

Laser frequency stabilization systems for laser  
cooling of trapped barium ions - senior thesis

Greg McGlynn  
Northwestern University

May 3, 2011

## **Abstract**

As part of Brian Odom's molecular ion trapping group, I have investigated the use of several methods of laser frequency stabilization for use in the laser cooling of trapped barium ions. The methods discussed are saturated absorption spectroscopy, the dichroic atomic vapor laser lock, and the temperature stabilized optical cavity. I review the basic principles of operation of each of these methods of laser frequency stabilization and describe implementations of each. I present data characterizing these implementations and discuss the successes and challenges encountered with each.

I thank Joan Marler and Yen-Wei Lin for a great deal of help throughout these projects.

# Contents

<b>Introduction</b>	<b>3</b>
1.1 Motivation and goals . . . . .	3
1.2 Ion trapping and cooling . . . . .	3
1.3 Laser frequency stabilization . . . . .	5
<b>Saturated absorption spectroscopy</b>	<b>8</b>
2.1 Principles . . . . .	8
2.1.1 Doppler broadening . . . . .	8
2.1.2 Saturated absorption . . . . .	10
2.2 Implementation . . . . .	13
2.3 Error signal and lock stability . . . . .	14
2.4 Challenges and conclusions . . . . .	17
<b>Dichroic atomic vapor laser lock</b>	<b>18</b>
3.1 Principles of operation . . . . .	18
3.1.1 The Zeeman effect and photon polarization . . . . .	18
3.1.2 Dichroism and the error signal . . . . .	19
3.2 Preliminary measurements . . . . .	21
3.2.1 Optogalvanic measurements . . . . .	22

3.2.2	Absorption measurements . . . . .	23
3.3	Implementation and characterization . . . . .	24
3.4	Conclusions . . . . .	29
<b>Temperature-stabilized optical cavity</b>		<b>30</b>
4.1	Principles . . . . .	31
4.1.1	Optics . . . . .	31
4.1.2	The influence of temperature . . . . .	34
4.2	Implementation . . . . .	35
4.2.1	Construction . . . . .	35
4.2.2	Alignment procedure . . . . .	36
4.3	Characterization . . . . .	38
4.3.1	Temperature stability . . . . .	38
4.3.2	Measurements of the cavity . . . . .	43
4.4	Current status . . . . .	47
<b>Summary</b>		<b>48</b>

# Introduction

## 1.1 Motivation and goals

The molecular ion trapping group at Northwestern University, under Brian Odom, is working toward cooling and trapping of molecular ions. Trapped molecular ions can be used to probe important questions such as whether left- and right-handed chiral molecules have observably different properties or whether physical “constants” like the electron-proton mass ratio actually exhibit time-variation.

Direct laser cooling of molecular ions is difficult, but they can also be cooled sympathetically by cooling the atomic ions in a gas mixture of atomic and molecular ions [4]. Toward this end, the group has trapped and cooled barium ions that can be used to sympathetically cool molecular ions. This thesis describes work on frequency stabilization systems for the lasers used in cooling these barium ions. In the introduction I describe the Doppler cooling method and explain the necessity for frequency stabilization of lasers. In the rest of the thesis I describe the principles and implementations of three methods of laser frequency stabilization.

## 1.2 Ion trapping and cooling

While electrically charged particles cannot be confined by static electric fields, they can be confined by oscillating electric fields. We trap barium ions with a Paul trap, in which both DC and radio-frequency AC voltages are applied to carefully arranged electrodes to set up an effective confining potential for




Figure 1.1: Each dot in these images is a cold trapped barium ion. The Coloumb repulsion of the ions causes them to naturally form into ordered "crystals"

the ions.

Ions confined in the trap must then be cooled. Ions can be cooled by lasers through the process of Doppler cooling. In this technique, a laser is tuned to a frequency just below an electronic resonance of the ion. An ion moving toward the laser perceives it to be blueshifted to a higher frequency and thus can absorb a photon. The ion is briefly excited to a higher energy state. It is also slowed down slightly because the photon carries momentum. After a brief interval the excited state decays, emitting a photon in a random direction, and the ion receives a slight kickback because of the emitted photon's momentum. On average such an interaction causes the ion to slow down. In this way the kinetic energy of the ions' motion—that is, the thermal energy of the gas of ions—is drained away. We cool barium in this way using laser light of about 493 nm, which is resonant with the transition from the  $6S_{1/2}$  state (the ground state) to the  $6P_{1/2}$  state. In this way the ions can reach quite low temperatures, settling into organized "crystals" in the center of the ion trap. Some crystals that have been trapped and imaged by the group are shown in Figure 1.1.

One hitch in the laser cooling process is that the excited state produced by absorption of a laser photon may not decay back to the ground state, but rather to some other excited state that is metastable and takes annoyingly long to decay to the ground state. This problem can be solved with

a “repump” laser which pushes ions in the metastable state back up into the original excited state. In Ba+, the  $6P_{1/2}$  state can decay to the  $5D_{3/2}$  state instead of the ground state; the corresponding light has a wavelength of about 650 nm. Accordingly we employ a repump laser at this wavelength in addition to the 493 nm cooling laser.

### 1.3 Laser frequency stabilization

For the cooling techniques described above to work, the frequency of the cooling lasers must be tuned very accurately relative to the transition frequencies of the trapped ions. Stability of the laser frequency of around a MHz over many hours is desired. Achieving this level of stability for a laser requires a dedicated feedback system that measures the laser frequency with the desired precision and stability and tunes the laser frequency up or down to keep it at the desired number. Our cooling lasers are diode lasers, whose frequency can be adjusted continuously by varying the current going to the laser diode.

Currently the lasers are stabilized by measuring their frequency with a wavemeter<sup>1</sup>. This device can only measure one laser at a time, so it switches between measuring the different lasers a few times per second using an optical switch<sup>2</sup>. This system has several drawbacks, including poor day-to-day stability, which requires frequently readjusting the laser frequencies manually, and slow response time of the lock since the frequency of any laser is only read a few times per second. Therefore we want dedicated frequency lock systems for the lasers.

In general a stabilization system measures the laser light to produce an “error signal” with the general features shown in Figure 1.2. The error signal is a function of frequency that has a zero crossing at the desired frequency of the laser. The error signal is fed into a controller (for example, a PID controller), which adjusts the current to the diode to cause the error signal to go to zero, stabilizing, or “locking” the laser at the frequency of the zero crossing of the error signal.

---

<sup>1</sup>HighFinesse WS-U wavelength meter

<sup>2</sup>From LIGHTech Fiberoptics, Inc.




Figure 1.2: A schematic error signal. Two numbers of interest are indicated. One is the "recapture range:" the extent of the nonzero wings to either side of the zero crossing, which determines how large a frequency jump the lock can tolerate and still bring the frequency back to the set point. The other is the slope of the zero crossing, which determines how precisely defined the zero crossing is in the presence of noise and thus the locking accuracy that is possible (a large slope at the zero crossing is desired). It is also important that the location of the zero crossing be tunable, so that the locked laser can be set to the right frequency.

This thesis deals with three different ways of obtaining an error signal for laser locking: saturated absorption spectroscopy, the dichroic atomic vapor laser lock (DAVLL), and the temperature-stabilized optical cavity. I review the basic principles of operation of each of these methods of laser frequency stabilization and describe implementations of each. I present data characterizing these implementations and discuss the successes and challenges encountered with each.

# Saturated absorption spectroscopy

Precision spectroscopy can be used to provide a very stable frequency reference for a laser. The laser can be stabilized at the frequency of a specific transition of some atom or molecule, which provides a fixed and unchanging lock point for the laser frequency. If the frequency of the transition is not exactly the desired frequency, the laser can be stabilized at a fixed offset from the transition frequency using a device like an acoustic-optic modulator to shift the frequency of the portion of the beam that is used in the lock.

We investigated saturated absorption spectroscopy of iodine gas ( $I_2$ ) as a way of stabilizing a 650 nm laser for cooling of  $Ba^+$ . Iodine has several hyperfine transitions within a few hundred MHz of the desired frequency for laser cooling. Figure 2.1 shows the iodine cell used.

## 2.1 Principles

### 2.1.1 Doppler broadening

In a gas at room temperature, transition frequencies are spread out by the random redshifts resulting from the thermal motion of molecules. If a laser beam passes through the gas, the velocities of the molecules it encounters will have random components along the beam axis distributed in a Gaussian of the form

$$\exp(-\frac{1}{2}mv^2/kT)$$



Figure 2.1: This cell containing  $I_2$  gas was used in the saturated absorption spectroscopy setup. (Thorlabs CQ19100-I)

Where  $v$  is the velocity along the beam axis,  $m$  is the mass of a molecule,  $k$  is the Boltzmann constant and  $T$  is the temperature of the gas. This velocity distribution has a full-width at half-maximum (FWHM) of

$$v_{FWHM} = \sqrt{(8 \ln 2)kT/m}$$

For  $I_2$  gas ( $m = 4.21 \times 10^{-25}$  kg) at room temperature ( $T = 298$  K), the FWHM of the line-of-sight velocity distribution is about 230 m/s. These molecular velocities along the beam direction result, in the molecules' frames of reference, in red- or blue-shifting of the laser frequency. For nonrelativistic velocities the frequency shift of light of frequency  $f$  from a line-of-sight velocity  $v$  can be approximated as

$$\Delta f = \frac{v}{c}f$$

Because of this effect, even when the laser frequency (in the lab frame) is not exactly the same as the transition frequency, there will be some molecules in the gas with the right velocity to be resonant with the laser. As the laser frequency is scanned past the transition, the laser interacts with molecules of various velocity classes, and the amount of absorption at a given frequency is proportional to the number of molecules with the correct velocity along the beam to be excited by the that frequency. Accordingly, the absorption spectrum of the laser will not peak sharply at the transition frequency but

will have a Gaussian shape with FWHM

$$f_{FWHM} = \frac{1}{c} \sqrt{(8 \ln 2)kT/m} \cdot f = f \sqrt{(8 \ln 2)kT/mc^2}$$

For 650 nm ( $f = 461$  THz) light, the FWHM of the I<sub>2</sub> absorption peak due to the Doppler effect at room temperature is about 350 MHz. When the desired locking accuracy is on the order of 1 MHz, these Doppler-broadened absorption features seem unsuitable as a frequency reference. While DAVLL, an approach discussed later, can lock to the desired precision using Doppler-broadened absorption signals, the approach of saturated absorption spectroscopy is to eliminate the Doppler broadening. With the effects of Doppler-broadening removed, individual hyperfine features of the absorption spectra can be resolved, with widths of a few MHz.

## 2.1.2 Saturated absorption

The optical setup for the saturated absorption technique is shown in Figure 2.2. The laser is split into two counter-propagating beams, called the pump and probe beams. Generally the probe beam is much weaker than the pump beam. The probe beam is passed through the gas and the emerging intensity is measured at a photodetector. If there is no pump beam, then the photodetector will see a Doppler-broadened absorption feature as the laser frequency is scanned across the transition.

The pump beam passes through the gas in the opposite direction and is made to overlap with the probe beam inside the gas. Doppler broadening is defeated by exploiting the interaction between the probe and pump beams through the molecules of the gas.

Let the frequency of the laser be  $f$  and the frequency of the transition be  $f_0$ . Because the pump and probe beams travel in opposite directions, in general they interact with different velocity classes of molecules. In the figure, if  $v$  is positive in the direction of the photodetector, then the pump beam interacts with molecules with line-of-sight velocity  $v$  satisfying

$$(1 + v/c)f = f_0$$

While the probe beam is travelling the other direction and so interacts with atoms satisfying

$$(1 - v/c)f = f_0$$




Figure 2.2: The saturated absorption spectroscopy optical setup described in the text.

A key element of saturated absorption spectroscopy is that the pump beam is made intense enough to deplete the ground state population of the molecules that it interacts with. In the velocity class of molecules with the correct velocity to interact with the pump beam, a large number of them are excited by the beam so that a smaller-than-normal proportion of them are in the ground state. This means that a different beam that tried to interact with these atoms would find fewer of them than normal to interact with, and so would suffer less absorption than normal.

Figure 2.3, then, shows the interaction between the probe and pump beams by looking at the ground state population of different velocity classes of molecules in the gas. The pump beam burns a hole in the distribution at the velocity class it interacts with, given its frequency. If  $f \neq f_0$ , that is, if the laser frequency does not equal the transition frequency, then the probe beam is interacting with a different velocity class and so does not notice this effect. However, when  $f = f_0$ , that is, when the laser frequency is right at the transition frequency, both probe and pump beams interact with the group of molecules with zero velocity along the beam axis. Because of the hole burning, the probe beam can excite fewer molecules than normal in its velocity class, and so suffers less absorption than normal. But this effect only happens in a very small frequency range around the transition frequency.

This is the effect that allows saturated absorption spectroscopy to achieve very sharp frequency resolution. If the laser frequency is scanned, a normal Doppler-broadened peak is observed, but with a sharp dip at the central transition frequency. This dip can be used to produce a steep error signal for frequency locking in the following way: take the derivative of the absorption

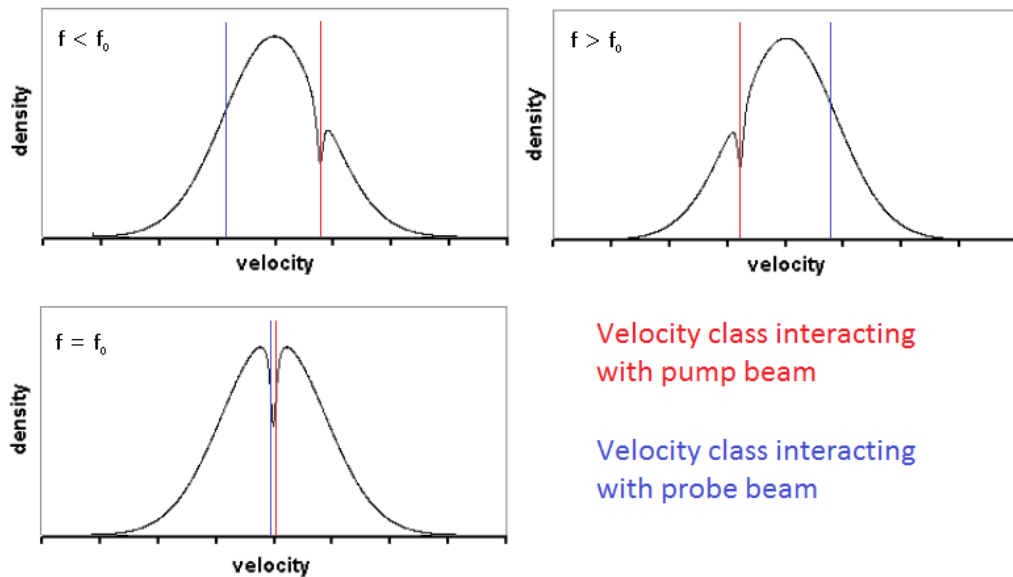


Figure 2.3: The Gaussians represent the velocity distribution of ground state molecules in the gas. Vertical lines show the velocity classes that are interacting with the probe or pump beams for various relationships between the laser frequency and the transition frequency. The pump beam "burns a hole" in the velocity class it interacts with, exciting many of these molecules out of the ground state.




Figure 2.4: The inclusion of an AOM into the saturated absorption setup, for frequency modulation.

signal by jittering the laser frequency slightly and observing the variation in light falling on the photodetector. This derivative of absorption with respect to frequency has a steep zero-crossing at the minimum of the dip, and can be used as the error signal for frequency locking.

## 2.2 Implementation

We implemented a saturation absorption spectroscopy setup as shown in Figure 2.2 with the addition of an acoustic-optic modulator<sup>3</sup> (AOM) that was used to modulate the frequency of the pump beam to get the derivative signal; see Figure 2.4.

This is slightly different from the situation described above, where both probe and pump beams were always at the same frequency. However, it achieves basically the same result. If the probe beam has frequency  $f$  and the pump beam has frequency  $f + \Delta f$  in the lab frame, then there is still some velocity class of molecules that sees them both at the same frequency by redshifting one and blueshifting the other. In the frame of this velocity class the frequencies are shifted toward each other by  $\Delta f/2$  and both beams have frequency  $f + \Delta f/2$ . If  $f$  is set properly, then varying  $\Delta f$  with the AOM will scan  $f + \Delta f/2$  over the transition, giving the same sharp signal as above when the special velocity class becomes resonant with both beams.

<sup>3</sup>Brimrose model TEM-85-2

One concern with this method might be whether this velocity class actually has any members—if the special velocity class is too far out of the room temperature velocity distribution we cannot observe the effect we are looking for. Our AOM produced a  $\Delta f$  of about 85 MHz, well inside the Doppler width calculated above for  $I_2$  at room temperature, so this concern is not problematic.

A slight practical annoyance is that the AOM necessarily deflects the beam that emerges from it by an angle proportional to  $\Delta f$ . In an acoustic-optic modulator, laser light passes through a crystal and interacts with sound waves of some wavelength  $\Lambda$  and frequency  $f_{AOM}$ . Several distinct diffracted beams emerge from the AOM, characterized the diffraction order  $m$ , an integer; for our purposes only the first diffracted order was useful ( $m = \pm 1$ , depending on whether the AOM is being used to upshift or downshift the frequency). Each emerging beam has its frequency shifted from the original laser frequency by  $\Delta f = mf_{AOM}$  and is deflected from the original beam axis by an angle  $\theta$  satisfying

$$\sin \theta = \frac{m\lambda}{2\Lambda}$$

where  $\Lambda$  is the wavelength of the laser light in the crystal.  $\frac{1}{\Lambda} \propto f_{AOM} \propto \Delta f$  and for small angles  $\sin \theta \approx \theta$ , so we have  $\theta \propto \Delta f$  as asserted above.

As a result, when jittering  $\Delta f$  to obtain the derivative signal, the pump beam moves around and the overlap between the probe and pump beams is affected. The change in the beam angle is fairly small, so that we were able to adjust the beams so that there was some overlap at all times as  $\Delta f$  was varied. It is possible to completely eliminate this worry by doing a double-pass through the AOM, so that the two deflections cancel out. Doing this shifts the frequency by twice as much, however, and this might cause problems due to the concern mentioned above about whether this would shift the special velocity class too far outside the room temperature velocity distribution.

## 2.3 Error signal and lock stability

Figure 2.5 shows the derivative signal obtained by the above procedure. By varying the frequency of the AOM, the pump beam frequency was modulated




Figure 2.5: A plot of the error signal. On the x axis is the wavelength according to the WaveMeter; on the y axis is the measured error signal from the lock-in. Each squiggle is an individual hyperfine transition of  $I_2$ ; the laser can be locked to any of these zero crossings. The x axis gives the laser wavelength in nanometers; at 650 nm a variation in wavelength of  $10^{-6}$  nm is a variation in frequency of 0.7 MHz. Accordingly each squiggle is about 15 MHz wide.

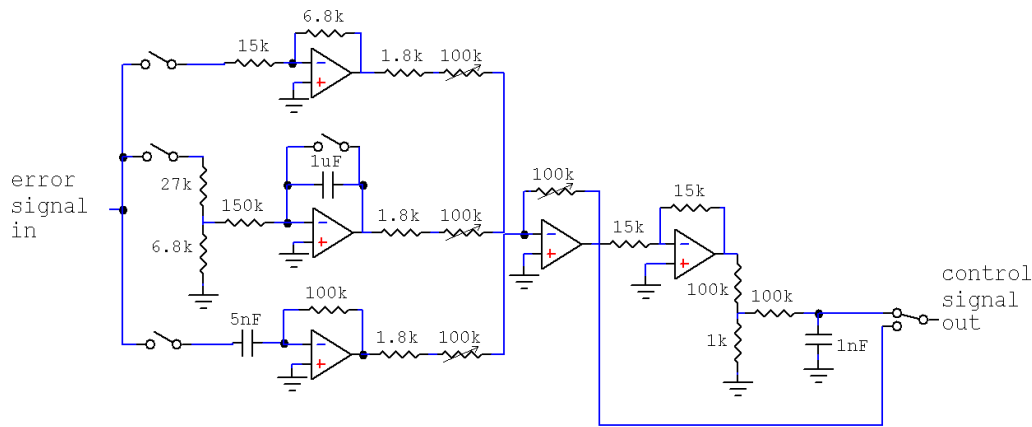


Figure 2.6: The circuit diagram of the PID controller used to stabilize the laser using the error signal from saturated absorption spectroscopy. Proportional, integral, and derivative terms of the error signal are computed by the first three op-amps, then summed and optionally inverted by the succeeding two. The variable resistors allow the tuning of the P, I, and D gains as well as the overall gain. Three switches allow the P, I, or D terms to be turned on or off. In practice only P and I were used. Another switch allows the I term to be zeroed. The last switch controls the sign of the output signal.

in a sine wave at 50 kHz; the amplitude of the modulation of the pump beam frequency was a few MHz. The error signal was found by measuring the resulting 50 kHz variation in detector photocurrent using a lock-in amplifier<sup>4</sup>.

Using this error signal, the laser frequency was locked by varying the laser diode current in order to keep the error signal at a zero crossing. A proportional-integral feedback loop was constructed to do this automatically. At first this was done with a Labview program to allow easy experimentation with the proportional and integral gains. However it was desired to have a separate electronics box to handle the feedback. Therefore, once good gains were found an analog PID circuit was constructed and used to lock the laser; its circuit diagram is shown in Figure 2.6.

The short-term stability of the saturated absorption lock was assessed using

<sup>4</sup>Stanford Research Systems model SR530

the wavemeter. Over a period of a few minutes the variation in frequency was less than or equal to about 1 MHz. The wavemeter's accuracy is not much better than this, so it was not possible to determine exactly how precisely the laser was locked.

To allow the set point of the lock to be varied, another AOM was added to the setup. Before arriving at the setup depicted in Figure 2.4, the light used for locking made a double-pass through this extra AOM. Varying the frequency shift produced by this AOM varied the location of the zero crossings of the error signal and thus the set point of the lock.

## 2.4 Challenges and conclusions

While above I have reported a degree of success in implementing a laser frequency lock based on saturated absorption of iodine, a few shortcomings have prevented this setup from being put into use.

First, the technique requires an expensive lock-in amplifier, currently the only one in the lab and needed periodically for other measurements. Some progress has been made by others in the lab toward a home-made lock-in amplifier that could be dedicated to this setup.

Second, the lock was unable to function properly when the RF power supply used for ion trapping was turned on. Presumably some part of the lock electronics was excessively influenced by the interference put out by the RF.

Finally, the above results were achieved using a rather large amount of laser power—over 10 milliwatts. Ideally the lock would use much less power so that as much as possible is available for the real purpose of the laser, which is ion cooling. However it was found to be much easier to obtain a good, clean, steep error signal with more laser power. Likely a significant reduction in laser power usage could be achieved by careful alignment of the AOMs to maximize diffraction efficiency and careful tuning of the rest of the optical and electrical setup to make the error signal as nice as possible even with a smaller laser power.

If the above challenges are overcome, the iodine saturated absorption lock seems well suited to locking the 650 nm cooling laser for the barium ion trap.

# Dichroic atomic vapor laser lock

The dichroic atomic vapor laser lock (DAVLL) is another method of frequency stabilization based on spectroscopy and is described in [2]. It has several attractive features. Its setup is very simple. It requires no modulation of the laser frequency, nor does it involve chopping the beam and recovery of a tiny signal with a lock-in amplifier. Further, its error signal has extremely long wings of hundreds of MHz, so that even if the frequency accidentally jumps the lock can be recovered over a large range.

There are a few disadvantages. DAVLL is sensitive to small and possibly poorly-understood temperature-induced variations in the optical behavior of the components used obtaining good stability may require putting part or all of the setup in a temperature-stabilized container. In addition, the slope of its error signal can be smaller than that of other techniques, since it does not eliminate Doppler-broadening of absorption signals.

## 3.1 Principles of operation

### 3.1.1 The Zeeman effect and photon polarization

The DAVLL makes use of the Zeeman effect: the splitting of electronic energy levels in a magnetic field. Neglecting hyperfine structure, an atomic energy level with total angular momentum  $J = L + S$  is  $(2J + 1)$ -fold degenerate. There are states of equal energy with values of the quantum number  $m_j = -J, -J + 1, \dots, J - 1, J$ . In a magnetic field this degeneracy is broken and these

$2J + 1$  states acquire slightly different offsets from the original energy. In a hydrogen-like atom the splittings are

$$\Delta E = \mu_B B m_j g_j$$

where  $\mu_B$  is the Bohr magneton and  $g_j$  is the Lande g-factor, which is a function of the value of  $L$  and  $S$  of the shell.

Next, the photon has spin 1, so if an atom absorbs a photon that is incident along the quantization axis it makes a transition for which  $\Delta m_j = \pm 1$ , depending on which way the photon angular momentum points. The direction of the angular momentum of the photon is given by its polarization. If the photon is right circularly polarized (RCP),  $\Delta m_j = +1$ , while if the photon is left circularly polarized (LCP),  $\Delta m_j = -1$ .

### 3.1.2 Dichroism and the error signal

The above facts imply that right and left circularly polarized photons are absorbed differently in a magnetic field; the different absorption of different polarizations is called “dichroism”. Consider only photons incident along the direction of the magnetic field. Recall that a magnetic field causes a shift in the energy of a state proportional to  $m_j$ . Then transitions with  $\Delta m_j = +1$  involve a larger energy change in a magnetic field than in free space, and shift to higher frequencies. On the other hand, transitions with  $\Delta m_j = -1$  involve a smaller energy change when the magnetic field is applied, and shift to lower frequencies. The effect of this is that at frequencies above the central transition frequency, RCP light has a higher absorption than LCP light, while the opposite is true at frequencies below the central transition frequency.

This effect can be used to construct an error signal for laser frequency stabilization. Taking the difference of the RCP and LCP absorption signals produces a function of frequency with a zero-crossing near the central frequency of the transition and with wide wings to either side (because these signals are Doppler-broadened). The optical setup that accomplishes this measurement is shown in Figure 3.1.

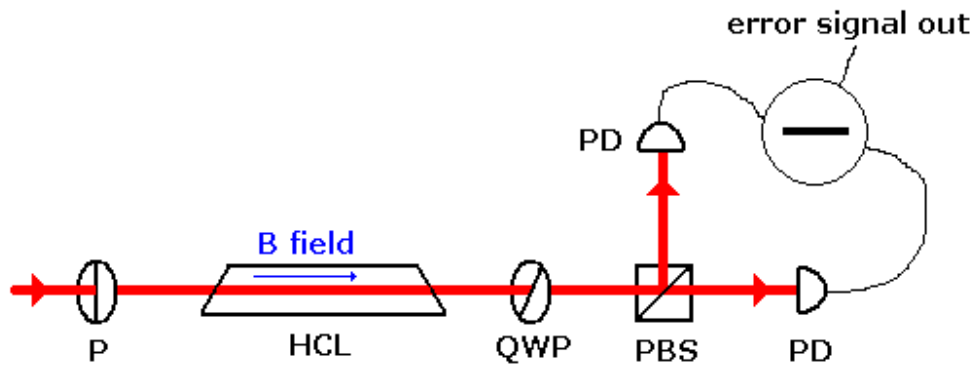


Figure 3.1: The optical setup for the dichroic atomic vapor laser lock. The vertical linear polarizer (P) produces an equal linear combination of LCP and RCP light. The hollow cathode lamp (HCL), described below, produces  $Ba^+$  ions. A magnetic field is applied so that these ions absorb different amounts of each circular polarization. The quarter wave plate (QWP) is set at 45 degrees to the linear polarizer, converting LCP and RCP light to horizontally and vertically polarized light. These polarizations are separated by the polarizing beamsplitter (PBS). The two photodiodes (PD) produce photocurrents that are differenced to get the differential absorption of the two polarizations.



Figure 3.2: The hollow cathode lamp used in the DAVLL setup. Shown also are two of the six ring magnets used to provide the B field.

## 3.2 Preliminary measurements

The use of DAVLL for the 650 nm transition of the Ba<sup>+</sup> ion was explored. To get the necessary gas of Ba<sup>+</sup> ions a hollow cathode lamp<sup>5</sup> (HCL), was employed and is shown in Figure 3.2. Also shown are the magnets<sup>6</sup> used to create the necessary axial magnetic field. In this device, a large voltage between an anode and a cathode is used to create an electrical discharge in a buffer gas. Ions in the resulting plasma accelerate toward the cathode and sputter off atoms from the cathode, which contains barium. Some of these barium atoms are ionized, and some ions become excited to higher energy states, including the  $5D_{3/2}$  state relevant to the 650 nm transition. The geometry of the electrodes is designed to allow a laser to pass through the plasma containing the ions of interest.

---

<sup>5</sup>Hamamatsu model L2783-56NE-BA

<sup>6</sup>McMaster part 5856K14




Figure 3.3: The signal from the 650 nm transition obtained via optogalvanic spectroscopy. The HCL discharge current was 15 mA. The width in frequency of the optogalvanic peak is about 1 GHz.

### 3.2.1 Optogalvanic measurements

First, measurements of the optogalvanic effect in the discharge were conducted to verify the presence of Ba<sup>+</sup> ions in the  $5D_{3/2}$  state. In the optogalvanic effect, laser light incident on the HCL discharge that is resonant with some transition of the ions in the discharge can excite additional electrons to higher, more loosely bound energy levels, slightly increasing the current through the discharge as it becomes easier to strip electrons from these ions. To measure this small effect, the HCL discharge was illuminated by the laser, which was chopped at 6 kHz. The resulting small 6 kHz variation in discharge current was measured with a lock-in amplifier. A diagram of the electrical circuit used for this measurement is shown in Figure 3.4.

Figure 3.3 shows the signal measured by the lock-in as a function of fre-




Figure 3.4: The circuit diagram for optogalvanic spectroscopy. The HCL can be treated as a resistor whose value varies slightly depending on whether the laser is shining through it. It forms a voltage divider with the 35k resistance, which is constructed from high-wattage resistors since there is a few hundred volts across it. The voltage across the 35k resistance thus varies slightly as the laser is chopped. A high-pass filter extracts the time-varying portion of this voltage, which goes to the lockin.

quency, showing a clear peak at the frequency of the Ba<sup>+</sup> transition. This measurement verified that the HCL was producing Ba<sup>+</sup> ions in the right state.

### 3.2.2 Absorption measurements

Next I attempted to detect measurable absorption of the laser light by the Ba<sup>+</sup> ions in the HCL discharge. The DAVLL operates by detecting the differential absorption of the two circular polarizations, so I wanted to verify that some reasonable amount of light was actually being absorbed.

Unfortunately the amount of absorption seemed to be very small. A first simple measurement in which a laser power meter was placed behind the HCL as the discharge was turned on and off established an upper bound of 1% absorption by the discharge. To try to detect some absorption, a more sensitive experiment was carried out, depicted in Figure 3.5. The beam was split in two and the two beams were sent to the two photodiodes of a




Figure 3.5: The optical setup for the absorption measurement. The combination of linear polarizer (P), half-wave plate (HWP) and polarizing beamsplitter (PBS) produces two beams whose relative intensities are adjusted to be exactly equal by turning the half-wave plate. One beam passes through the HCL. The beam intensities are measured by photodiodes and the difference in photocurrents give a measure of the absorption in the HCL.

differencing photodetector (circuit diagram shown in Figure 3.7). On the way to the photodetector, one beam passed through the HCL, while the other was used as a reference. The large gain of the photodetector allowed the detection of small differences in the received power in the two beams as the laser frequency was scanned over the 650 nm transition.

This method detected a clear bump in the signal from the differencer, and thus a bump in the absorption spectrum of the HCL discharge, at the correct frequency. Figure 3.6 shows the detected absorption as a function of frequency. The observed maximum absorption was about 0.1%.

### 3.3 Implementation and characterization

Having detected small but finite absorption of 650 nm light by the HCL discharge, we next set up the full DAVLL setup of figure 3.1 and attempted to detect a DAVLL error signal. Forewarned by the literature that high-quality optical components were necessary to obtaining a workable error signal ([2], [8]) we used a Glan-laser calcite polarizer<sup>7</sup> for the linear polarizer and a

<sup>7</sup>Newport part 10GL08




Figure 3.6: In the depicted measurement, each beam had a power of 490 microwatts, while the maximum observed absorption was 0.5 microwatts. As expected, the absorption signal has the same width as the optogalvanic signal: about 1 GHz.




Figure 3.7: The circuit diagram of the photodifferencer used to detect absorption and the DAVLL signal. Based on a design in [5]. The purpose of the potentiometer is to compensate for differences in the series resistance of the photodiodes; in practice it didn't have an effect on measurements.

Wollaston prism<sup>8</sup> for the polarizing beamsplitter.

We experienced marginal success in obtaining an error signal. A typical error signal obtained is shown in figure 3.8. The DAVLL error signal is discernable, but exhibits a number of serious deficiencies as a signal for laser locking.

1. The signal amplitude is very small. As a result its slope is also small, and simple noise would lead to large fluctuations in the laser frequency.
2. The error signal is superimposed on a background that is not flat. Often it was observed that the DAVLL signal was superimposed on a signal that varied like a sine wave with frequency, as in the figure. The hypothesized origin of this sort of sine-wave background is an unintentional Fabry-Perot cavity somewhere in our optical setup producing interference effects that vary like a sine wave with frequency. In addition to the sine-wave variation there was also generally an overall slope that is even more poorly understood.
3. The vertical displacement of the signal is not stable with time. As a result, the zero-crossing of the signal, which is the frequency set-point, is not stable, as it must be for the technique to be useful. This may be due to temperature fluctuations causing changes in the optical properties of the components, which is a commonly reported problem in DAVLL locks. No attempt was made to temperature stabilize any part of the setup, and it was observed that breathing on the components could produce vertical displacements in the signal, suggesting that temperature fluctuations can account for the observed time-variation.

The most serious deficiency is the first. Because our signal is so tiny, it is easily overwhelmed by other weak effects such as (2) above, which would not be a problem if the differential absorption signal were larger. We cannot simply amplify the signal with, say, a larger photodetector gain, because that would also amplify the effects of (2) and (3) in the same proportion. One partial remedy would be to obtain stronger magnets, which would move the LCP and RCP absorption peaks farther apart and so increase the differential absorption of the two polarizations. However no matter how strong the magnets we would not obtain a greater differential absorption than about 0.05%. This is because our

---

<sup>8</sup>Thorlabs part WP10




Figure 3.8: A DAVLL signal obtained with a laser power of 500 microwatts and a discharge current of 18 mA. The desired signal is superimposed on an unwanted and poorly understood background signal. The red dashed line is a fit to this particular background of the form  $f(\lambda) = A + B\lambda + C \sin D\lambda$ . The desired signal is the deviation from the dashed red line. While the background was consistent from minute to minute, it changed from day to day or when the setup was modified.

peak absorption with no magnetic field was 0.1%, so the peak absorption of each polarization is half that, or 0.05%. So we could not obtain a greater difference absorption between LCP and RCP light than 0.05%.

### 3.4 Conclusions

The main result of this exploration of the DAVLL frequency stabilization technique for the 650 nm line of Ba<sup>+</sup> was that the HCL we used did not create a discharge with enough absorption to be useful for laser locking. DAVLL is usually implemented using an atomic vapor, but we wanted to use Ba<sup>+</sup> ions in an electrical discharge since that is the species we are trapping. A recent paper [10] reported successful implementation of a DAVLL lock very similar to our attempt. They used an HCL by the same manufacturer (with ytterbium instead of barium); however they measured 6% absorption of the laser by their transition of interest, while we measured only 0.1%.

It might well be feasible to use a DAVLL lock for a different laser line, if it exhibited more absorption in the HCL discharge. The 493 nm line of Ba<sup>+</sup> might experience more absorption. Alternatively, it might be feasible to use an iodine vapor cell, as with saturated absorption spectroscopy. Neither of these were attempted.

In addition to this the problem of temperature dependence was not confronted. While noticeable, it was secondary in this instance to the problem of a small signal. Even if better absorption were obtained, it might be necessary to enclose the DAVLL setup in some sort of temperature-insulated or -stabilized container to achieve good frequency stability.

# Temperature-stabilized optical cavity

An optical cavity is formed by two parallel partially reflective surfaces. In this cavity only certain frequencies of light are resonant, and this effect can be used as a frequency reference for a laser. An attractive feature of optical cavities is that the width, in frequency, of the resonances can be made very small: on the order of hundreds of kilohertz. This means that a cavity can be used to make a very precise frequency reference.

One problem with using a cavity as a frequency reference is that the resonant frequencies depend very sensitively on the length of the cavity (the distance between the two mirrors). Since this length is liable to change slightly due to thermal expansion as the ambient temperature fluctuates, it is important to stabilize the temperature of the cavity very carefully in order to keep the resonant frequency stable. This thermal dependence can also be used as a tool, by using the temperature to tune the length of the cavity to place its resonance at the desired frequency.

This section describes an in-progress effort to construct a temperature-stabilized cavity for the 493 nm Ba<sup>+</sup> cooling laser. This laser operates by frequency-doubling light from a 987 nm laser, and the cavity will actually work with this undoubled 987 nm light to indirectly stabilize the frequency of the 493 nm light.

This work has relied heavily on the experiences reported in [6] and [1].




Figure 4.1: In every encounter with a mirror, a fraction of light is transmitted and a fraction reflected. Interference effects between photons that have bounced across the cavity different numbers of times leads to a nontrivial relation between the frequency of the light and how much is transmitted through the cavity.

## 4.1 Principles

### 4.1.1 Optics

The simplest optical cavity is composed of two identical parallel planar mirrors of some amplitude reflection coefficient  $r$  and amplitude transmittance coefficient  $t$  separated by a distance  $L$ . Here  $r$  and  $t$  are the amplitudes to reflect and transmit electromagnetic waves; if a beam of intensity  $I$  is incidence on a mirror then the reflected beam has intensity  $|r|^2 I$  and the transmitted beam has intensity  $|t|^2 I$  is transmitted. A lossless mirror has  $r^2 + t^2 = 1$ .

Consider light incident from the left on such an optical cavity as in Figure 4.1. Call the incident amplitude  $E$  (so the incident intensity is  $|E|^2$ ). Then we can calculate the amplitude for the light to be transmitted through the cavity by considering the possible paths it could take through the cavity. Suppose that the phase delay the light experiences on one crossing of the cavity is  $\delta$ . Then we can sum the amplitudes of the various paths that end up with the light exiting the cavity to the right to find the amplitude transmission

coefficient of the cavity,  $t_{cav}$ .

$$\begin{aligned}
 t_{cav}E &= Ete^{i\delta}t + Ete^{i\delta}re^{i\delta}re^{i\delta}t + \dots \\
 &= Et^2e^{i\delta} \sum_{n=0}^{\infty} (r^2e^{2i\delta})^n \\
 &= \frac{Et^2e^{i\delta}}{1 - r^2e^{2i\delta}}
 \end{aligned}$$

The intensity transmission coefficient for the cavity,  $T_{cav} = |t_{cav}|^2$  can be written

$$T_{cav} = \frac{t^4}{(1 - r^2)^2} \cdot \frac{1}{1 + F \sin^2 \delta} \text{ with } F = 4 \frac{r^2}{(1 - r^2)^2}$$

In the case of ideal lossless mirrors,  $r^2 + t^2 = 1$  and so the leading constant above is unity.

Highly reflective mirrors have  $r$  close to 1, giving cavities with large  $F$ . Large  $F$  means that there is only appreciable transmission in the narrow regions where  $\sin^2 \delta$  is very small, corresponding to narrow peaks in the transmission spectrum of the cavity. So for large  $F$ , the transmission through the cavity is only non-negligible in small peaks around  $\delta = n\pi$ ,  $n$  an integer.

For light of frequency  $f$ , the phase delay  $\delta = 2\pi fL/c$ . Thus the condition for a frequency to have non-negligible transmission through the cavity is

$$2\pi fL/c = \pi n \text{ or } f = n \frac{c}{2L}$$

So there is a set of resonant frequencies with a constant spacing between them. The spacing  $c/2L$  is called the free spectral range (FSR).

For small  $\delta$ ,  $T_{cav}$  can be approximated as a Lorentzian function:

$$T_{cav} \propto \frac{1}{1 + F\delta^2}$$

This form makes it easy to find the width of the resonances. In terms of the phase  $\delta$  the width is given by

$$\delta_{FWHM} = 2/\sqrt{F}$$

or in terms of frequency by

$$f_{FWHM} = \frac{c}{\pi L \sqrt{F}}$$

The ratio of the FWHM to the FSR is a dimensionless number written  $\mathcal{F}$  called the finesse of the cavity, which is determined solely from the mirror reflectivity as

$$\mathcal{F} = \frac{\pi \sqrt{F}}{2} = \pi \frac{r}{1 - r^2} = \pi \frac{\sqrt{R}}{1 - R}$$

where  $R = r^2$  is the intensity reflection coefficient of the mirrors. The finesse quantifies the “quality” of the cavity in terms of its ability to produce a sharp resonance.

Above we have considered a cavity composed of two planar mirrors, with plane waves of infinite transverse extent incident on it. This is a simplification; real laser beams have a finite transverse width, which causes them to diverge as they bounce back and forth in a planar cavity. A finite-width beam of light therefore cannot be exactly resonant with a planar cavity.

This is no longer the case if we allow one or both mirrors to be curved. Our cavity has one planar and one spherical mirror. Such a cavity supports resonant modes of finite transverse width. For instance the basic TEM<sub>00</sub> mode, which we like because it has a Gaussian cross section like our laser beam, has a definite spot size on the planar mirror of [9]

$$w_0 = \left( \frac{L\lambda}{\pi} \right)^{\frac{1}{2}} \left( \frac{g}{1 - g} \right)^{\frac{1}{4}} \text{ where } g = 1 - L/R$$

Here  $\lambda$  is the wavelength of the light and  $R$  is the radius of curvature of the spherical mirror. If the incident beam is non-Gaussian or has a different waist size it will excite other transverse modes of the cavity. The different transverse modes have different resonant frequencies, given by the formula [7]

$$f = \frac{c}{2L} \left( n + \frac{a + b + 1}{2} \cos^{-1} \sqrt{g} \right) \text{ for the } n\text{th TEM}_{ab} \text{ mode.}$$

Thus if we scan the laser frequency we will see, not one peak in each FSR, but a multitude of peaks, repeating each FSR, as we pass over the resonant frequencies of different transverse modes. In general we want to achieve good

“mode matching”—that is, tweak the incident beam’s shape so that we excite the TEM00 mode as much as possible and the others as little as possible.

The above has not touched on how to generate a frequency stabilization error signal from a cavity. A simple way to do this would be to take the derivative of the transmission through the cavity, just saturated absorption spectroscopy produced an error signal by differentiating the absorption signal. We plan to use the Pound-Drever-Hall method described in [3] to lock the laser, which is essentially a more sophisticated version of this simple idea.

### 4.1.2 The influence of temperature

As was mentioned above, the cavity modes are sensitive to temperature variations because these change the length of the cavity. Suppose we lock the laser to the frequency of the  $n$ th resonant mode of the cavity, which is given above as having the frequency

$$f = n \frac{c}{2L}$$

where for simplicity we are again considering the case of the planar cavity. Then if the cavity length changes slightly the laser frequency will vary also according to

$$\frac{df}{dL} = -n \frac{c}{2L^2} = -\frac{f}{L}$$

The thermal expansion or contraction of a material of length  $L$  is given by

$$\frac{dL}{dT} = \alpha L$$

where the coefficient of thermal expansion,  $\alpha$ , has units of Kelvin<sup>-1</sup>. Thus for a cavity subjected to temperature variations we have

$$\frac{df}{dT} = \frac{df}{dL} \frac{dL}{dT} = -\alpha f$$

which gives the variation of the laser frequency with temperature in terms of the coefficient of thermal expansion of the cavity. This equation can be used to find the required temperature stability of the cavity needed for a desired frequency stability of the laser.

## 4.2 Implementation

### 4.2.1 Construction

The cavity was built from two mirrors<sup>9</sup> at each end of a quartz tube of length  $L = 25$  cm. One mirror was planar, the other spherical with a radius of curvature of  $R = 30$  cm. This radius of curvature is partially constrained by the fact that the cavity only forms a stable resonator if  $R > L$ .

The length of the tube gives a calculated FSR of 600 MHz. Quartz was chosen for the tube material because of its low but nonzero coefficient of thermal expansion of  $5 \times 10^{-7}$  Kelvin<sup>-1</sup>. The cavity is to be used to stabilize 983 nm (305 THz) infrared light giving a temperature dependence  $df/dT = -150$  kHz/millikelvin. The desired frequency stability is less than one MHz, requiring temperature stability to a few millikelvin. A temperature change of 4 Kelvin will shift the frequency of a resonance by one full FSR, so that it will not be difficult to tune the frequency of the locked laser to the desired number by changing the temperature.

To allow temperature stabilization, the cavity was closely wrapped in resistive wire. The particular wire used<sup>10</sup> was chosen because of its convenient resistance per unit length of 0.33 ohms/foot. If the resistance of the heating wire is too low or too high, a typical 30V/1A power supply will either not be able to supply enough voltage or enough current to deliver a reasonable amount of heat to the cavity. The total length of wire wrapped around the tube was measured to have a resistance of about 6 ohms. So an amp of current, which is reasonable to expect from a power supply, will deliver 6 watts of heat to the cavity, which is enough that the close-wrapped wire becomes slightly warm to the touch.

To sense the temperature of the cavity, two strings of eight thermistors<sup>11</sup> were glued along the length of the cavity.

To hold the mirrors on each end of the quartz tube, two mirror mounts were constructed. Each mirror is held in the mount by threaded ring. Figure 4.2 shows the cavity at this stage of construction, during the process of aligning

---

<sup>9</sup>CVI parts LDM-980-0537-0 and LDM-980-0537-0-0.30CC

<sup>10</sup>17 Gauge Nichrome 60 7/25 Stranded w/Kapton from Pelican Wire Company, Inc.

<sup>11</sup>Betatherm part 10K3A1B




Figure 4.2: The cavity during alignment. The heating wire is visible along most of the length of the quartz tube. Black tape applied lengthwise along each side of the tube provides electrical insulation for two strings of eight thermistors. The stainless steel mirror mounts are visible at each end, held in tilt mounts.

the mirrors.

### 4.2.2 Alignment procedure

The cavity is currently being aligned. The mirror mounts do not slide tightly on to the end of the cavity; there is deliberately some play to allow for adjustments of the mirror mounts to perfect their alignment. The following procedure has been used so far to align the cavity, with the objectives of obtaining good mode matching (high transmission through the cavity in the TEM<sub>00</sub> mode and little excitation of higher transverse modes). The optical setup used for alignment is shown in Figure 4.3.

1. A beam from a fiber was sent through a 3:1 telescope to reduce its size

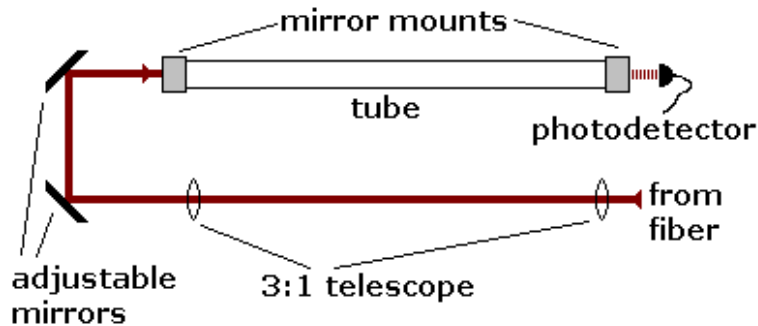


Figure 4.3: The alignment optical setup. The angles of both the two external mirrors and the two mirror mounts are adjustable to optimize mode matching.

to closer to the the spot size of the TEM<sub>00</sub> mode on the flat mirror.

2. The quartz tube was clamped solidly in place with two V-clamps.
3. Two circular pieces of paper were cut to the same diameter as the quartz tube. Small holes were pricked in their centers, and they were temporarily taped on each end of the cavity. The two mirrors were adjusted to get as much light passing through the two holes as possible. This aligned the beam along the central axis of the tube.
4. The curved mirror was clamped to a pan-tilt mount and slid onto the far end of the cavity (closer to the photodetector). Its angle was adjusted so that it reflected the beam directly backward.
5. The flat mirror was clamped to a pan-tilt mount and slid onto the near end of the cavity (farther from the photodetector). Its angle was adjusted so that it reflected the beam directly backward.
6. The laser was made to scan back and forth over a frequency range corresponding to one FSR, and the signal from the photodetector was observed on an oscilloscope, showing many peaks. The angles of the cavity mirrors and the external mirrors and the size of the iris were adjusted to maximize the height of the largest peak and minimize the rest.

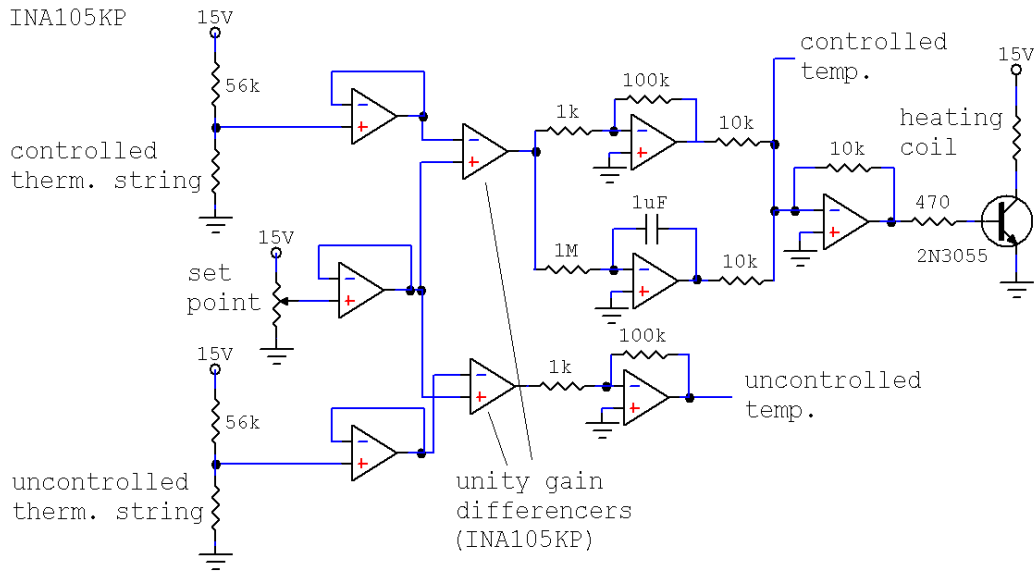


Figure 4.4: The temperature box circuit diagram. The two thermistor strings are made part of voltage dividers to create voltages that vary with temperature. These voltages and the set point voltage are sent through simple op-amp buffers so that they are not influenced by the rest of the circuit. Then the thermistor voltages are differenced from the set point to generate error signals. The controlled error is the input to a PI circuit that feeds back into the heating coil current to stabilize the temperature. The two error signals (after amplification by 100) are outputs that can be recorded by computer to assess temperature stability over time.

## 4.3 Characterization

### 4.3.1 Temperature stability

Before the mirrors were glued on to the cavity, an electronics box for temperature control was constructed and tested to ensure that good temperature stability could be achieved. The circuit diagram of the temperature control box is shown in figure 4.4.

In a simple test of temperature stability, the temperature of the tube was

sensed with two strings of eight thermistors in series glued to the heating coil on opposite sides of the tube. These strings of thermistors were used to obtain average temperatures along the length of the tube: a single thermistor would only have provided the temperature at a single point. Two strings of thermistors were used to answer the following question: when the temperature controller stabilized the average temperature on one part of the tube, how large were the temperature fluctuations in other parts of the tube? There will be small independent temperature fluctuations in different parts of the tube. And presumably the part of the tube whose temperature is being actively stabilized will be the part of the tube with the smallest temperature fluctuations. So measuring the temperature stability using the same thermistors that are being used in the temperature control feedback loop is likely to underestimate the typical magnitude of temperature fluctuations along the tube.

Thus one string of thermistors—the “controlled string”—provided the feedback to the temperature control box. These electronics stabilized the temperature indicated by controlled string with a simple analog PI feedback loop. The other string—the “uncontrolled string”—was not part of the temperature control feedback loop.

The tube with heating coil was wrapped in aluminum foil to help equalize the temperature across the tube, then inserted in a tube of foam thermal insulation. This was placed in a cardboard box in a cardboard box to provide shielding from air currents. The temperature variations in both strings of thermistors were then recorded over a span of 24 hours, and the results appear in Figure 4.5. A test was also made to see how quickly the temperature could stabilize again after a change in the set point, shown in Figure 4.6.

During the stability test, the controlled string’s temperature remained very steady, as expected, varying by 1 or 2 millikelvins at most, and was completely insensitive to changes in the room temperature. However it was observed that the uncontrolled string’s temperature varied to some extent with the ambient room temperature, as shown in Figure 4.5. A variation in the room temperature produced a variation in the coil temperature about 0.5% as large. Ideally we would likely to totally eliminate this effect. It means that although we are attempting to stabilize the tubes temperature, in fact we are only stabilizing the temperature of one part of it, and we will not achieve the frequency stability we desire. Still, the observed temperature stability is




Figure 4.5: Top: temperature fluctuations in the uncontrolled thermistor string over a 24-hour test. The tube was stabilized at about 30 degrees Celsius. Bottom: fluctuations in room temperature over the same period, divided by a factor of 200. Temperature stability is better than 5 mK over 24 hours. A very strong correlation between variations in the coil temperature and the room temperature is apparent.




Figure 4.6: Recovery of the temperature controller after a change in the temperature set point (top: controlled string; bottom: uncontrolled string). Temperatures of both thermistor strings stabilize after about 15 minutes.

not bad. Using the value  $df/dT = -150$  kHz/millikelvin found above, the observed temperature stability would correspond to a frequency stability of better than 1 MHz if the cavity were in full operation.

### A simple model of temperature variations

In an effort to understand the origin and remedy of this variation of the uncontrolled temperature with room temperature, I constructed a very simple model of how this variation occurs. I believe that the reason that the uncontrolled temperature varied when the controlled one was kept steady is that the two thermistor strings had different strengths of couplings to the external ambient temperature. Approximate the two regions (call them 1 and 2) of the tube sensed by the two strings of thermistors as separate objects with definite temperatures  $T_1$  and  $T_2$  that are in thermal contact with each other and with the room. There is some thermal conductivity  $b$  (with units of Watts/Kelvin) between the two regions. Each region also has some thermal conductivity with the environment; call these values  $a_1$  and  $a_2$ . The room is at some ambient temperature  $T_{amb}$ . Finally, to prevent useless complication, suppose regions 1 and 2 are the same size and so receive equal amounts of heat per unit time from the coil; call the power delivered to each region  $P$ .

In the steady state the power delivered to each side is equal to the heat conducted away:

$$\begin{aligned} P &= a_1(T_1 - T_{amb}) + b(T_1 - T_2) \\ P &= a_2(T_2 - T_{amb}) + b(T_2 - T_1) \end{aligned}$$

Now,  $T_{amb}$  may vary with time. If this happens we can vary  $P$  to keep  $T_1$  (the temperature of the controlled string) constant. We ask: when we do this, how does  $T_2$  vary with  $T_{amb}$ ? Subtracting the above two equations,

$$\begin{aligned} 2b(T_1 - T_2) + a_1(T_1 - T_{amb}) - a_2(T_2 - T_{amb}) &= 0 \\ (-a_2 - 2b)T_2 - (a_1 - a_2)T_{amb} &= \text{constant} \\ (a_2 + 2b)dT_2 &= (a_2 - a_1)dT_{amb} \\ \frac{dT_2}{dT_{amb}} &= \frac{a_2 - a_1}{a_2 + 2b} \end{aligned}$$

This last equation gives the variation of the temperature of region 2 when the ambient temperature varies but region 1 is kept at a stable tempera-

ture by adjusting the current through the coil. We see that even though region 1 is temperature stabilized, region 2 is not necessarily kept at a stable temperature.  $T_2$  will vary with  $T_{amb}$  if there is a difference in the thermal conductivities  $a_1$  and  $a_2$  between each region of the tube and the environment. This is intuitive: if each region has a different coupling with the environment, then the heat input needed to stabilize  $T_1$  will not be the same as the heat input needed to stabilize  $T_2$ . We also see that this unwanted variation is suppressed by  $b$ , the thermal conductivity between the two regions. This also makes sense: good conductivity of heat between the regions pulls them toward the same temperature, preventing  $T_2$  from varying too much if we keep  $T_1$  constant.

For this reason, the cavity was wrapped in aluminum foil before being enclosed in thermal insulation in order to increase  $b$ ; that is, to try to keep the entire cavity at the same temperature so that stabilizing the temperature of one part stabilizes the temperature of the entire cavity.

To further reduce this effect, the working cavity will have even more thermal insulation, which will reduce the magnitudes of both  $a_1$  and  $a_2$ , and thus hopefully their difference. Not only will the cavity be immediately surrounded by thermal insulation, but the box in which it resides will have a layer of thermal insulation on it as well.

### 4.3.2 Measurements of the cavity

During the mode-matching process, oscilloscope traces were recorded showing the mode structure of the cavity. Figure 4.7 shows two traces that resulted from scanning across 1 FSR of the cavity, demonstrating the importance of the beam size to achieving good mode matching.

The FSR of the cavity and the width of the TEM00 resonance peak were measured in the following way. A 20 MHz AC signal was superimposed on the current to the laser diode, creating frequency sidebands 20 MHz to each side of the laser's central frequency. As the laser frequency was scanned, these sidebands showed up as extra peaks spaced 20 MHz to each side of each resonance of the cavity. Since the spacing of these peaks was known to be 20 MHz, the measured spacing of these peaks in an oscilloscope trace allowed the scan rate of the frequency to be determined. Then the FSR and




Figure 4.7: Resonances visible in the cavity transmission spectrum for a rather wode beam (bottom, waist about 1.25 mm) and beam narrowed by an iris to get better mode matching (top, waist about 0.4 mm). The smaller beam excites fewer transverse modes, giving a cleaner mode structure in the top trace. The x axis spans one FSR; the large peaks at the extreme left and right in each scan are the TEM00 mode.

resonance widths were calculated from the oscilloscope trace by converting time into frequency via the known laser frequency scan rate.

The FSR was determined to be 575 MHz, with an uncertainty of about 5%. The uncertainty comes from the fact that the laser scan rate was not uniform across the range of the scan; at one end it scanned slower and at the other end faster. To account for this somewhat, the spacing of the sidebands around several resonances in different parts of the scan were measured and averaged to try to calculate an average scan rate.

The FWHM of the TEM00 resonance was determined to be 300 kHz, with an uncertainty of about 20%. Here the uncertainty comes from the fact that the cavity was constantly perturbed by vibrations of the optical table. These vibrations resulted in constant slight variations in the length of the cavity, resulting in random distortions of the oscilloscope trace. As a result the measured width of the resonance was not constant from scan to scan. The effect of these vibrations can be diminished by increasing the scan rate so that the measurement is accomplished faster than the timescale of the vibrations, but our photodetector's finite response time placed a limit on how fast we could go. Figure 4.8 is a typical depiction of the horrible distortion of a single peak caused by vibrations during a very slow scan.

To reduce the effect of vibrations on the working cavity, it will be placed atop a small (18 inch square) table supported by vibration-dampening feet and will be surrounded by acoustic insulation to prevent sound from disturbing it.

The finesse—the ratio of the FSR to the resonance FWHM—corresponding to the above measurements is around 1500 to 2500. From equation for the finesse, this range corresponds to an intensity reflectivity coefficient of 99.79%-99.87%

Using a power meter, the intensity transmission coefficient  $t^2$  of the mirrors was measured at about 0.03%, though the flat mirror's transmission varied by  $\pm 50\%$  across its surface. During alignment a peak of 7% intensity transmission through the cavity at the TEM00 resonance was achieved. From the expression derived above for transmitted intensity through the cavity these values correspond to a intensity reflection coefficient of  $r^2 = 99.89\%$ , which is not too far from agreement with the estimate from the crudely measured finesse.




Figure 4.8: In this oscilloscope trace, the frequency of the laser was scanned over a single resonance peak, but vibrations disrupted the smooth scan so that what should be a single Lorentzian peak appears as a multitude of peaks.

## 4.4 Current status

The temperature stabilized cavity project is still in progress. As reported above, good thermal stability has been already been achieved with somewhat less insulation than will surround the working cavity. Decent mode matching has been achieved; the next step is for the mirror mounts to be glued on permanently in a good alignment. The cavity will then be enclosed in a thermally insulated box on top of its vibration-isolated table. The frequency locking electronics for the Pound-Drever-Hall locking scheme must also be constructed.

# Summary

I have investigated three different techniques for laser frequency stabilization of the 493 and 650 nm lasers used for cooling of trapped barium ions in the molecular ion trapping group at Northwestern University. Laser locking with saturated absorption spectroscopy of iodine with good stability was demonstrated but has not been put in day-to-day use because it was disrupted by the RF power supply used for trapping and because a dedicated lock-in amplifier for it needs to be constructed. A dichroic atomic vapor laser lock using barium ions in an electrical discharge was found to be infeasible for the 650 nm laser due to too little absorption in the discharge, but possibly DAVLL would experience greater success with the 493 nm line or by using an iodine vapor cell. Finally a temperature stabilized cavity is being constructed to stabilize the 493 nm laser. The cavity resonances have been measured to be quite narrow and decent coupling of light into the cavity has been achieved. In separate tests, good temperature stabilization of the tube has been observed. The frequency locking electronics that generate the error signal for this scheme still need to be constructed.

# Bibliography

- [1] Lauge Christensen. Bachelor project: Construction of a reference cavity with variable length for frequency locking to the d2 line in caesium., 2009.
- [2] Kristan L. Corwin, Zheng-Tian Lu, Carter F. Hand, Ryan J. Epstein, and Carl E. Wieman. Frequency-stabilized diode laser with the zeeman shift in an atomic vapor. *Appl. Opt.*, 37(15):3295–3298, May 1998.
- [3] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B: Lasers and Optics*, 31:97–105, 1983. 10.1007/BF00702605.
- [4] Nguyen et al. Challenges of laser-cooling molecular ions. *arXiv:1102.3368v3*.
- [5] Philip Hobbs. *Building Electro-Optical Systems*. Wiley, 2009.
- [6] Frank Jensen. Laser frequency stabilization for use in stirap experiments. Master’s thesis, 2004.
- [7] P. W. Milonni and J. H. Eberly. *Lasers*. Wiley, 1988.
- [8] Jessica M. Reeves, Ofir Garcia, and Charles A. Sackett. Temperature stability of a dichroic atomic vapor laser lock. *Appl. Opt.*, 45(2):372–376, Jan 2006.
- [9] A. E. Siegman. *Lasers*. University Science Books, 1986.

- [10] E. W. Streed, T. J. Weinhold, and D. Kielpinski. Frequency stabilization of an ultraviolet laser to ions in a discharge. *Applied Physics Letters*, 93(7):071103–071103–3, Aug 2008.