LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T2400157-v1** 09/27/2024

Frequency Locking a Laser to ^{87}Rb

Briana A. Chen
Mentors: Ian MacMillan and Lee McCuller

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW17-161
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

1 Abstract

To prevent frequency drifts over time, lasers are frequency-locked to reduce unwanted degrees of freedom in a system. This has applications in projects ranging from improving filter cavities in LIGO and GQuEST to atom trapping experiments. Lasers are typically locked, as in LIGO, using the Pound Drever Hall technique, which passes modulated light through a cavity to form a feedback loop tuning the laser frequency back onto resonance. However, the cavity only acts as a relative reference because its constituent mirrors are susceptible to fluctuations in distances, reflectivity, and transmissivity. In contrast, atoms have absolute transitions that enable them to be an absolute locking reference. In this project, the 780 nm laser was locked to the D2 780 nm transition of ^{87}Rb via a similar feedback loop as PDH locking. By passing a modulated probe beam and counterpropagating pump beam through a vapor cell of ^{87}Rb, the system's frequency discrimination capabilities are improved. Future work includes implementing second harmonic generation to lock a 1560 nm laser to the vapor cell and incorporating a specialized controller into the feedback loop to correct frequency drift with minimal overshoot.

2 Introduction

The GQuEST (Gravity from the Quantum Entanglement of Space-Time) experiment aims to detect a quantum gravity signal by using correlated interferometers to measure fluctuations in spacetime. If successful, this would be an exciting step in developing the theory of quantum gravity.

GQuEST has the same fundamental setup as LIGO (Laser Interferometric Gravitational-Wave Observatory). Light travels down and back the two arms of a Michelson interferometer and recombines. The phase difference between the two light beams is measured and analyzed in search of a statistically significant signal. The presence of any additional motion of the mirrors (in LIGO for example, a gravitational wave with frequency ϵ) will modulate the carrier (i.e. laser) frequency ν_c, adding sidebands of frequency $\nu_c \pm \epsilon$ [1]. In LIGO, a DC readout method extracts the interference between these sidebands and the original laser frequency. In GQuEST, photon counting will be used to extract the photons of a specified frequency from the interferometer setup. The optical filters, which are a series of optical cavities, then allow the photons of interest to pass through. All of these cavities have the property that light can only transmit through if it is on resonance (i.e. of a certain frequency ν).

Ideally, you only want light of frequency ν to pass. However, in reality, the transmitted light from the optical cavity will have a Lorentzian spectrum (see Fig. 5 for the transmission shape). One can obtain a peak with greater attenuation away from resonance by placing N cavities one after another such that the Lorentzian shape is raised to the Nth power, narrowing the linewidth and improving signal-to-noise. An even better option would be to use atomic cells instead of optical cavities as the resulting transmission lineshape will take on

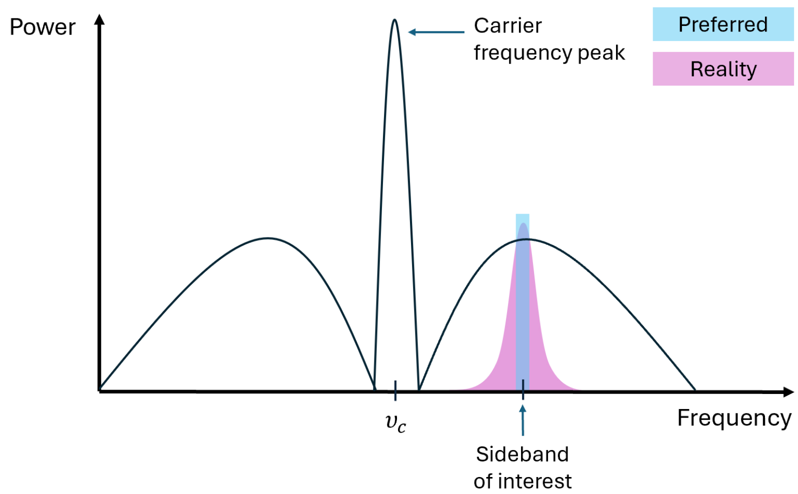

Figure 1: Preferred and actual behavior of optical filters. The power spectrum of the recombined light from the interferometer will display a single prominent peak around the carrier frequency along with a smooth curve on either side of the peak that represents the sidebands formed from a stochastic quantum gravity signal. Noise is not shown on top of this power spectrum. The optical filter should solely isolate the intensity from the frequency of interest but in reality, will pick up signals of other frequencies.

other forms like the Gaussian lineshape, which could significantly improve the signal-to-noise ratio.

Designing better optical filters is necessary to reduce noise and consequently amplify the signal-to-noise ratio at the sideband of interess, suppressing the carrier. They must be tested and adjusted to minimize leakage outside their intended frequency range. A magneto-optical trap (MOT) is in development for future design and testing of optical filters in the GQuEST experiment.

The simple principle behind the MOT is as follows. Upon absorbing a photon, an atom will gain some recoil momentum. Although it will emit the photon later due to the state's finite lifetime, the atom's momentum is dominated by absorption due to the isotropic nature of spontaneous emission. However, at a high enough power, stimulated emission will produce an opposing recoil to the absorption, causing the atom to theoretically have a zero net change in momentum as it is ”stopped” by the stimulated emission. To cool the atoms further, the laser beams operate at a frequency below resonance to take advantage of Doppler shifting, which will increase the absorption probability of an atom traveling towards the laser beam and subsequently slow it down. If an atom is travelling away from the laser beam, it has a lower absorption probability and would thus not be slowed down. The counterpropagating beams enable the slowing of atoms moving with different velocities in 3D space and the detuned laser produces this damping force on the atoms [2].

The addition of a magnetic field and circularly polarized laser beams causes the atom to undergo a different transition depending on the direction of circular polarization. The radiation pressure force pushes the atom towards the center of the trap and is unbalanced due to the resulting Zeeman effect [3]. This imbalance enables the atoms to be confined beyond the Doppler cooling limit [2].

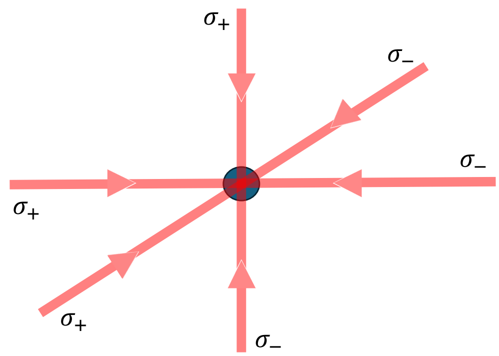

Figure 2: Basic diagram of a MOT. The atom is trapped in the middle and the 6 counter-propagating beams are of different handedness circular polarizations. [2].

The MOT requires frequency-stabilized laser beams to maintain the detuned laser frequency for performing Doppler cooling. The goal of this project is thus to produce a laser frequency-locked to ^{87}Rb suitable for use in the MOT, which will trap ^{87}Rb atoms. In the MOT, an acousto-optic modulator (AOM) can be used to adjust this locked frequency to the desired detuning.

2.1 How does a laser's frequency drift?

Lasers work by exciting electrons in some medium that amplifies light through population inversion. This inversion produces a higher population of atoms in the excited state. Undergoing the process of stimulated emission, the atoms drop down in energy level and release a photon. This photon is reflected back into the medium to produce more photons of the same energy in an optical cavity. The greater amount of reflections in unison produces more photons and ultimately, the energy of the photons enables them to pass through the mirror, producing the laser beam [4].

The gain medium will have different properties with shifts in temperature or the laser's mechanical properties. For example, changes in the medium's temperature shift the laser's gain curve, changing the laser output power and emission properties [5]. Other effects on laser emission include pressure changes, environmental effects, age, and electromagnetic fields, all of which contribute to frequency drifting [5].

2.2 Pound Drever-Hall Locking

The basic process behind locking a laser's frequency is to measure the laser's frequency, determine how far away this frequency is from a cavity's resonant frequency, and then send the error back into the laser input to tune the laser. Fig. 3 depicts the experimental feedback cycle of Pound Drever-Hall (PDH) locking, which has been a standard method of laser locking with an optical cavity [5].

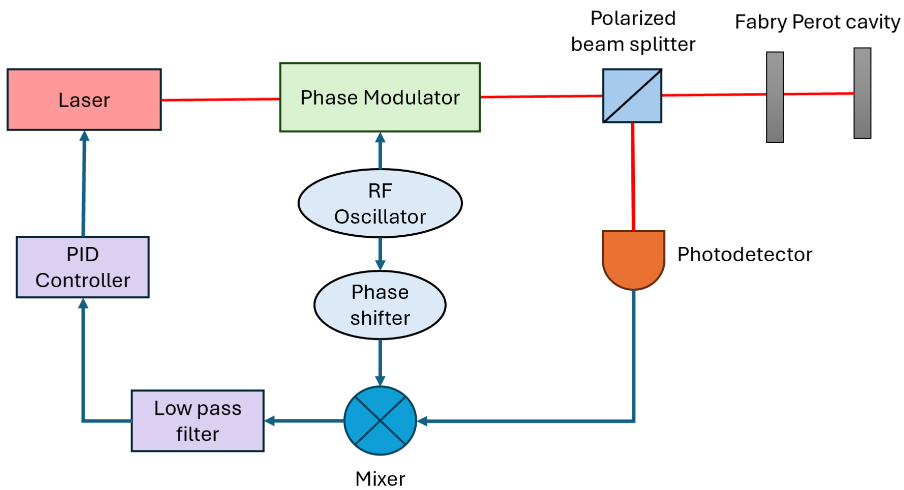

Figure 3: Basic schematic of PDH locking feedback loop.

To measure the frequency difference, the PDH technique uses a Fabry-Perot cavity, which has the important property that the total reflected and transmitted intensity is dependent on the frequency of light sent into the cavity [6]. The cavity consists of two partially transmissive and reflective surfaces that cause higher-order reflections and transmissions.

Assuming the surfaces have uniform reflection and transmission coefficients, the material between the two surfaces has a higher refractive index than those outside, and the beam enters with normal incidence, the electric field of the transmitted light is given by:

$$E_t = \frac{E_o t^2}{1 - |r|^2 e^{i\Delta\phi}} \tag{1}$$

where $\Delta\phi$ is defined as the phase acquired by the light from one round trip in the cavity, i.e. $\Delta\phi = \frac{\omega}{\Delta f_{fsr}}$ [7].

The transmitted intensity is thus:

$$I_t = |E_t|^2 = \frac{I_o T^2}{|1 - Re^{i\Delta\phi}|^2} = \frac{I_o}{1 + F\sin^2(\Delta\phi/2)} \tag{2}$$

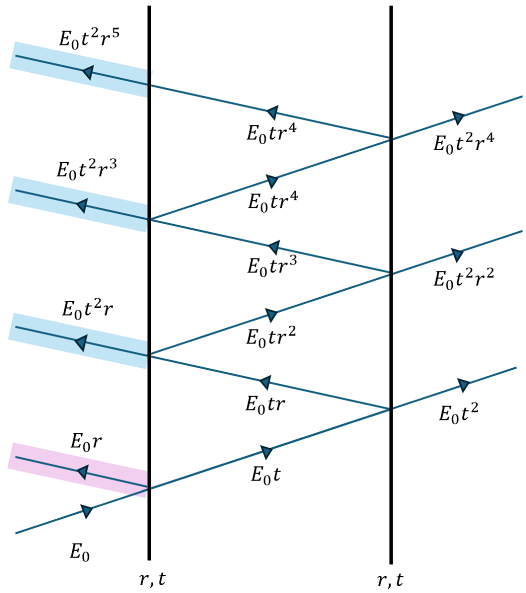

Figure 4: Simple Fabry-Perot cavity. Vertical black lines represent the cavity's surfaces. Dark blue lines represent the light's path and, assuming normal incidence, are drawn skewed to distinguish beams. Typically, light comes into the cavity at normal incidence, meaning the beams overlap. Notice that the electric field of the reflected and transmitted beams can be written as a geometric series, which allows us to derive the subsequent equations. The primary reflected beam is highlighted in pink while the leakage reflected beams are highlighted in blue.

where the coefficient of finesse, a property of the cavity, is $F = \frac{4R}{(1-R)^2}$.

The electric field of the reflected light is similarly given by

$$E_r = E_o r \frac{1 - e^{i\Delta\phi}}{1 - Re^{i\Delta\phi}} \tag{3}$$

with the intensity being:

$$I_r = |E_r|^2 = \frac{I_o F \sin^2(\Delta\phi/2)}{1 + F\sin^2(\Delta\phi/2)} \tag{4}$$

From the intensity plots, when the laser's frequency is in perfect resonance with the Fabry-Perot cavity (i.e. the laser frequency is an integer multiple of the cavity's free spectral range), the reflected intensity will be zero. This corresponds to a phase difference between the primary and leakage-reflected beams of 180 degrees, as shown in Fig. 4.

However, if there is a shift in wavelength, a nonzero reflected intensity will be measured. Because the intensity curve is symmetric about resonance, the derivative of the reflected intensity with respect to frequency must be used to determine if the laser frequency is above

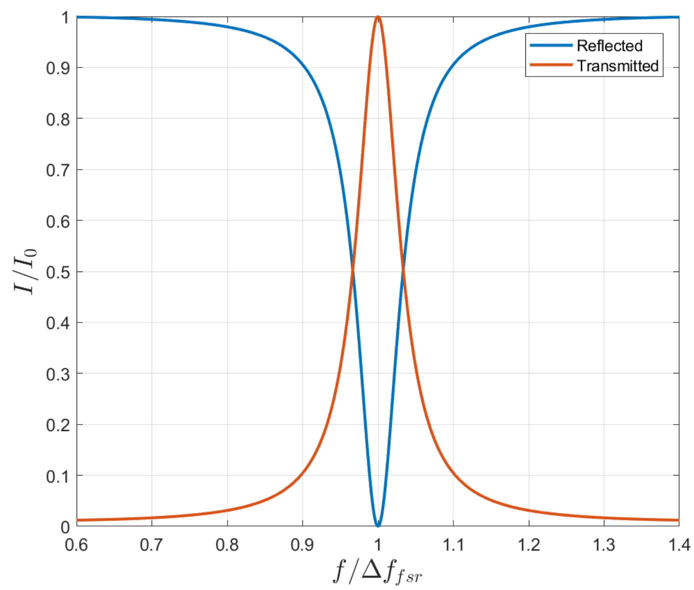

Figure 5: Fabry-Perot Intensity (transmitted/reflected beam). On resonance ($\frac{f}{\Delta f_{fsr}} = 1$), the transmitted intensity is at a maximum while the reflected intensity is zero.

or below the resonant frequency. This is done by measuring the phase difference between the primary and leakage reflected beams, which will constitute the error signal to be sent back to the laser [6].

To measure the phase difference between the reflected beam and leakage beams, the incident beam is first modulated by $\beta \sin(\Omega t)$, where β is the modulation depth, ω is the carrier frequency, and Ω is the phase modulation frequency. This gives the following incident field entering the Fabry-Perot cavity [6]:

$$E_i = E_0 e^{i(\omega t + \beta sin(\omega_m t))} \approx E_0(e^{i\omega t} + \frac{\beta}{2} e^{i(\omega+\Omega)t} - \frac{\beta}{2} e^{i(\omega-\Omega)t}) \tag{5}$$

The first term describes the carrier (laser) frequency while the second and third term constitute the upper and lower sidebands, respectively. These sidebands (frequencies above or below the carrier frequency) are a byproduct of modulation.

Using properties of Fabry Perot cavities, the transfer function below defines the relationship between the input (incident beam) and output (reflected beam) from the cavity.

$$R(\omega) = \frac{E_r}{E_i} = r(\frac{e^{i\Delta\phi} - 1}{1 - r^2 e^{i\Delta\phi}}) \tag{6}$$

This transfer function is applied to the incident light E_i to obtain the electric field of the reflected light, whose magnitude squared yields the power of the reflected light, P_r [6]:

$$
\begin{aligned}
P_r =& P_0|R(\omega)|^2 + P_0\frac{\beta^2}{4}(|R(\omega+\Omega)|^2 + |R(\omega-\Omega)|^2) + \\
& P_0\beta[Re(\chi(\omega))\cos(\Omega t) + Im(\chi(\omega))\sin(\Omega t)] + (\text{terms in } 2\Omega)
\end{aligned} \tag{7}
$$

where $\chi(\omega)$ and P_0 are defined to be $\chi(\omega) = R(\omega)R^*(\omega+\Omega) - R^*(\omega)R(\omega-\Omega)$ and $P_0 = |E_0|^2$. The terms involving $\chi(\omega)$ represent the interference of the sidebands with the original laser beam (also referred to as the carrier beam), informing us about the carrier's phase.

The photodetector measures a voltage V_r that is proportional to the reflected power P_r. Subsequently, a mixer multiplies the measured voltage with the modulated signal $\cos(\Omega t + \phi_m)$ to demodulate the beam and isolate the interference terms between the sidebands and carrier [6].

$$
\begin{aligned}
V_r' =& V_r\cos(\Omega t+\phi_m) \propto P_r\cos(\Omega t+\phi_m) \propto \\
& P_0[|R(\omega)|^2 + \frac{\beta^2}{4}(|R(\omega+\Omega)|^2 + |R(\omega-\Omega)|^2)]\cos(\Omega t+\phi_m) + \\
& P_0\beta[Re(\chi(\omega))\cos(\Omega t) + Im(\chi(\omega))\sin(\Omega t)]\cos(\Omega t+\phi_m) + \\
& (\text{terms in } 2\Omega)\cos(\Omega t+\phi_m)
\end{aligned} \tag{8}
$$

Setting the modulation phase ϕ_m to zero and selecting Ω such that $\chi(\omega)$ is purely real, the expression can be simiplified:

$$
\begin{aligned}
V_r' \propto & P_0[|R(\omega)|^2 + \frac{\beta^2}{4}(|R(\omega+\Omega)|^2 + |R(\omega-\Omega)|^2)]\cos(\Omega t) + \\
& P_0\frac{\beta}{2}\chi(\omega)\cos(2\Omega t) + P_0\frac{\beta}{2}\chi(\omega) + (\text{terms in } 2\Omega)\cos(\Omega t)
\end{aligned} \tag{9}
$$

The low pass filter isolates the constant term $P_0\frac{\beta}{2}\chi(\omega)$, which informs us of the phase difference between the reflected and leakage beam and therefore the side of resonance the laser frequency is on. Note that this will tell us whether the laser frequency is on the positive or negative side of resonance as $\chi(\omega)$ is an asymmetric function (Fig. 6). This error will be fed back into the laser to tune the frequency, completing the feedback loop.

In summary, the PDH technique consists of the following steps:

1. Send modulated light into a Fabry-Perot cavity.
2. Measure the reflected intensity and demodulate the signal using a mixer.
3. Extract the error term using a low-pass filter, which shows which side of resonance the laser wavelength is on.
4. Send the error into the laser to tune it back onto the cavity resonant frequency (i.e. desired frequency).

This method allows correction of the laser's frequency in real-time.

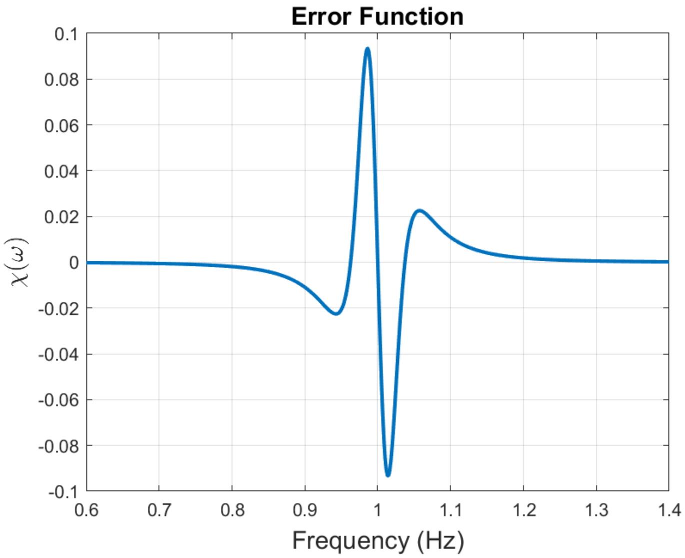

Figure 6: Real component of $\chi(\omega)$ over frequency.

2.3 Locking a Laser to a Vapor Cell

While optical cavities yield good results for laser locking at a wide range of frequencies, they are a relative, not absolute, reference. The cavity itself must be stabilized very precisely against mechanical shifts, temperature changes, material expansions, and other environmental effects that could change the properties or distance between mirrors [5]. They must also be properly coupled to the laser's Gaussian modes. This begets another way of performing laser locking, which is to use atomic transitions, an absolute reference. This method is not perfect as the atoms are affected by external electromagnetic fields and surrounding atoms, but overall, it permits more accurate locking due to the precise nature of atomic transitions [5]. In this specific application, the laser will be locked to a transition of ^{87}Rb because the laser will ultimately be used for interaction with ^{87}Rb in the MOT. In a well-constrained system where the atoms used to lock the laser have the same properties as those trapped in the MOT, the laser's interaction with the atoms will not change based on the laser locking scheme. If a cavity was used, changes in the laser's interaction with the MOT will occur if the cavity falls out of resonance.

First, consider a probe laser beam passing through a cloud of atoms. For simplicity, assume there are no external fields and there are only two energy levels (the ground and excited state). If the photons have resonant frequency ν_0, the energy of photons in the laser beam is equal to the energy difference between the ground and excited state and thus the atom is energized from the ground to the excited state. If the atoms do not move and only absorb photons at resonant frequency ν_0, absorption will only occur when the photons are all at resonant frequency ν_0. Shockingly, atoms do move, so they will absorb radiation and become excited even if the photon is not at the resonant frequency due to Doppler shifting [8]. For example, if the atom is moving towards the laser, photons with frequencies below the resonant frequency will get absorbed because the atom "sees" the photons as being of higher frequency. Similarly, if the atom is moving away from the laser, photons with frequencies above the resonant frequency will get absorbed because the atom sees the photons as being of lower frequency. This is what gives the absorption curve its line shape and spread. This effect is known as inhomogeneous broadening, where absorption is dependent on the frequency of the incoming photon and the velocity of the atom. In homogeneous broadening, the decay of atoms from the excited to ground state produces the same fluctuations in frequency, such as stimulated emission. Typically, the resulting lineshape is a convolution of homogeneous and inhomogeneous broadening profiles (Lorentzian and Gaussian, respectively) [3]. Other broadening causes include collisions with other atoms, limitations of the laser bandwidth, and residual Doppler broadening if the pump-probe beams are not exactly antiparallel [5].

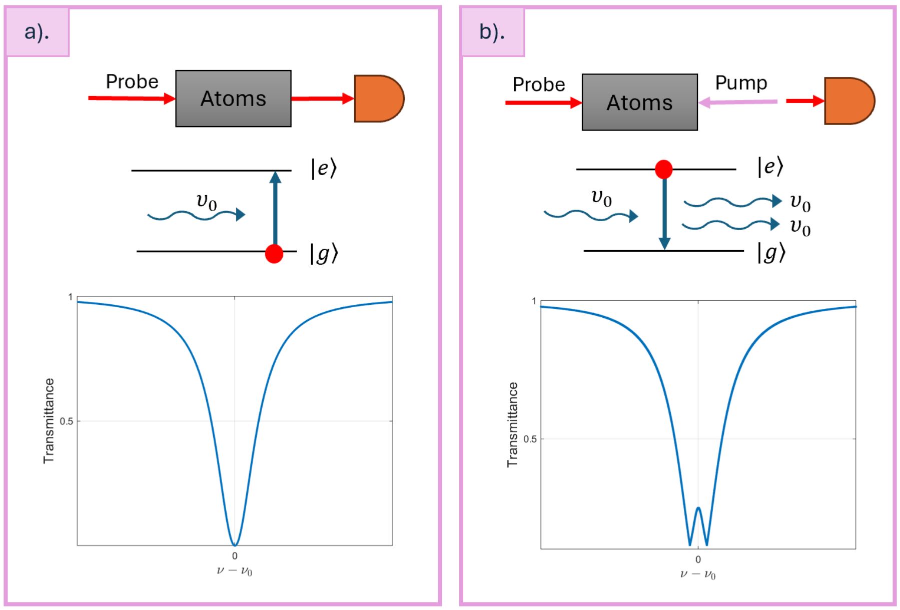

Figure 7: Graphs of the normalized transmittance as a function of laser frequency. a). Basic scheme of a laser passing through a vapor cell. If on resonance, the detector (orange) will measure no transmitted intensity because the energy provided by the incoming beam will be absorbed by all the atoms, energizing them to the excited state. The transmittance shows a Doppler-broadened lineshape. b). The addition of a pump beam will cause stimulated emission at resonance, which knocks the atom down to ground state and produces two photons. This produces a small, narrow peak in transmittance symmetric around the atomic resonance frequency, which has a smaller linewidth than the overarching curve in (a).

As with the PDH locking scheme, measuring the transmittance (in addition to measuring the derivative of the transmittance) will detail the amount of frequency detuning. However, in any locking scheme, a high-frequency discrimination (ability to distinguish between frequencies) is desirable [5]. To understand this, consider the graph in Fig. 7a. A narrower line width yields stronger frequency discrimination because for a fixed change in frequency, the change in intensity is steeper than that of a wider line width. Thus, to sharpen the Doppler broadening effect and thereby increase frequency discrimination, a pump beam is sent into the vapor cell anti-parallel to the probe beam.

The pump-probe works by first exciting atoms in the ground state with the probe beam. Then, the pump beam, which is at some potentially detuned frequency ν, will excite the atoms traveling with a certain velocity due to the Doppler effect, carving a "hole" in the population of atoms in the ground state (known as Lamb dips, where the dips refer to the dip in the ground state population) [3]. The pump is typically stronger than the probe because its higher intensity will produce a stronger peak, although the exact width of the Lamb dips are dependent on factors like beam diameter, surrounding temperature, and external fields [3]. Lamb dips do not reduce transmission to zero because the pump only excites atoms of a certain velocity and not the entire population of atoms: the pump simply causes less of the probe beam to be absorbed.

The pump and probe are antiparallel to take advantage of this Doppler effect as off-resonance, the pump and probe beams excite different populations of atoms (traveling at different velocities in opposite directions) and thus act independently of each other. The probe will then primarily experience absorption [3]. However, if the probe and pump are near resonance, both beams will interact with atoms at zero velocity. As a result, those atoms will be excited by either the pump or probe beam. If the atom is excited by the pump, the atom is rendered "transparent" to the probe, which will enable it to transmit through the medium. As a result, a peak in transmittance will occur when both beams are on resonant frequency [3].

Another way of thinking about this is that when the probe beam passes through, stimulated emission occurs at resonance. If the pump has excited an atom to the excited state, the probe will knock the atom back down to the ground state, releasing one photon that travels coherently with the stimulating photon toward the detector. If off-resonance, the probe will experience more absorption from exciting an atom moving with certain velocity based on Doppler shift.

As evidenced in Figure 7, it is not possible to determine which side of resonance the laser frequency is on due to the symmetry of transmittance peak around resonance. To find the error signal within this narrow range, a similar process of error term extraction as with PDH locking must be performed. Modulating the probe beam produces a carrier frequency ν_c and sideband frequencies of $\nu_c \pm \Omega$. Upon passing through the nonlinear atomic medium, the previously unmodulated pump beam will become modulated, and the pump and probe will engage in four-wave mixing (FWM) [9]. FWM occurs when two or more beams indirectly interact within a nonlinear medium (in this case, the vapor cell), producing a new frequency due to scattering of incident photons. Typically, this occurs with three frequencies, which produces another frequency that is some sum/difference of the three initial frequencies. In this case, two pump photons (at some frequency ν) and a probe photon (at frequency $\nu \pm \Omega$) interfere and produce a probe photon at sideband frequency $\nu \mp \Omega$. The interference between the probe and these sideband photons creates a modulation signal which is received by the photodetector [9].

The signal is demodulated to generate an error signal as shown in Fig. 8, which is a measure of the phase difference between the probe and its generated FWM sidebands (in PDH locking, this was the interference between the carrier and sidebands of the reflected beam). This

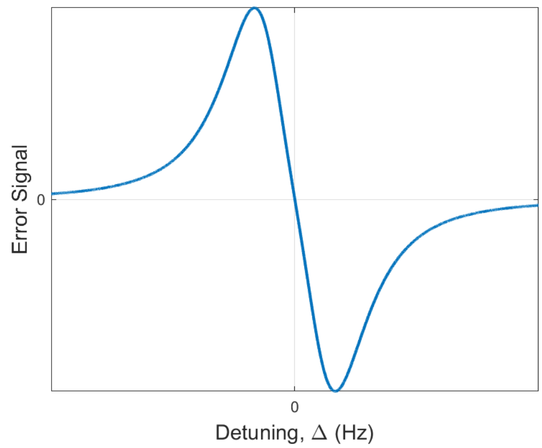

Figure 8: Example of error signal from frequency spectroscopy method. Note the asymmetry just like the error signal from PDH locking.

signal is given by [9]:

$$S(\Omega) = \frac{C}{\sqrt{\Gamma^2+\Omega^2}} J_0(\delta) J_1(\delta)((L_{-1} - L_{-1/2} + L_{1/2} - L_1)\cos(\Omega t + \phi) + (D_1 - D_{1/2} - D_{-1/2} + D_{-1})\sin(\Omega t + \phi)) \tag{10}$$

where $L_n = \frac{\Gamma^2}{\Gamma^2+(\Delta-n\Omega)^2}$ and $D_n = \frac{\Gamma(\Delta-n\Omega)}{\Gamma^2+(\Delta-n\Omega)^2}$, where Γ is the linewidth and Δ is the amount of frequency detuning [9].

Similar to PDH locking, Ω and ϕ are chosen such that the signal is maximized and only the cosine or sine term of the error signal is left. The resulting error signal is asymmetric, as for PDH, which gives information about which side of resonance the laser frequency is on. This error can be fed back into the laser input to correct frequency drifts.

In general, the primary advantage of using this modulation method (also known as modulation transfer spectroscopy, or MTS) is that the zero-crossing point of the error signal corresponds to resonance peaks and is relatively independent of polarization, temperature, and intensity in comparison to other methods [5]. Some limitations apply and should be considered. For example, because modulation frequencies are smaller than natural linewidth, there is a bandwidth limit to the locking scheme. This also means that the sidebands can be affected by intensity fluctuations in the absorption dip because they occur at frequencies where there are still contributions from the absorption dip. MTS also only applies for atoms with cycling transitions, meaning that if the atom emits a photon going from excited to ground state, sending another photon of the same frequency in will energize it back to the same excited state. Finally, there is a limitation in the region of dithering because, as shown in Fig. 7b, there is additional symmetry outside of the transmission peak. If the frequency is too far from resonance, the first-order derivative will no longer satisfy the degrees of freedom in the graph and the laser will not be locked.

The objective of this project is to lock a 780 nm laser to a vapor cell of ^{87}Rb atoms and verify the apparatus' ability to reduce frequency drifts and noise in the system.

3 Methods

To lock the lasers, the schematic in Fig. 9 was built. Fig. 10 shows the tabletop setup that implemented laser locking with the probe and pump beam. A vapor cell of the ^{87}Rb isotope was chosen to lock the laser to. Aside from its intended use in the MOT, ^{87}Rb is also advantageous because of its heavier molecular mass, which narrows atomic Doppler widths, and well-characterized properties.

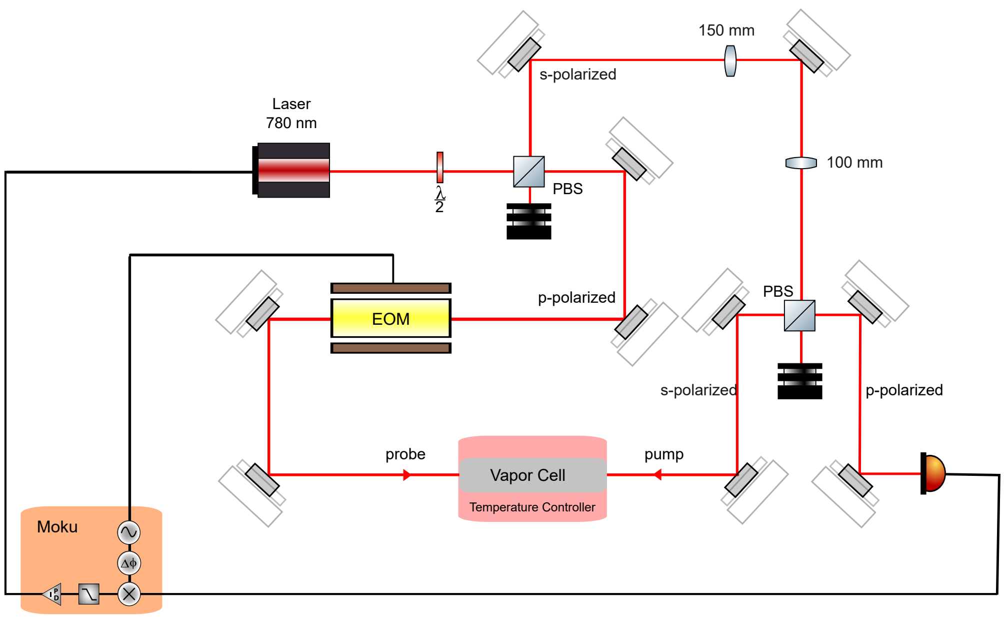

Figure 9: Schematic of laser locking to the ^{87}Rb vapor cell. The 780 nm laser beam was of sufficiently small size to enter and exit the EOM's 2 mm aperture so additional lenses were not needed.

The 780 nm distributed Bragg reflector laser (DBR780PN from Thorlabs) was used for the setup. The power and wavelength of light from the laser are adjusted by a temperature and current controller. By holding the current constant and scanning over the laser temperature, the dominant absorption dip of ^{87}Rb can be located, which corresponds to about 780 nm as expected [10]. The temperature and current controllers are ultimately the actuators that control the laser's wavelength.

The laser is split into the probe and pump beam using a polarized beamsplitter immediately after the first half waveplate. Rotating the half waveplate in front of the beamsplitter adjusts the power of the probe and pump beam. The pump and probe beam initially have

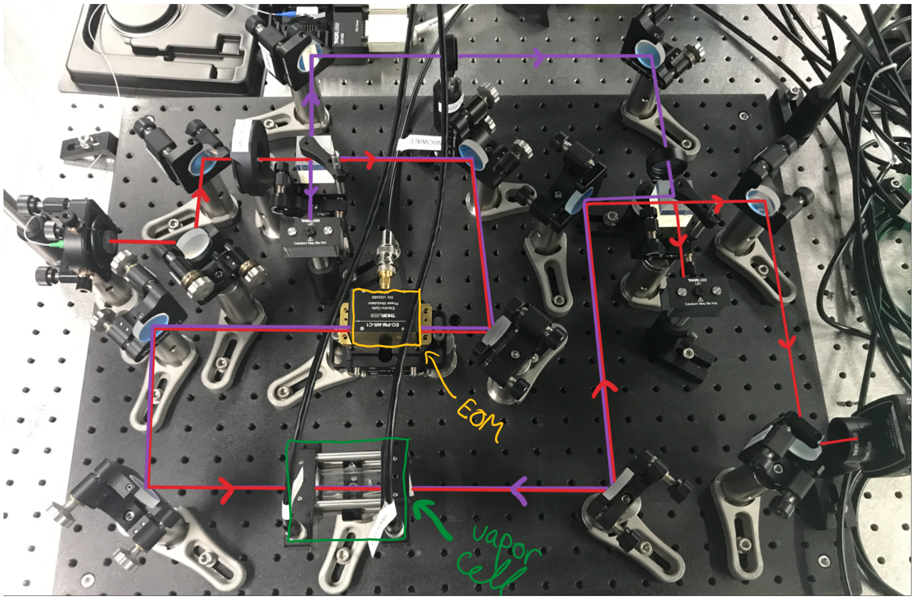

Figure 10: Experimental setup in the lab. The purple and red paths are the pump and probe beam respectively.

the same frequency and detuning as they are derived from the laser output beam. The probe beam then undergoes phase modulation by an electro-optic modulator (EOM), which changes the phase of the beam and produces the sidebands discussed in the previous section. After passing through the vapor cell, the intensity of the probe will be measured by a photodetector. The probe's transmittance determines the distance from resonance via the demodulation sequence, which produces an asymmetric error signal that is fed into the laser. The counterpropagating pump beam overlaps with the probe beam as much as possible using mode-matching. Using the JamMt software and a Nanoscan Photon beam profiler, two lenses (150 mm and 100 mm) were added to the pump path to reduce the beam size. By concentrating the pump to match the probe, the absorption dips obtained are sharper.

The vapor cell is heated using a Thorlabs TC300 temperature controller whose PID controller was auto-tuned to maintain the vapor cell at a stable temperature to reduce temperature fluctuation effects. The Thorlabs PDA10A2 photodetector was used to detect the probe beam as it has a sufficiently high bandwidth (150 MHz) to detect higher frequency signals. Once the signal is measured, the Moku Pro from Liquid Instruments (a field-programmable gate array, or FPGA) is used to perform the demodulation and error signal extraction as well as implement control on the laser's temperature and current controller. In the Moku Pro, the parameters of the phase modulator (such as the EOM phase and frequency) are tuned to maximize error signal and, consequently, frequency discrimination.

A PID controller was used to implement both slow and fast control. The controllers send control signals to the laser's current and temperature controller. This PID controller must be optimized, which was done through trial and error by checking which parameter produced the lowest noise spectra, which will be discussed in Sec. 3.6.

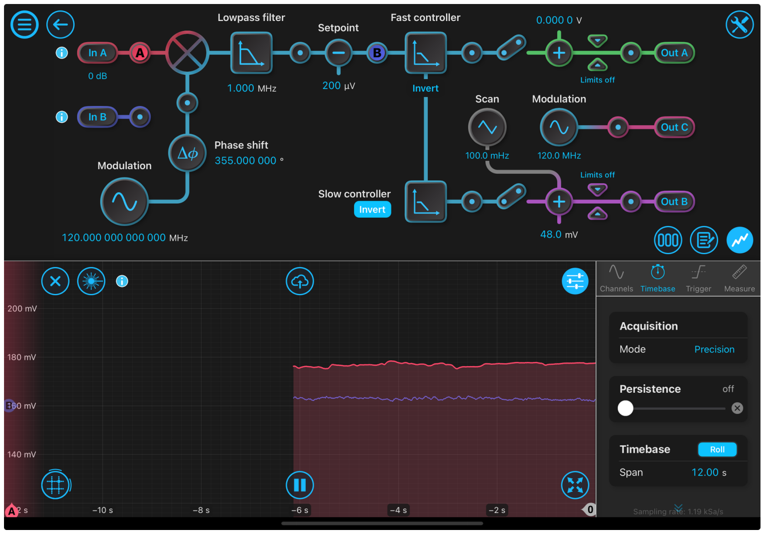

Figure 11: Moku Pro setup for laser locking scheme. Notice the demodulation process with the error term being extracted at setpoint B and the subsequent fast and slow controller control signals being output back to the current and temperature controller, respectively.

3.1 Maximizing Error Signal

Prior to locking the laser, the absorption dip strength must be characterized to maximize the error signal. Maximizing the error signal will increase the slope of the linear region of the asymmetric error signal, as shown in Fig. 12. Notice that with a greater slope, a small detuning will result in a larger change in the error signal within the linear region. This refines sensitivity to frequency shifts, improving the system's frequency discrimination abilities. Sec. 3.2 and 3.4 demonstrate two parameters optimized to maximize the error signal.

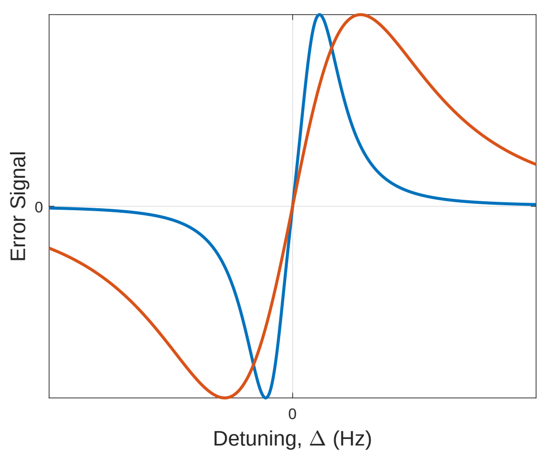

Figure 12: Example of two error signals, one with a higher slope in the linear region (red) and one with a lower slope in the linear region (blue). The error signals are normalized but it intuitively makes sense that the slope is greater for an error signal of greater amplitude.

3.2 Effect of Vapor Cell Temperature

In this setup, a laser passes through the vapor cell and into a photodetector. As a note, the lower bandwidth photodetector (Thorlabs PDA36A2) was used at this point because higher frequency signals did not need to be detected.

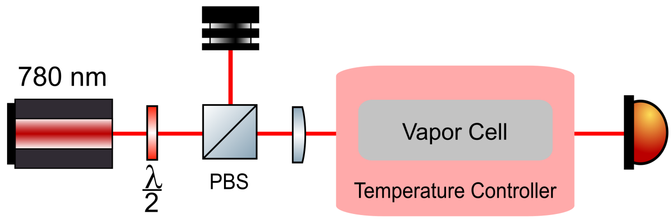

Figure 13: Simplified setup for experiments maximizing dip depth. The lens was used for changing beam size and was not included for tests on the effect of vapor cell temperature on dip depth.

The absorption dip depth was obtained by scanning over laser wavelength for different vapor cell temperatures. The effect of vapor cell temperature is useful to know because to increase the absorption dip depth, as many atoms as possible should be excited into the ground state ($5s_{1/2}$) for the D2 780 nm transition by initially heating the vapor cell. For each vapor cell temperature, the laser power will be tuned to find the optimal power that excites the most atoms into the $5p_{3/2}$ transition state (i.e. produce the greatest dip depth). These two variables can maximize the strength of the absorption dip. Upon performing the experiment, at different currents, the dip depth is larger at a higher temperature as can be seen in Fig. 14, so the vapor cell was operated at 40 degrees Celsius as this is a high temperature below the operating limit of the vapor cell heater (50 degrees Celsius). Maintaining a stable temperature will pack the states, ensuring as many atoms are in the $5^2S_{1/2}$ level as possible [10]. Note that while the $5^2S_{1/2}$ is the true ground level of ^{87}Rb, a higher temperature can excite the atom to the necessary hyperfine level needed for the transition being locked to. Another reasoning for this phenomenon is that a higher temperature increases the atomic density of the vapor, which increases the likelihood of absorption [11].

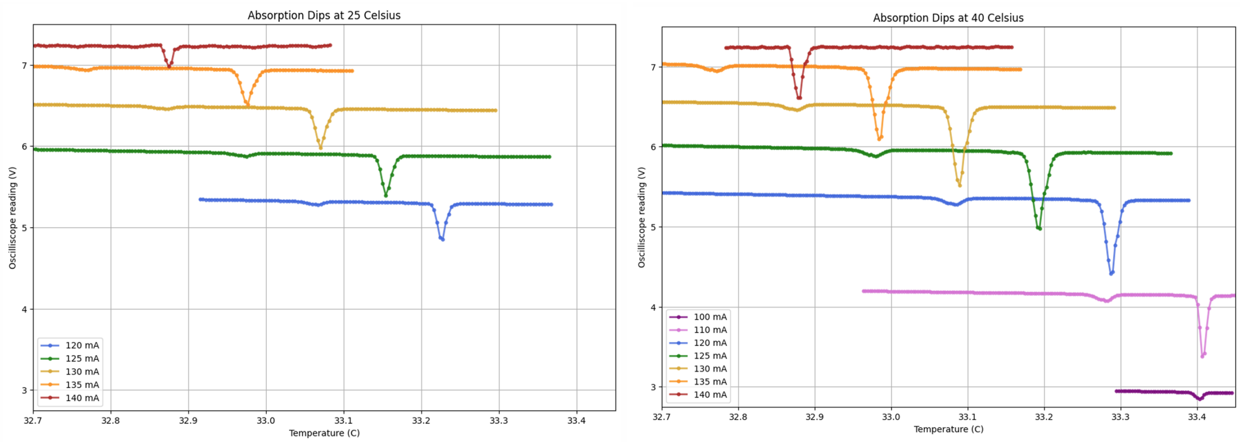

Figure 14: Absorption dips at a chosen temperature. For consistency, the dips corresponding to the rising section of the temperature scanning signal were plotted. Notice that at different currents, the absorption dip is greater for the vapor cell at higher temperatures. At 25 degrees Celsius, the 100 mA and 120 mA trials were not performed as the dips were not clearly visible. Thus, in Fig. 15, only the dip depth for currents of 120, 125, 130, 135, and 140 mA are plotted.

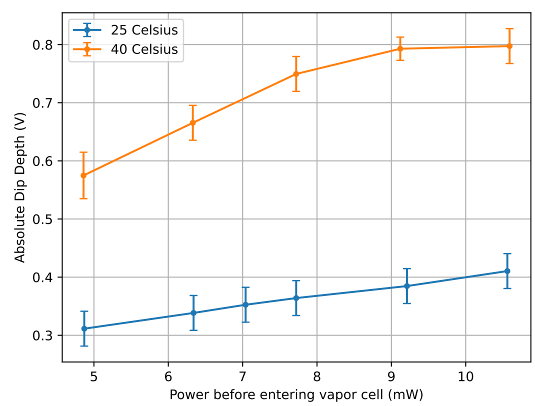

Figure 15: Plotting the absolute dip depth from Fig. 14 for both temperatures. These values were extracted by subtracting the linear background bracketing the absorption dip (from temperature scanning) and finding the subsequent minimum value of the absorption dip. Error bars were established by how far the background fit deviated from the data points.

3.3 Effect of Laser Beam Size

The effect of beam size through the vapor cell on the absorption dip depth was also tested. The method for testing how beam size changes throughout the vapor cell was as follows. First, the beam was profiled. The laser power should not affect the beam profile but for consistency, the laser current was set fixed at 150 mA. The position of the lens and vapor cell were set using the JamMt software to the desired beam size. The positions of where the vapor cell begins/ends and where the lens will be located was recorded. A beam profile was taken at the beginning of the vapor cell without the lens in place and the location recorded. The vapor cell was then put in place. A beam profile after the vapor cell without the lens was measured and the location recorded. The location of the photodetector was determined by ensuring the beam profile at that location was within the photodetector's aperture size. The photodetector was placed at that location. A wavelength scan was performed with the lens not in place. Afterward, the lens was put in place. Then it was confirmed that the light falling onto the photodetector was the same as without the lens by checking alignment. A wavelength scan was performed with the lens in place. Afterward, a beam profile was measured at the location of the photodetector (keeping the lens in place), which also required removing the photodetector. Then, the beam profile was measured after the vapor cell and recorded. Finally, the vapor cell was removed and a beam profile was measured right before the lens.

As a note, the laser beam was beam profiled with a 150 mm lens to produce a reasonably small beam profile that did not converge to a waist near the vapor cell. A half-waveplate and polarizing beamsplitter was used for power adjustment. The power incident on the vapor cell was measured with a Thorlabs PM100D photodetector.

The results from this experiment suggest that a larger beam size increases the absorption dip depth. Thus, lenses were not used to reduce beam size through the vapor cell. That a smaller beam size reduces dip depth intuitively seems reasonable because photons will interact with a lower number of atoms, so less absorption occurs. In contrast, a larger beam means photons are more likely to be absorbed because they interact with more atoms. Once you reach a high enough power, you would be saturating the atoms by bombarding them with photons at a rate faster than the lifetime of the excited state. This phenomenon of eventual decreasing absorption with increasing optical field power is known as saturation [3, 12].

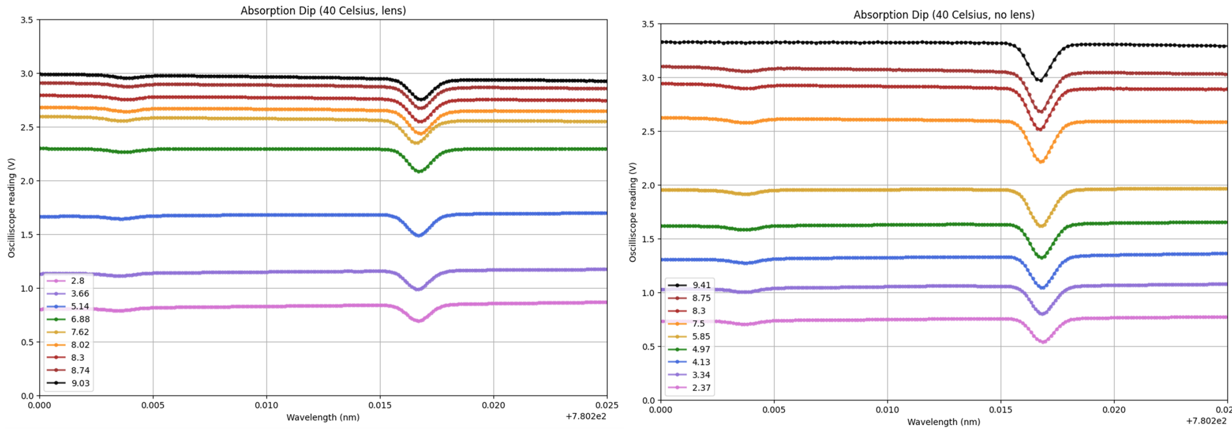

Figure 16: Dip depth at a set current for varying power (see legend for power in mW). Visually, the smaller beam size (left, with lens) appears to have a lower dip depth compared to the larger beam size (right, without lens). The smaller and larger beam size were 500-1200 micron and 1400-1500 micron respectively.

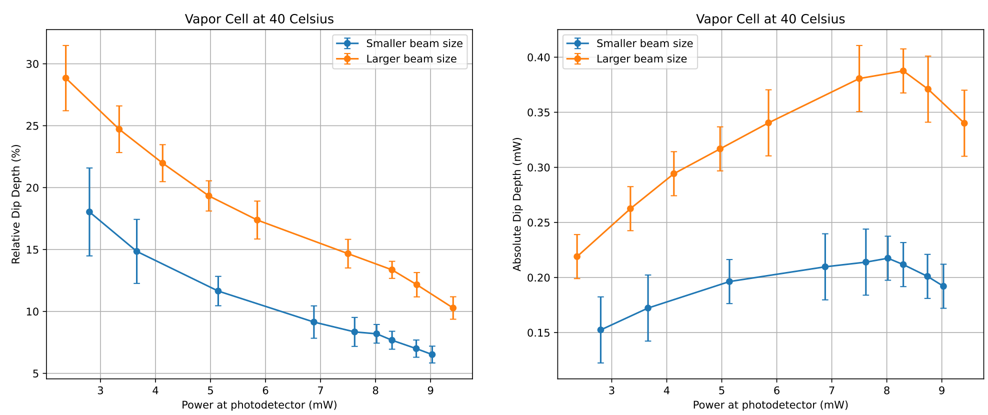

Figure 17: Comparison of the relative dip depth and absolute dip depth at different incident powers for different beam sizes.

3.4 Locking to a Single Absorption Dip

After choosing these parameters of vapor cell temperature and beam size to maximize the error signal, an initial lock to the vapor cell was performed using just the probe beam. The scheme for this is found in Fig. 18.

When scanning over the laser's wavelength range, which produces the expected absorption dip from Fig. 7 and modulating the signal with an 80 MHz signal, the asymmetric error signal seen in Fig. 25 and Fig. 20 is obtained. The phase shift and modulation frequency were tuned to further maximize the error signal. Within the Moku Pro, a zero-crossing point can be chosen to lock to in the error signal shown in Fig. 19. The location of the zero point of the error signal corresponds to the minimum of the absorption dip as seen in Fig. 21.

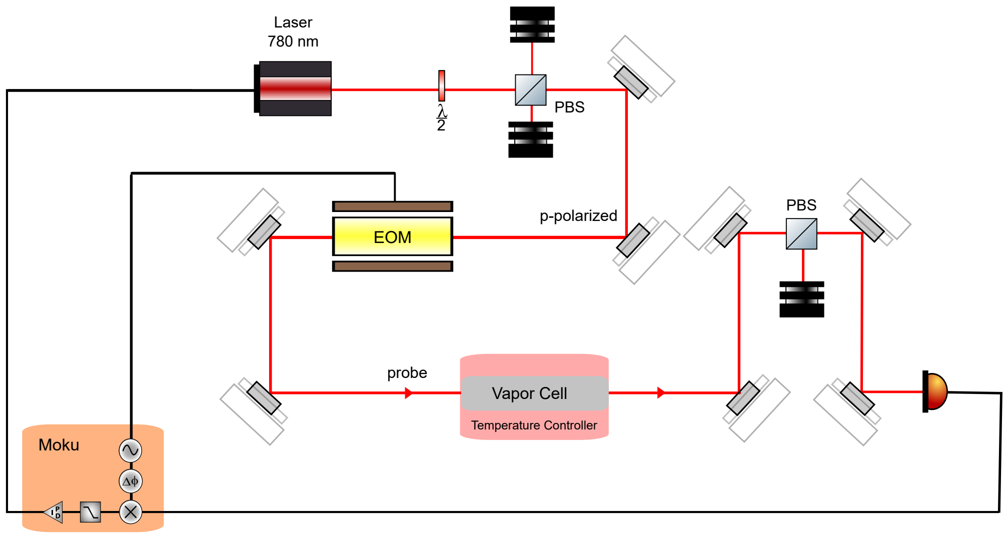

Figure 18: Schematic of laser locking to just the probe beam, which passes directly through the vapor cell.

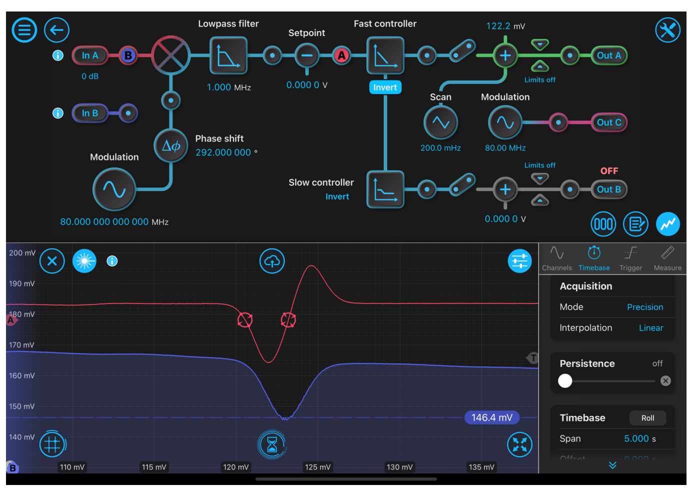

Figure 19: As seen in red, the error signal is produced from modulation. The absorption dip after scanning is shown in purple.

As shown in Fig. 22, the laser can be confirmed as locked by verifying that the photodetector signal is stable around its minimum value (corresponding to the zero-crossing point) and that the error signal begins to deviate more once the feedback loop is broken (i.e. the controllers are disconnected from the system).

After locking, the error signal is first measured for 30 seconds. Then, the feedback loop was broken by disconnecting the slow and fast controllers. The error signal value was then

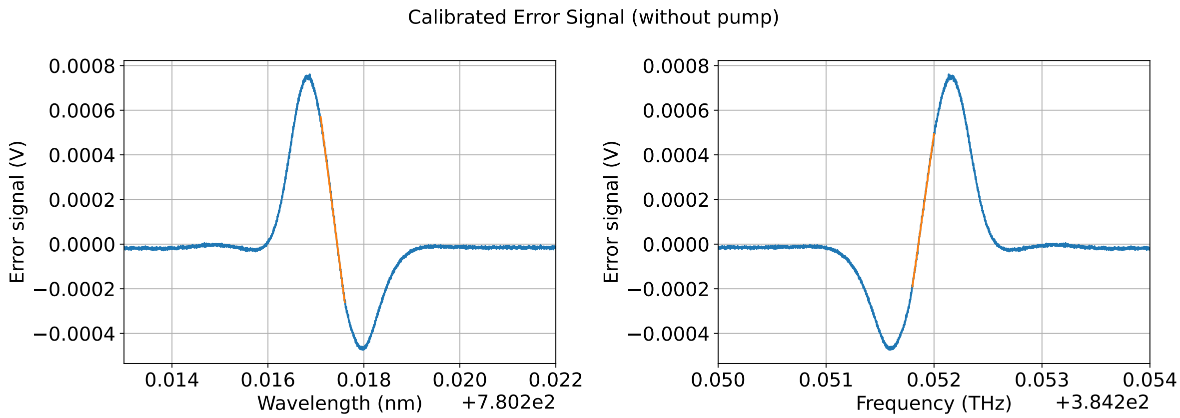

Figure 20: Error signal with wavelength (left) and frequency (right) axes (calibrated according to the DBR780PN laser specifications from Thorlabs). As will be discussed later, the orange line on the two graphs shows the portion that were used for a linear fit to convert from V^2/Hz to $Hz/\sqrt{Hz}$ when comparing noise spectra across different setups.

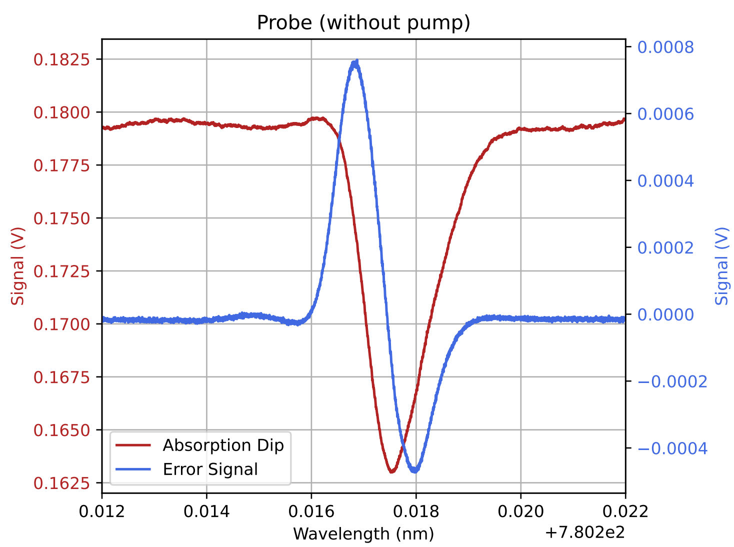

Figure 21: Overlaid absorption dip and error signal (without pump) from scanning wavelength.

measured for 30 seconds. These values are converted into a power spectral density using the Welch method. As can be seen in Fig. 23, which shows the noise spectrum with just the probe beam, the system has greater noise when the laser is unlocked compared to when the laser is locked. This matches the expectation that noise is reduced due to frequency drift by locking the laser.

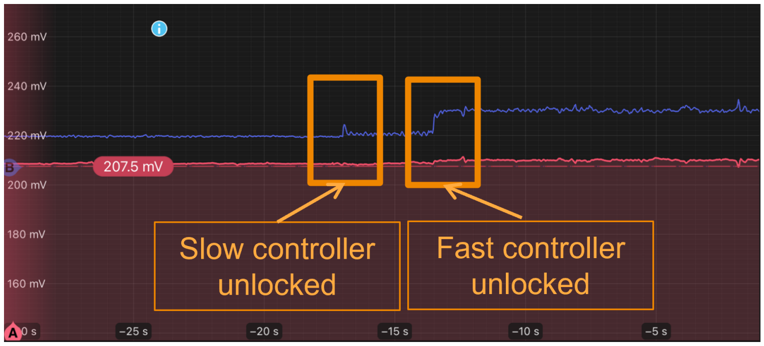

Figure 22: Error (purple) and photodetector (red) signal as lock is broken. Initially, the photodetector signal is located near the minimum value seen from the absorption dip. Note that when the fast and slow controller are unlocked, the error signal becomes much noisier and the photodetector signal deviates more from the zero-point locking value.

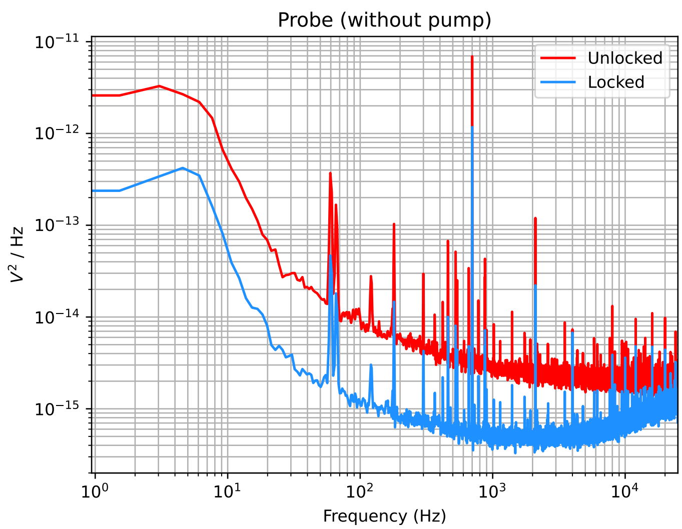

Figure 23: Noise spectrum for the locked and unlocked laser without the pump beam.

3.5 Locking to a Saturated Absorption Dip

The counterpropagating pump was added to the setup shown in Fig. 9. The same procedure as above was repeated to lock and obtain the noise spectrum. Notice that the photodetector signal contains new peaks due to the pump. The multiple peaks likely correspond to different hyperfine transitions and subsequent crossover resonances. Because these peaks are also apparent in the error signal, the error signal has multiple linear regions that cross the zero point. The laser was locked the zero-crossing of the steepest slope to produce the greatest frequency discrimination. However, a tradeoff with this is that if the amplitude of the linear region is not large enough, a small detuning in frequency will break the lock as the frequency will exit the linear regime. Thus, locking with the pump has been more difficult than without the pump and various improvements can alleviate this pain, as will be discussed in Sec. 4.

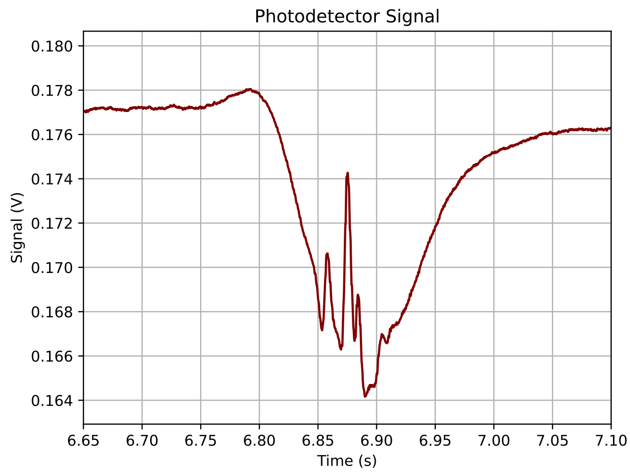

Figure 24: Photodetector signal with the inclusion of the counterpropagating pump beam. Note the new peaks in the absorption dip caused by the pump.

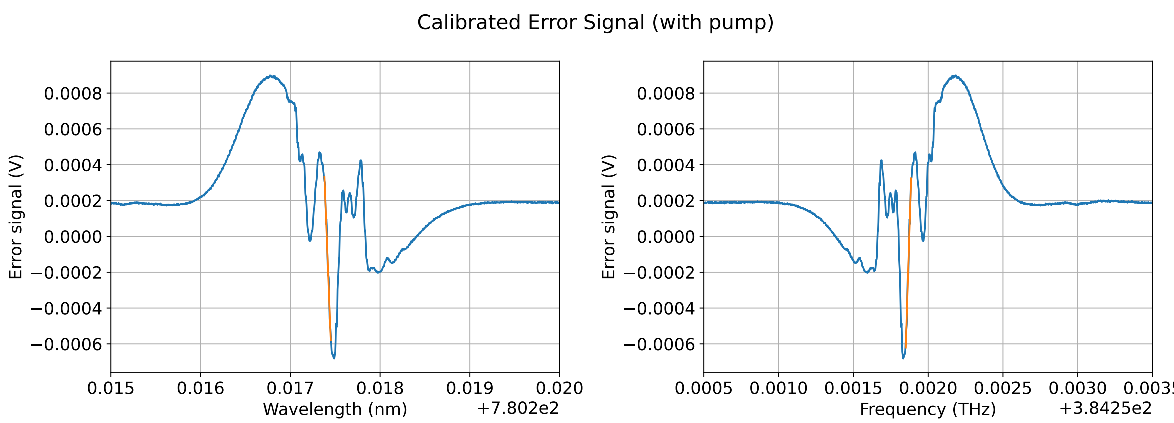

Figure 25: Error signal with wavelength (left) and frequency (right) axes (calibrated according to the DBR780PN laser specifications from Thorlabs). The orange line on the graphs shows the portion that is chosen to lock too as it has the largest slope. This slope will be used to convert from V^2/Hz to $Hz/\sqrt{Hz}$ when comparing noise spectra.

From the noise spectrum in Fig. 26, there is less noise in the locked laser than when the laser is unlocked.

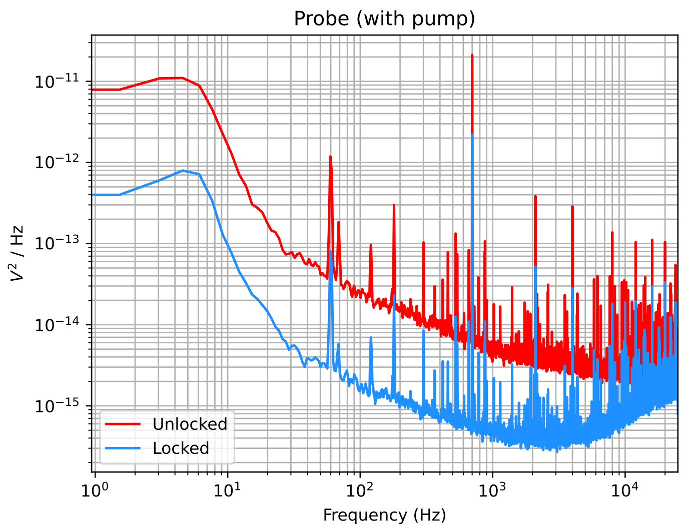

Figure 26: Noise spectrum for the locked and unlocked laser with the pump beam

To compare the noise spectra of the probe with the pump, the probe without the pump, and the unlocked laser, the noise spectrum must be calibrated. The PSDs between Fig. 26 and Fig. 23 are not comparable in their current power units V^2/Hz as the different voltage readings are based on different detunings from resonance. Thus, the slope of the error signal is used, which corresponds to the linear region of the signal, to convert the voltage into frequency units (e.g. $Hz/\sqrt{Hz}$), which gives us frequency noise. Because this slope passes through the zero crossing point of resonance, only the slope is needed as the "linear equation" in terms of detuning is slope $\times$ $\Delta\lambda$, where $\Delta\lambda$ is the distance from resonance). The linear region is used because beyond the two asymmetric peaks, the controller will continuously push the laser frequency away from resonance (in the linear region the slope at the same error signal value would have the opposite derivative compared to the value outside of the peaks, so it pushes the frequency towards resonance instead of away).

After calibration, the noise spectra comparison in Fig. 27 is obtained. This suggests that locking with the pump yields the lowest noise in the system and can maintain the laser frequency with the greatest stability. This is followed by the locking without the pump, which still exhibits less noise than the system when the laser is unlocked. These results are promising as a proof-of-concept that the probe-pump scheme of laser locking to a vapor cell can successfully reduce frequency drift and noise in the system.

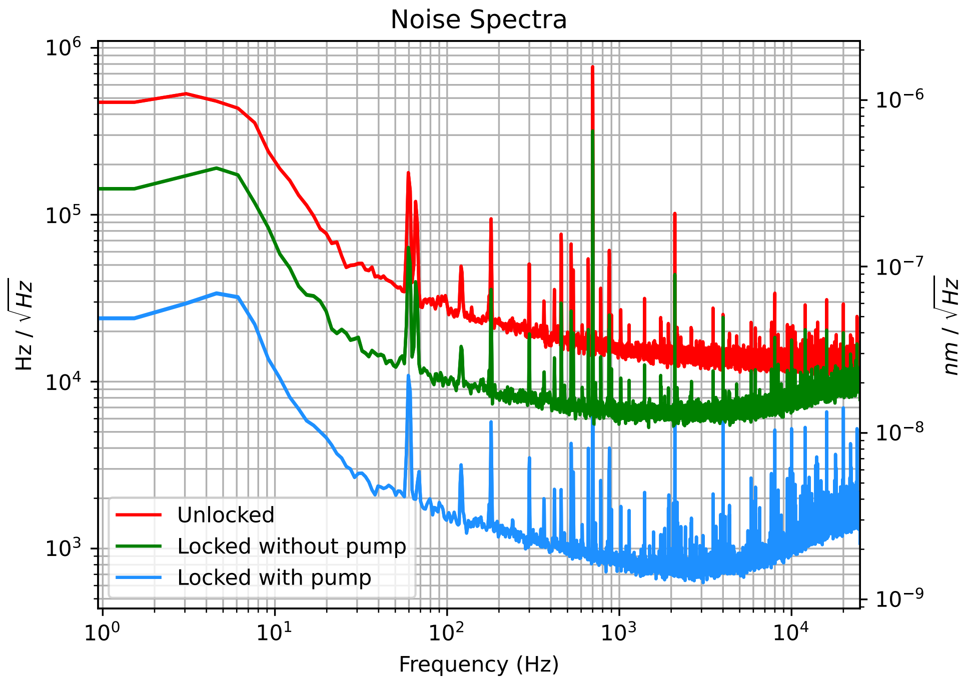

Figure 27: Noise spectra (calibrated) comparison when the laser is unlocked, locked without the pump, and locked with the pump.

3.6 Controller Optimization

To determine the optimal parameters for the slow controller to reduce noise, manual trial and error was performed. The PSD of the locked and unlocked error signal over the timescale of about two minutes was taken to see which controller properties reduced noise the most. In this case, the slow controller corresponded to the slow actuator, which was laser temperature.

Fig. 28 shows a few of the parameters tested for the slow controller in an effort to optimize the controller. Note that there is a peak around 40 Hz introduced by the locked laser, likely due to the controller. If the unity gain frequency and the gain of the integrator controller are reduced, this peak gets damped down, which makes sense because the gain will be lower at that frequency. This peak could be a result of the driven 40 Hz frequency being in resonance with something in the system. This can be thought of in terms of a vertically oscillating spring. If the spring is driven at the right frequency, the amplitude will increase significantly. Anything at lower or higher frequencies will have a lower amplitude than the resonant amplitude. Thus, the resonant peak is introduced because the controller is driving some resonance in the system.

Observing these noise spectra again indicates that a fast controller and actuator are necessary to affect the noise spectra at higher frequencies (>1 Hz). Trying to increase the bandwidth (integrator's unity gain frequency) of the slow controller to extend the frequency regime that the controller acts in is ultimately ineffective because the actuator itself is slow. Thus,

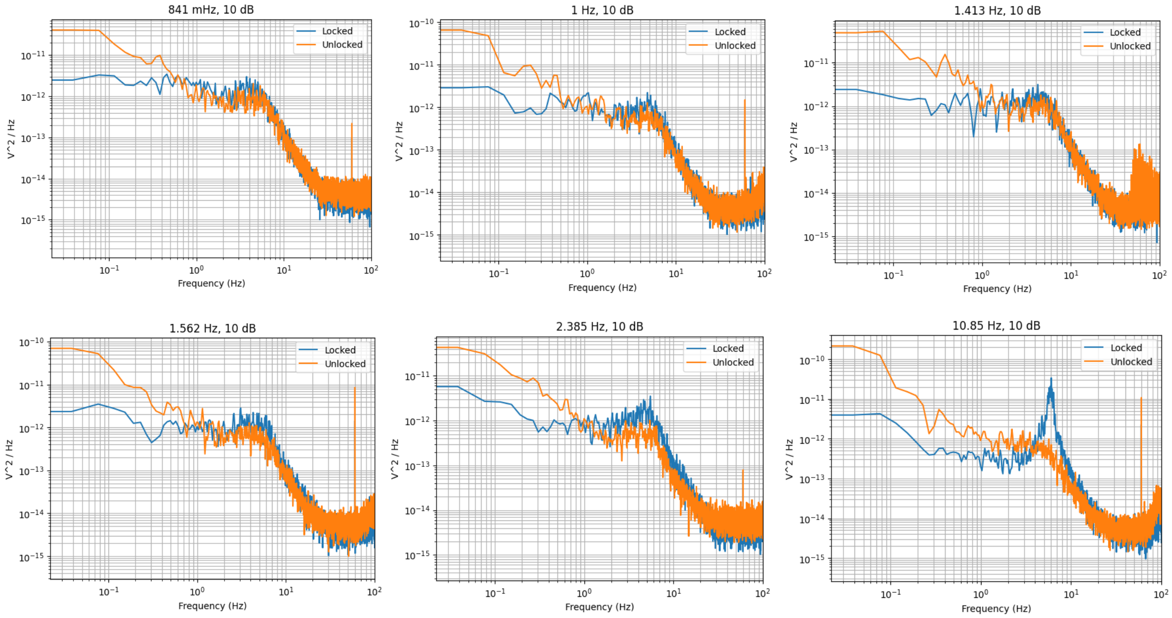

Figure 28: Various noise spectra taken for different slow controllers. The first number in the graph title is the unity gain frequency for the integrator control and the second number is the initial gain of the controller.

using a current controller as a fast actuator allows us to get the improved noise spectrum shown above that removes noise at higher frequencies. The fast controller was not optimized because of time constraints.

The controller parameters used in the end for optimal performance were an integrator unity gain frequency of 995.9 mHz and gain of 6.2 dB.

4 Future Work

These results are a good sanity check but more optimization and testing are necessary before they can be fully trusted. In particular, the parameters for the slow and fast controller must be further optimized because for certain gains and corner frequencies, the locked laser can actually display greater noise than the free-running laser.

The laser wavelength has to be sufficiently close to the absorption transition to lock fully, so manual tuning is sometimes required. To avoid the issue of lock being broken for a small detuning, which was especially prominent when including the pump beam, the error signal should be amplified with a low-noise amplifier and tank (resonance) circuit. This will make it easier to lock the laser and maintain lock.

There are additional issues within the system. First, after placing in a polarizing beamsplitter and half waveplate after the 780 nm laser to control power (without changing the temperature and current controllers), the laser fluctuated in power over time, as seen in Fig. 29. The source of the issue is the polarization maintaining fiber, which is not functioning properly. In this fiber, the laser input needs to be aligned with the slow axis, but because the input fiber to the DBR780PN laser does not seem to be adjustable, another solution is needed. Overall, these intensity fluctuations will not only add noise to the system but also affect the robustness of the lock.

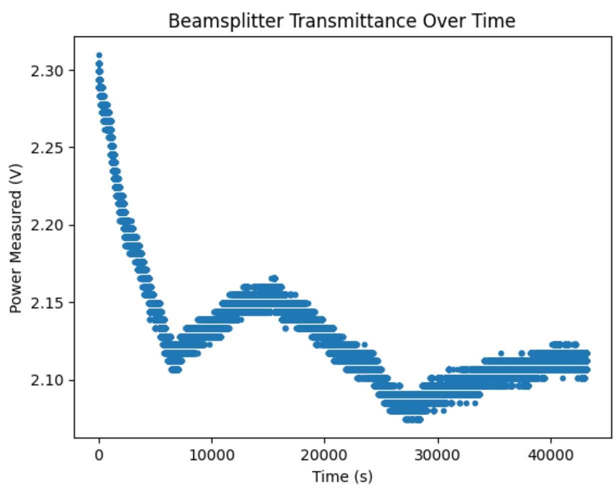

Figure 29: Drift in transmitted signal through the polarizing beamsplitter cube overnight. Discretized values are due to the longer time between taking data points and the Moku Pro's signal range setting.

In the vapor cell schematic, an additional beamsplitter should be added to control the laser power, whereas the subsequent beamsplitter will control the power of the probe and pump beam. Additional materials such as a half waveplate are needed to optimize efficiency of the EOM. A Faraday isolator, used to prevent unwanted reflected light from reentering the laser, will also need to be added due to the affected polarization from these waveplates, which prevents the p and s-polarization lights from being beam dumped entirely at the second polarizing beam splitter (PBS). This means that the ideal setup would be as shown in Fig. 30.

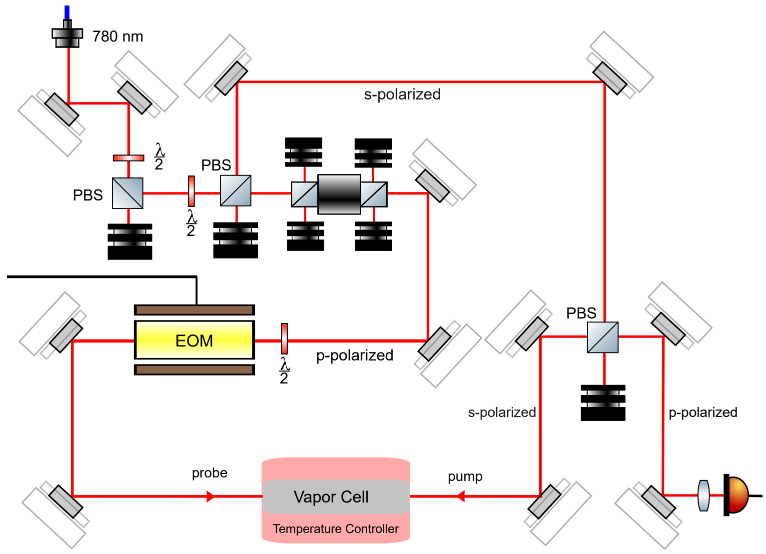

Figure 30: Schematic of laser locking to the rubidium vapor cell. A half waveplate in front of the EOM would be added to maximize the EOM efficiency. A lens is added in front of the photodetector to ensure the beam is enclosed in the active area of the photodetector.

Another issue is the varying dip depth when scanning in increasing or decreasing laser temperature, which can be attributed to the changing timescale of cooling the laser rather than heating it as well as uncertainty in the laser response time. This shift in the location of the absorption dip was dependent on the direction of scanning (i.e. scanning by increasing or decreasing wavelength). Over a slower scan frequency (10 mHz), this shift was not as notable, so scans were performed consistently at 10 mHz for the dip depth experiments 31. However, when implementing laser locking, scans were done at a frequency on the order of 100 mHz for efficiency, so the offsets in the location of the absorption dip could affect the location of the zero-crossing point over time.

Finally, to expand the laser locking to a 1560 nm laser, second harmonic generation (SHG) will be implemented, which uses a nonlinear crystal to double the original laser frequency. This process produces 780 nm light, which will be used for locking. The resulting control signal will then be fed back into the 1560 nm laser. The setup would then be split into two separate sleds for the SHG and the vapor cell portions respectively. This is for ease of transportation and swapping laser sources as a 780 nm laser can be easily connected to the vapor cell portion of the SHG. SHG would require extensive mode-matching due to the size and numerical apertures specified by the nonlinear crystal, which will determine the needed focal length of the lens such that the beam becomes narrow enough to enter the crystal.

Although the power spectral density is currently to quantify the effectiveness of the lock in reducing frequency drift, the success of the locking mechanism can also be quantified using the Allan variance, which measures the long-term variance in the laser frequency and displays different noises in the system, which are distinguishable by the power laws followed over time [5]. It would additionally be interesting to extract the phase shift term from the error signal to demonstrate where the sidebands originate from.

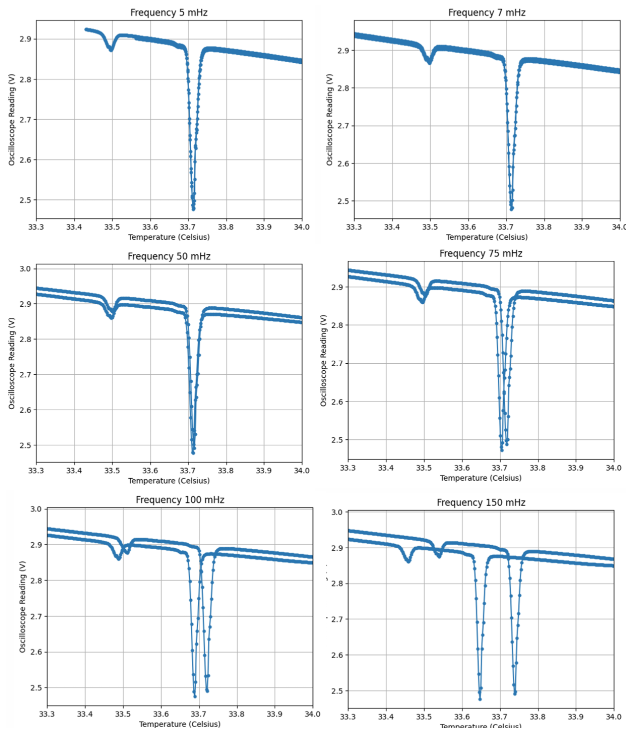

Figure 31: Notice that when scanning temperature at higher frequencies, the temperature (and consequently, calibrated wavelength) at which the absorption dip occurs changes significantly based on whether the dip corresponded to whether the ramp signal (i.e. temperature control) was decreasing or increasing. For lower-frequency scans, this difference is minimized.

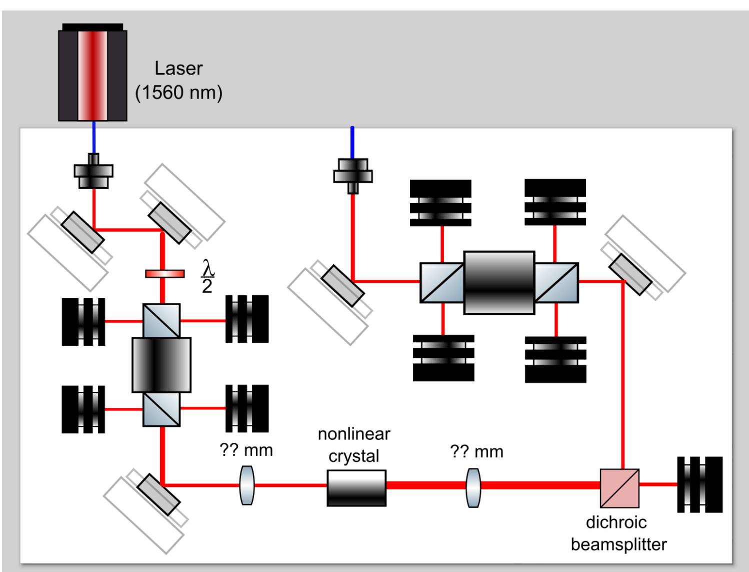

Figure 32: Second harmonic generation (SHG) schematic.

Acknowledgements

Thanks to my mentors Ian Macmillan and Lee McCuller and the members of the McCuller lab for all their help (Torrey Cullen, Daniel Grass, Alex Ramirez, Sander Vermeulen, Jeffrey Wack). I would also like to thank the National Science Foundation and the Caltech LIGO SURF program for this opportunity.

References

[1] S. M. Vermeulen, T. Cullen, D. Grass, I. A. O. MacMillan, A. J. Ramirez, J. Wack, B. Korzh, V. S. Lee, K. M. Zurek, and C. Stoughton, "Photon counting interferometry to detect geontropic space-time fluctuations with gquest", arXiv preprint arXiv:2404.07524 (2024).

[2] K. Kowalski, C. Van, D. Khoa, M. Głódź, B. Nguyen, and J. Szonert, "Magneto-optical trap: fundamentals and realization", Computational Methods in Science and Technology **Special Issue**, 115–129 (2010).

[3] C. J. Foot, *Atomic physics* (Oxford University Press, 2005).

[4] *Nif's guide to how lasers work*, `https://lasers.llnl.gov/education/how-lasers-work`.

[5] M. Mausezahl, M. F., and R. Low, "Tutorial on laser locking techniques and the manufacturing of vapor cells for spectroscopy", American Journal of Physics **69**, 79–87 (2001).

[6] E. D. Black, "An introduction to pound–drever–hall laser frequency stabilization", American Journal of Physics **69**, 79–87 (2001).

[7] J. Peatros and M. Ware, *Physics of light and optics* (Brigham Young University, 2023).

[8] U. of Florida Department of Physics, *Saturated absorption spectroscopy*, 2020.

[9] D. McCarron, S. King, and S. Cornish, "Modulation transfer spectroscopy in atomic rubidium", Measurement Science and Technology **19**, 105601 (2008).

[10] D. A. Steck, *Rubidium 87 d line data*, available online at `http://steck.us/alkalidata` (revision 2.1.4, 23 December 2010.

[11] A. Kaewko, P. Parkprom, K. Ketaiam, P. Chaiwongkhot, N. Chattrapiban, P. Sompet, and N. Thaicharoen, "The effect of temperature on absorbance of rubidium atoms in a vapor glass cell", Journal of Physics: Conference Series **2653**, 012036 (2023).

[12] R. W. Boyd, "Chapter 6 - nonlinear optics in the two-level approximation", in *Nonlinear optics (third edition)*, edited by R. W. Boyd, Third Edition (Academic Press, Burlington, 2008), pp. 277–328.