The downside to this is that it is more difficult to tune them individually. One of the improvements we plan for future versions of the electronics is to separate the gain controls—for feedback to the grating and for feedback to the current—in order to better control the laser.

Chapter 3

Results

3.1 Ion probe

The ion probe is a relatively simple system, compared to the cooling laser. We have successfully tuned the laser to the correct frequency and scanned it over a range of about 5GHz. We won't be able to lock it to the ionic transition, however, until both the cooling laser and the other (405nm) ionization laser are ready.

To test the performance of the laser and the electro-optic modulators (a key component of the locking setup), we analyzed the spectrum of the laser in an optical cavity. The cavity has a free spectral range (FSR) of about .5GHz. As the cavity length is scanned the laser comes in and out of resonance and the spectrum repeats. Plotting the transmission through the cavity against the scanning signal gives a slice of the spectrum of the laser equal in frequency range to the FSR of the cavity. Figure 3.1 shows two such plots. The first is a scan of the laser alone. The second shows the spectrum of the laser with the electro-optic modulator (EOM) turned on. The EOM modulates the laser frequency by a few MHz, which is on the order of about 1 part in 10^8 of the overall frequency

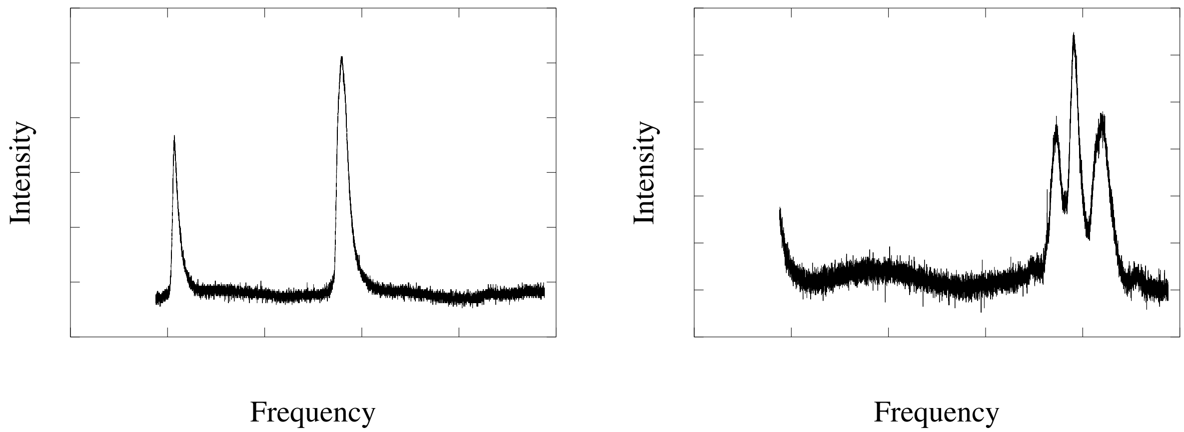

Figure 3.1 Spectra of the ion probe laser (units are arbitrary). These spectra were generated by plotting the transmission through an optical cavity. The first is of the laser by itself—the repeated peaks are due to scanning over more that one free spectral range of the cavity. In the second plot, the electro-optic modulator was turned on, and the resulting side bands are seen to the left and right of the central peak.

of the laser. In the spectrum of the laser, this shows up as two peaks a few MHz to the left and right of the central peak.

The laser fulfills at least two of the three requirements. It is a single mode with a narrow linewidth and tunable to within the amount we need. The linewidth of the laser is likely less than the finesse of the cavity allows to resolve, but we do know from the scans that it is at least narrow enough for our purposes. In addition, the laser is semi stable over a short period of time, but will not be as stable as it needs to be over long time periods precise until the full PID lock can be implemented.

3.2 Cooling Laser

The cooling laser (461nm), unfortunately, has presented more difficulty than the ion probe. The laser system works, and it is tunable to within the wavelength we need, but it is too unstable to lock

Current (A)	1st mode wavelength (cm^{-1})	2nd mode wavelength (cm^{-1})	Fluorescence
99.184	10848.72	10849.60	no
102.998	10848.62	10849.53	no
106.813	10848.54	10849.41	no
110.628	10848.48	10849.30	no
111.582	—	10849.27	yes
115.969	10848.35	10849.18	no
115.015	10848.35	10849.20	no

Table 3.1 Two of the frequencies at which the doubling crystal would produce light for a given current through the laser. Adjusting the laser into resonance involves checking each of the modes at which the laser lases with each adjustment of laser current or temperature.

to the Sr transition. Amplitude noise, mode instability and limited scanning range are the primary obstacles to establishing a good lock.

The problems with the laser are partly due to the number of variables involved, which make it difficult to tune the laser to run at its optimum setting. The process of doing so involves changing one setting slightly, then adjusting various others to maximize output, and then repeating the process.

We can start tuning the laser by roughly setting the grating angle and the current through the diode. Then we adjust the feedback (which goes to both) to bring the laser in resonance with one of the modes of the doubling cavity. These are spaced apart at regular intervals, but the laser can only reach two of these modes without making rougher adjustments to the current or grating angles. If neither of these modes is at the correct frequency, then we have to go back, make an adjustment, and then again find the frequencies at which we can lock the laser to the doubling cavity. This process is demonstrated in Table 3.1, which contains the data taken during one such iterative process. Even when the laser and cavity are tuned together to produce the correct frequency (at which the laser causes fluorescence), however, there are other variables which may not be very well optimized, including power output and scanning range.

Tuning the laser is tedious but can be done. The laser will lock to the doubling crystal at the correct frequency; however, we had difficulty getting a clean absorption signal from the vapor cell due to amplitude noise and mode instability. Scanning the frequency (by slowly manipulating the doubling cavity) does not cause the laser to completely lose it's lock, but it does cause a significant amount of fluctuation in the amplitude of the output blue light, which interferes with the transmission readings from the vapor cell.

When the laser is scanned, it will stay in a single mode for up to ~1GHz, but the amplitude oscillates erratically while it is being scanned. With this amplitude noise, it is very difficult to detect the absorption dips in the signal as we scan over the Sr resonant frequency. This is one of the primary reasons we've been unable to obtain a good error signal from the vapor cell to lock to the transition.

At best the scanning range is barely sufficient to do what we need, which is to scan over the width of the Doppler broadened transmission dip. This, along with the noise introduced in the amplitude as we scan the laser, makes it difficult to see the transmission drop at all when the laser is in resonance with atoms in the cell.

Despite the problems encountered, we have been able to achieve some important things with the cooling laser. It is tuned to the right wavelength, and it has a narrow enough linewidth. We were able to get the right wavelength of light, in a single mode, and cause fluorescence in the cell (see Fig.3.2). The obstacles that we have run into, however, indicate that we may need to consider an alternative method of generating the laser needed for cooling.

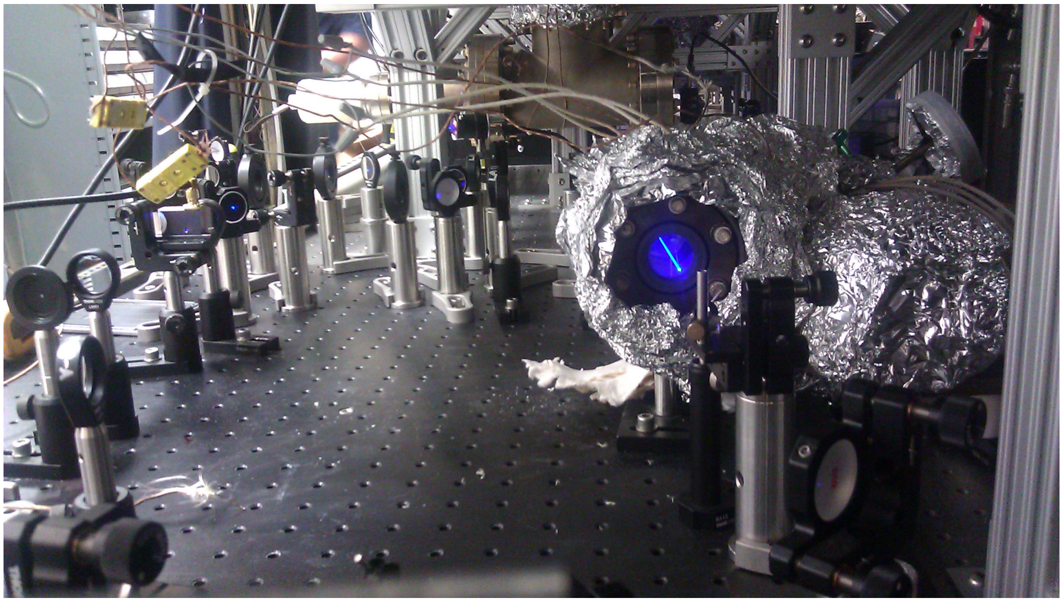

Figure 3.2 The cooling laser causes fluorescence along it's path through the vapor cell when it is tuned to the Sr resonant transition.

At the time we started this project, there were not any laser diodes available at the frequency that we needed for this laser, which is why we chose to use a doubling crystal. Currently, however, there are some commercial providers of diodes at the wavelength we need. Given the problems that we have with the doubled laser, it may be worth the time and investment to set up a new laser system that doesn't require the doubling stage.

3.3 Conclusion

Despite some of their inherent problems, we can make diode lasers work within the constraints needed for experiments in atomic physics. This is accomplished through the use of feedback control loops and external cavity optical feedback. When these methods are used, a diode can be controlled and tuned, with the required linewidth and mode stability necessary to probe and manipulate the inner structure of an atom.

The complexity of a laser system—and the corresponding instability—increases with the number of variables involved. As such we have learned that it is important to keep the system as simple as possible. While it can be difficult to maintain such a balance, a well designed and implemented diode laser system can enable us to study and manipulate atomic features on such a fine scale as to be truly incredible.

Appendix A

Wavelength

Wavelength is a measurement of the distance between peaks of a wave. Period is a measurement of the time it takes for one full wavelength to pass by a point. Frequency is defined as the inverse of the period, and wavenumber is defined as the inverse of wavelength. All of these variables are related by the velocity of the wave. Light has the unique characteristic of having a single universal value for velocity (when in a vacuum), commonly denoted as $c = 3 \times 10^8 m/s$. Because of this the variables mentioned above are all determined once one of them is given, and so they are all often used interchangeably. The relationships between frequency(f), wavelength(λ), period(T) and wavenumber(k) are summarized below:

$$T = \tfrac{1}{f}, \quad k = \tfrac{1}{\lambda}, \quad f = \tfrac{c}{\lambda}.$$

In addition the following is given as an approximate correspondence for small changes *at the wavelengths we are discussing* (461nm):

$$\Delta f = 3\text{GHz} \rightarrow \quad \Delta k \approx .1\text{cm}^{-1}, \quad \Delta\lambda \approx 2.13\text{pm}.$$

At first the use of wavelength, wavenumber, or frequency may seem arbitrary. In some cases it is, but often there are practical or historical reasons for doing so. In particular, when small differences are being compared, frequency is generally used in place of wavelength. This is partly because units of gigahertz are more familiar than picometers, and partly because these differences are usually directly measured electronic frequencies, so experimentalists just stick to the units of frequency they are working with.

The wavenumber historically has been used in spectroscopy, and so is often used when referring to the wavelength of an atomic transition. In part, this is because data on such things is usually given in terms of wavenumber and because wavenumber is proportional to energy.

Appendix B

Laser Spectra

Light is composed of electromagnetic waves. The key defining characteristics of a pure sine wave are its frequency and wavelength. If a wave (electromagnetic or otherwise) is a pure sine wave, then these are very well defined. If it is not, however, it can be still be expressed as the sum of sine waves at different frequencies. When the amplitudes of the various frequency components that make up the wave are plotted, we call this a spectrum. The closer a wave is to perfect sine wave, the narrower its spectrum will be. See figure B.1.

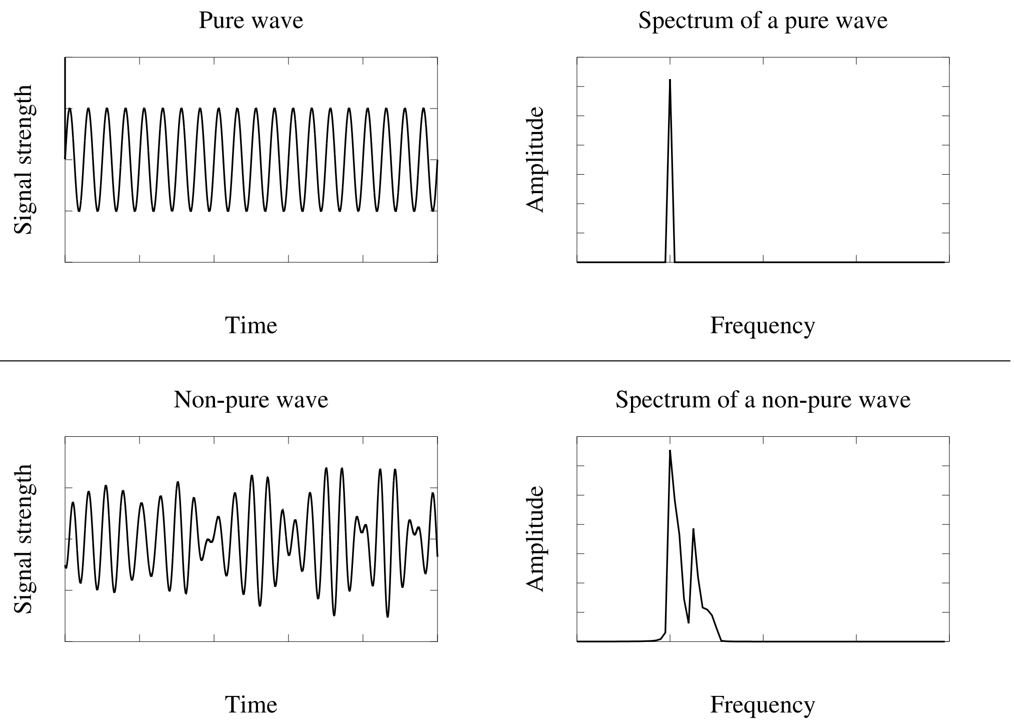

Figure B.1 The above illustrations demonstrate the relationship between an oscillating function (i.e. wave) and its frequency spectrum. The top graphic shows a pure sine wave. Note that the corresponding spectrum is a single narrow peak. A single mode laser would have a similar spectrum. The bottom function is a less simple oscillating function. Note that its corresponding spectrum is spread out over a small range of frequencies, with more than one peak. This is more representative of what the spectrum of a multi-mode laser could look like.

Appendix C

Laser Locking Methods

Any technique that can measure a lasers frequency directly, or relative to some reference can be used to derive an error signal for PID control. The methods that I am familiar with are briefly introduced below, with a corresponding summary in table C.1.

Cavity Scan

This method works by comparing the transmission peak locations of two separate lasers in an optical cavity as the cavity length is scanned. If the frequency of the two lasers moves relative to each other, then the relative distance of the peaks will move by a proportional amount. Software analysis can normalize the data, adjust for scanning differences, and then derive an error signal based on movement of the peaks relative to each other.

Stabilized Cavity

The stabilized cavity method is similar to the cavity scan method, in that both lasers are overlapped in an optical cavity. In this method, however, the length of the cavity is locked (using feedback to piezoelectric crystals) to the stable reference laser. The second laser is then locked to the cavity. This method is likely to be more expensive to implement, as it requires modulators for both lasers as well as equipment, electronics, etc for two PID locks. The bandwidth of the lock is limited to the mechanical response of the cavity, which is generally in the tens of kHz.

Method	Requires Ref. Laser	Lock Bandwidth	Relative Cost	Notes
Cavity Scan	yes	10 kHz	Low	Requires computer to process data
Stabilized Cavity	yes	100KHz	High	Bandwidth limited to mechanical response of cavity piezo crystals
Offset Reference	yes	100MHz	Low	Reference laser must be within 5GHz of target frequency
Saturated Absorption Spectroscopy	no	100Mhz	High	

Table C.1 Comparison of the pros and cons of various common locking techniques, including bandwidth, whether or not a stable reference is necessary, and the relative cost of all the needed equipment.

Offset Reference

This method works by directly comparing the frequency of one laser to another. When the lasers are overlapped on a photodiode, the electric fields mix such that the intensity oscillates at frequencies equal to the sum and difference of the two lasers' frequencies. Typically, these would both be much too high for any electronic equipment to detect. However, if the second laser doesn't need to be further than a few GHz from the first, then the difference frequency is detectable, and an error signal can be derived directly from it.

Saturated Absorbtion Spectroscopy

Saturated absorption is discussed in Section 2.4.2

Appendix D

Acousto and Electro-Optic Modulator Oscillator Circuit

The oscillator circuit was developed to drive acousto or electro-optic modulators (AOM and EOM, respectively). For the AOM we need a modulated frequency signal, this adds a frequency offset to the laser in addition to modulating it's frequency. The EOM just needs a simple unmodulated signal, and modulates the laser without adding a frequency offset. The purpose of the oscillator circuit is to provide the driving signals for these components in the locking setup.

The circuit consists of an oscillating op-amp and a filter. The op-amp generates a signal that is approximately a square wave. Since a square wave is composed of a sum of sine waves at odd multiples of the fundamental frequency, it can be turned into a sine wave by filtering out all the frequencies higher than the fundamental one. The circuit shown below does essentially this. The square wave is filtered, then amplified before being sent to a linear power amplifier to generate a signal capable of driving the AOMs.

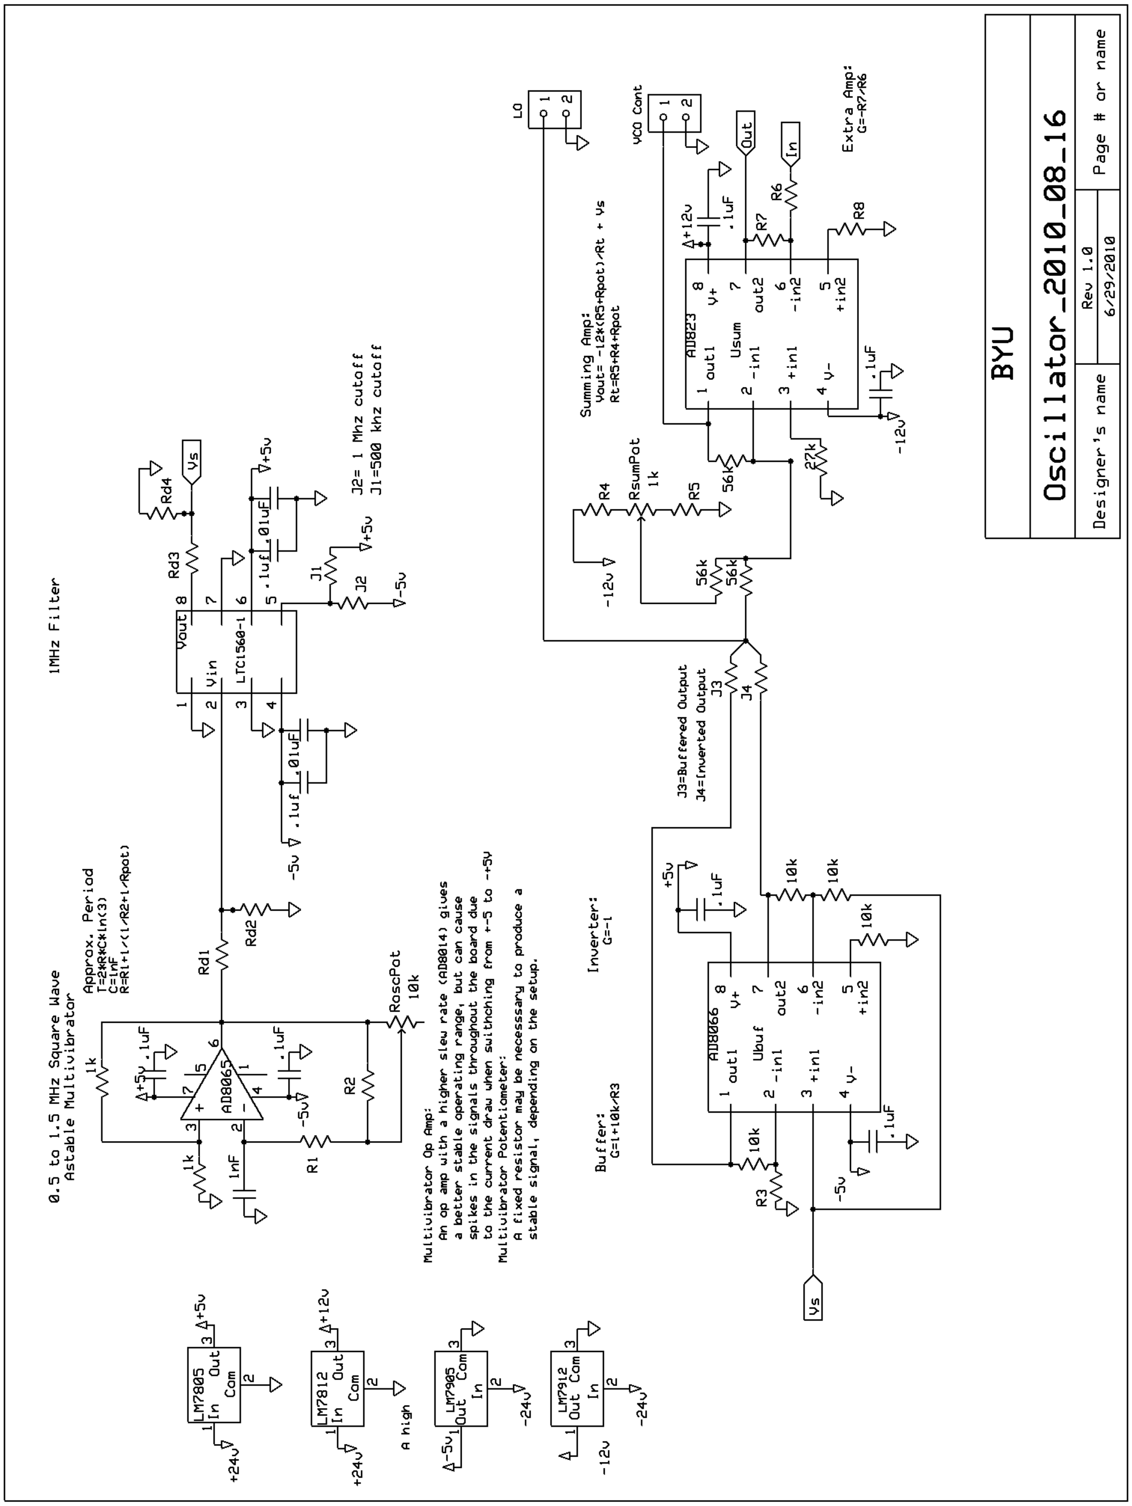

Figure D.1 Oscillator circuit for AOM driver

Bibliography

[1] J. B. Fixler, G. T. Foster, J. M. McGuirk, and M. A. Kasevich, "Atom Interferometer Measurement of the Newtonian Constant of Gravity," Science **315,** 74–77 (2007).

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

[3] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of Neutral Sodium Atoms with Radiation Pressure," Phys. Rev. Lett. **59,** 2631–2634 (1987).

[4] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard, "Optics and interferometry with atoms and molecules," Rev. Mod. Phys. **81,** 1051–1129 (2009).

[5] B. Neyenhuis, D. Christensen, and D. S. Durfee, "Testing Nonclassical Theories of Electromagnetism with Ion Interferometry," Phys. Rev. Lett. **99,** 200401 (2007).

[6] A. Kramida, Y. Ralchenko, J. Reader, and N. A. Team, "NIST Atomic Spectra Database (version 5.0)," Available: http:\\ physics. Nist. gov/asd. National Institute of Standards and Technology, Gaithersburg, MD. (2012).

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

[8] M. Zhuang and D. Atherton, "Automatic tuning of optimum PID controllers," Control Theory and Applications, IEE Proceedings D **140,** 216–224 (1993).

Index