UNIVERSITY OF COPENHAGEN
FACULTY OF SCIENCE

Locking an Optical Cavity using a Microcontroller

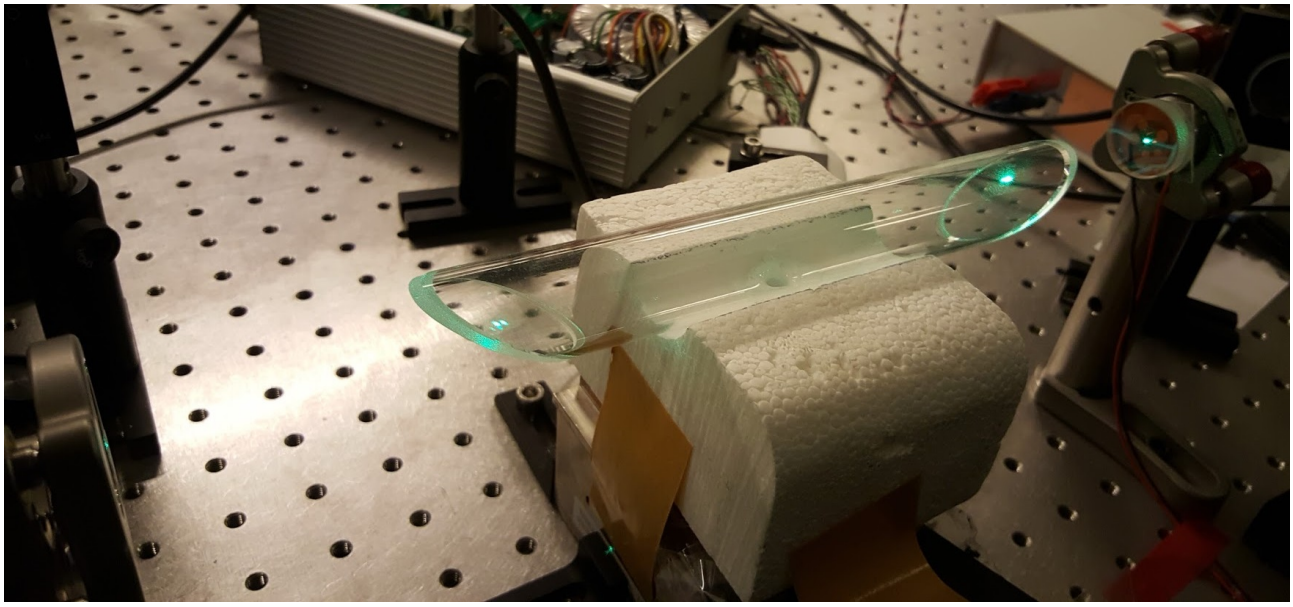

Bachelor's Thesis in Physics

Adam S. Knorr & **Ida L. Stoustrup**

Date: 16/1 - 2018

Locking an Optical Cavity using a Microcontroller

Prepared by:
Adam S. Knorr & Ida L. Stoustrup

Academic supervisor:
Jan W. Thomsen

Lab supervisors:
Asbjørn A. Jørgensen
Martin R. Henriksen

Quantum Metrology and Ultra Cold Atoms Group
Niels Bohr Institute
University of Copenhagen
Blegdamsvej 17
København Ø
Denmark

Acknowledgements

We would like to thank all the people who helped us produce this thesis.

Jan for inspiring and motivating us to learn optics, and bearing with us even when we would clear the front page of "Noob times".

Arvad for always being there to answer our questions, for many, many hours of guidance and for spending his scarce free time correcting this thesis.

Martin for sharing his hard-earned, magical lab insights with us, for fixing our set-up when we would mess it up and for helping us in the quest for the Lamb dip.

Axel for performing black magic on our electronics when they misbehaved.

And of course the group in general - we had a blast doing this project!

Honourable mentions go to:
Prof. Stoustrup, for (re)teaching us control theory.
Jimmy, for performing first aid on Ida's computer when she thought it might like to taste some of her tea.

Abstract

Stable and precise frequency standards are important in many areas of science. A molecular clock using iodine has potential to become cost-efficient, compact and simple and could be used in e.g. spacecraft, where these qualities are essential. In this thesis, a Teensy 3.6 microcontroller has been set up and programmed to act as a PID controller. The microcontroller has successfully locked an optical cavity to a narrow linewidth laser. Saturated Absorption Spectroscopy was performed on an iodine cell placed in the cavity. A small Lamb dip was observed, which proves the controller's usability. The Allan deviation and Fourier transform of the in-loop measurement of the system performance, revealed a large resonance, which must significantly limit the quality of the servo system. In future work on the project, the effect of the resonance could be reduced to significantly improve system stability.

Resumé

Stabile og præcise frekvens standarder er vigtige i mange områder af videnskaben. Et molekylært jod-ur har potentiale til at blive billigt, kompakt og simpelt, og kunne bruges i for eksempel rumfartøjer, hvor netop disse egenskaber er vigtige. I dette projekt er en Teensy 3.6 microcontroller sat op og programmeret til at fungere som en PID controller. Microcontrolleren har succesfuldt låst en optisk kavitet til en laser med smal linjebredde. Satureret absorptions spektroskopi blev benyttet på en jodcelle inde i kaviteten. Et lille Lamb dip blev observeret, hvilket beviser controllerens brugbarhed i projektet. På Allan afvigelsen og Fourier-transformationen af en lukket løkke måling af systemet, kan man se en stor resonans, hvilket betydeligt må begrænse servo systemets kvalitet. I fremtidigt arbejde på projektet kan resonansens indflydelse begrænses for at forbedre systemets stabilitet.

List of Figures

List of Abbreviations

ADC: Analogue to Digital Converter
DAC: Digital to Analogue Converter
EOM: Electro-Optic Modulator
FSR: Free Spectral Range
FWHM: Full-Width Half-Maximum
FWTM: Full-Width Tenth-Maximum
GUI: Graphical User Interface
OP-AMP: Operational Amplifier
PD: Photo Detector
PDH: Pound-Drever-Hall
PID: Proportional Integral Differential (In regards to control theory)
SAS: Saturated Absorption Spectroscopy

Contents

1 Introduction

Mankind has been using frequency references since the first apes looked to the Sun to determine the day and night cycle. As they evolved, time measurement evolved with them. From directly measuring natural phenomena, like sundials, to the man-made frequency standards like the pendulum clock, mankind was on the hunt for the best way to define and measure time. For centuries we improved our methods of measuring frequencies, until 1955 where Louis Essen produced an atomic clock from a transition in Caesium-133, which was the most accurate frequency standard the world had seen [2]. This atomic clock was used to define the second. Since then, atomic clocks and quantum metrology have been focused on creating the best frequency reference possible. However, it is not simple to define what certifies the best clock. Some atomic clocks (e.g. strontium and yttrium clocks) can become extremely precise and accurate, while the experimental set-up is quite complex and requires a lot of work to function. Other clocks (e.g. acetylene and iodine), do not require the same amount of upkeep but they can not become as precise. When designing an atomic clock, trade-offs between different qualities are always made, these being: Precision, accuracy, stability, maintenance, compactness and cost. Iodine is especially good when it comes to compactness, maintenance and cost. When comparing these qualities to how good precision, accuracy and stability it can achieve, it blows the competition out of the water. This is in large due to iodine being a gas at room temperature. This allows us to just fill a glass tube with iodine gas, which means we do not need a vacuum, a magneto optical trap, a Zeeman slower etc. for our set-up to function.

To achieve the relatively good precision, accuracy and stability with the iodine set-up, we want to lock a lasers frequency to a hyperfine transition in iodine.

To optimize the laser's interaction with the iodine, an optical cavity is used. However optical cavities only allow very specific frequencies to enter, so it is necessary to control the cavity length to allow the laser light to enter.

In this thesis, we will focus on locking the optical cavity length to the laser wavelength. To lock the cavity we are using a Teensy microcontroller to run a digital PID control loop. Once the cavity is locked we will calculate the stability of the cavity using Allan deviation, and test if the lock is good enough to find these hyperfine transitions.

2 Theory

2.1 The Optical Cavity

The basic principle of an optical cavity is trapping light at a specific frequency between two highly reflective mirrors causing it to build a large electromagnetic (EM) field.

To assemble an optical cavity one aligns two highly reflective mirrors, M_1 and M_2 with reflective and transmission coefficients $(t_1, t_2 < r_1, r_2)$, facing each other a distance L apart (see fig 1). Light inside the cavity will reflect off the mirrors, which flips the phase of the light. If the phases do not align throughout the cavity the light will destructively interfere. Since the light is reflected multiple times, this destructive interference reduces the occurrence of non resonant frequencies to practically zero. The waves that constructively interfere must be those with a wavelength being a multiple integer of the cavity length, since their phases overlap upon reflection, i.e. standing waves. The frequency difference between two such standing waves is called the Free Spectral Range (FSR) and can be calculated: $\Delta k = 2\pi/L \Rightarrow FSR = c/2L$ [3] where Δk is the difference in wave vector of the two waves and c is the speed of light. In this thesis we use an optical cavity with a length of L = 30 cm giving us an FSR of roughly:

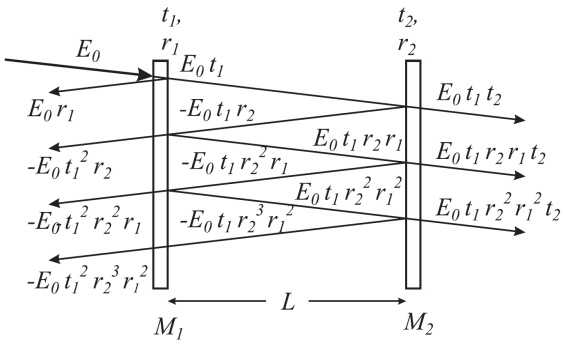

Figure 1: A close up of two cavity mirrors M_1 and M_2 a distance L apart, and with (r_1, r_2) and (t_1, t_2) as their respective reflection and transmission coefficients. A resonant electrical field E_0 interrogating the cavity shows the field being reflected, transmitted and built up. Note that this cavity is not aligned, to easily display the light-mirror interactions. When the cavity is aligned all the lines would be place on top of each other. [1].

$$FSR = \frac{c}{2L} \approx \frac{(3 \cdot 10^8\ m/s)}{2 \cdot (0.3\ m)} = 500\ \text{MHz} \tag{1}$$

The allowed frequencies have a shape of a Lorentzian in frequency space and the FWHM of this resonance curve is called the optical linewidth of the cavity:

$$\delta\nu = \frac{(1 - r^2)}{\pi r} \frac{c}{2L} \tag{2}$$

The ratio between the FSR and the optical linewidth is called the finesse of the cavity (F). The finesse is used to determine how well the cavity is aligned, and

the amount of loss in the cavity.

$$F = \frac{FSR}{\delta\nu} = \frac{\pi r}{(1-r^2)} \tag{3}$$

To study the reflection and transmission from the optical cavity, when interrogated by a laser beam at resonance, we approximate the beam as an infinite plane wave: $\bar{E}(t,\bar{r}) = E_0 \exp(i(\omega t - \bar{k}\bar{r}))$. When aligned with the cavity, the beam will interact with M_1 with transmission and reflection coefficients r_1 and t_1. The transmitted beam is then reflected and transmitted in M_2 (with coefficients r_2 and t_2). The beam is reflected between the two mirrors transmitting a bit of power each time it interacts with a mirror. The full interaction has been illustrated in fig. 1.

In this thesis we will be using an optical cavity with two identical mirrors where $r = r_1 = r_2$ and therefore $t = t_1 = t_2$ Following the idea in fig. 1 we can then write the total reflected field (left side of M_1) as:

$$E_R = E_0 r - E_0 t^2 r e^{-i\omega 2L/c} - E_0 t^2 r^3 e^{-i\omega 4L/c} - E_0 t^2 r^5 e^{-i\omega 6L/c}... \tag{4}$$

A more rigorous derivation for mirrors with different reflectivity can be found in: [1]. Rewriting the expression leads to:

$$E_R = E_0 r - E_0 t^2 r e^{-i\omega 2L/c}\left(1 + r^2 e^{-i\omega 2L/c} - r^4 e^{-i\omega 4L/c}...\right) \tag{5}$$

where we realize that we can use the geometrical sum: $\sum_{n=0}^{\infty} q^n = \frac{1}{1-q}$ with $q = r^2 \exp(\frac{-i\omega 2L}{c})$ and then arrive at:

$$E_R = E_0 r - E_0 r\left(\frac{t^2 \exp(-i\omega 2L/c)}{1 - r^2 \exp(-i\omega 2L/c)}\right) \tag{6}$$

Extending the first term of the right side with 1 in terms of: $1 - r^2 \exp(-i\omega 2L/c)$:

$$E_R = E_0 r\left(\frac{1 - (r^2 + t^2) \exp(-i\omega 2L/c)}{1 - r^2 \exp(-i\omega 2L/c)}\right) \tag{7}$$

Using that $1 = r^2 + t^2$ allows us to calculate the reflection coefficient r_0, the ratio of light reflected from the cavity.

$$r_0(\omega) = \frac{E_R}{E_0} = r\left(\frac{1 - \exp(-i\omega 2L/c)}{1 - r^2 \exp(-i\omega 2L/c)}\right) \tag{8}$$

Likewise we can calculate the transmission coefficient. The derivation is somewhat similar. Realizing we can write the terms of transmission as a geometrical sum, and calculating the ratio between the incoming and transmitted light we find:

$$t_0(\omega) = \frac{E_T}{E_0} = \frac{t^2 \exp(-i\omega L/c)}{1 - r^2 \exp(-i\omega 2L/c)} \tag{9}$$

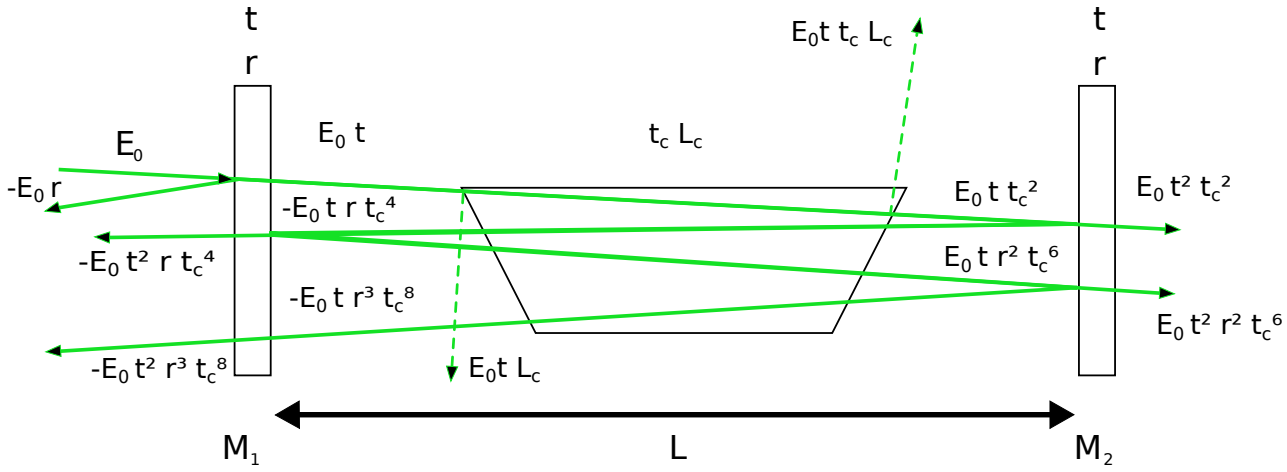

Figure 2: A linear cavity with an inter-cavity iodine cell. The iodine cell introduces loss (L_c) in the shape of both its glass windows reflecting light out of the cavity, and the iodine gas absorbing and spontaneously emitting the light. All light propagating though the cavity is introduced with a transmission coefficient(t_c) on each side of the cavity. L_c and t_c are defined such that $1 = L_c^2 + t_c^2$ for each of the cells surfaces.

In our set-up we will be working with an iodine cell inside the cavity. A cell is a glass tube of a pure pressurized substance, in this case iodine. However introducing two new surfaces and a gas into the cavity introduces new reflections and losses. In our case the cell's windows are Brewster angled, meaning the surfaces are angled such that maximum transmission is obtained. This also means that all reflections are pointed out of the cavity leading to loss. Another loss factor is the iodine gas. The laser we use is (at times) on resonance with the iodine atoms causing us to lose some light due to absorption immediately followed by spontaneous emission. Spontaneous emission happens in random directions and with random phase, so we can note this effect as loss from the cell. We combine these two loss factors to a single loss coefficient of the cavity surfaces L_c. We also define the transmission coefficient tc such that $1 = L_c^2 + t_c^2$ for each surface (see fig. 2). We can now look at the cavity transmission and reflection again:

$$E_{Rc} = E_0 r - E_0 t_c^4 t^2 r e^{-i\omega 2L/c} \left(1 + r^2 t_c^4 e^{-i\omega 2L/c} - r^4 t_c^8 e^{-i\omega 4L/c} ...\right) \tag{10}$$

And once again using the geometric sum we obtain:

$$r_C(\omega) = \frac{E_{Rc}}{E_0} = r\left(\frac{1 - t_c^4 \exp(-i\omega 2L/c)}{1 - r^2 t_c^4 \exp(-i\omega 2L/c)}\right) \tag{11}$$

Following this example we obtain the transmission coefficient with a cell as well:

$$t_C(\omega) = \frac{E_T c}{E_0} = \frac{t^2 t_c^2 \exp(-i\omega L/c)}{1 - r^2 t_c^4 \exp(-i\omega 2L/c)} \tag{12}$$

In addition to the the change in reflection and transmission coefficients, we also see a change in the cavity linewidth and finesse. The introduction of the iodine

gas effectively changes the reflection coefficient (r) in eq. 2 and 3. The change is proportional to the pressure of the iodine gas. The pressure can be indirectly controlled by the temperature of the cell. Reducing the temperature of the cell condenses more of the gas, reducing the pressure. Reducing the pressure reduces loss due to absorption, which increases the finesse, and reduces the linewidth.

Throughout this section we have used straight lines to depict the cavity mirrors and laser light, which is a good approximation in many cases. However, in reality laser light propagates as a Gaussian beam due to diffraction. To account for this property most optical cavities use curved mirrors to match the beam's wave front. In addition to matching the wave front one also has to consider matching the mode of the beam. If the laser beam is diverging while entering the cavity the wave front will not match the curvature of the mirror. For a Gaussian beam the waist is defined as the point in space where the light is most focused. In a cavity such as ours, it is very important to have the waist placed in the middle of the cavity to match the mode best.

2.2 PID Control theory

In the previous section, we saw that the cavity only allows light to enter, if the length of the cavity is an integer multiple of the wavelength of the laser light. Controlling a mirror to always be a set distance from another on a scale of less than a nanometer is impossible to do by hand, so we are using a piezo-electric crystal. Piezo-electric crystals expand when a voltage is applied, and the change in their length is small enough to allow us to adjust the length of the cavity on the nanometer scale. However, we still need to know what voltages to apply. Using a photodetector (PD) and the cavity we can scope out how much voltage equals an FSR. But we need to be able to control the piezo so that when it is disturbed in one direction, it will be pushed right back into place.

Imagine we have some electrical error signal that shows how far we are from resonance with the cavity. We then want to feed the piezo a control signal that is *Proportional* to this error signal, so if the error signal is large we will apply a large voltage to the piezo to move the mirror back into place. The mathematical expression would be: $P(t) = k_P Err(t)$ where P is the proportional control signal sent to the piezo, k_p is the proportional constant, and $Err(t)$ is the error signal.

Using only this technique one would soon realize that the error signal will oscillate around the set-point. This is due to the proportional control overshooting the set-point. In our case if the piezo is moving towards the set point, it can not instantly terminate its velocity at the set point. This makes the piezo overshoot - making the control and error signal oscillate around the set value.

To overcome these unwanted oscillations we could look at what the previous

values of the error signal were and add that to the signal. We do this by adding an *Integral* term *I(t)* to the proportional term: $I(t) = k_I \int_0^\tau Err(t)d\tau$ that is slowly going towards the right value. This also allows us to follow any changes to the offset of the control signal which is important to counteract drift.

However, this is not always good enough, if we want our mirror to be locked in place. If we were able to look at how fast the error signal is moving in a given direction, we would be able to counteract big disturbances better. Therefore we add a *Differential* term to our calculations of the control signal: $D(t) = k_D \frac{d}{dt} Err(t)$.

These three components together constitute PID control and the complete control signal becomes:

$$S(t) = P(t) + I(t) + D(t) = k_P Err(t) + k_I \int_0^\tau Err(t)d\tau + k_D \frac{d}{dt} Err(t) \quad (13)$$

2.2.1 Analog or digital PID

A PID controller can be either analog or digital.

Analog controllers produce continuous signals and are mainly limited by the operational amplifiers (OP-AMP) needed to build them. They are however prone to drift on a day to day basis, and if the lock disengages, it has to be manually re-engaged. The drift means that many analog systems must be tuned every day which can be quite time consuming.

Digital PID control is mainly limited by the clock frequency of the processor and the speed and resolution at which analog signals can be converted to digital (and vice versa). What makes the digital controller desirable is the amount it can be customised to a specific system. A digital servo system can automatically re-engage, optimize PID constants, function as an oscilloscope and much more. Some of these functionalities are time-saving and some of them can improve the servo system. Of course, every extra utility occupies the processor, so depending on the hardware of the microcontroller, there might be a limitation to how much functionality can be added, without making the corner frequency lower than the required one.

2.3 Pound-Drever-Hall Locking Technique

To use the PID control mentioned in the previous section we will have to generate an error signal. Unfortunately we cannot simply look at the amplitude of either the reflection or the transmission of the cavity, although these will depend on how close the laser is to a cavity resonance (see eq. 8 and 9). It is possible, but it requires somewhat advanced techniques because two cavity length deviations of the same size will produce the same error, even if they have

different signs. Thus we want a simple way to know which way to change our cavity length. There is such a technique, and it is named Pound-Drever-Hall (PDH) after the scientists who developed it. In the technique, light is first phase modulated to add sidebands to it. It then interacts with the cavity and is picked up by a photodetector. After that, it is demodulated at sideband frequency, and the DC part is filtered out and sent into the servo system. Each step will be explained in detail below.

First the light is sent through an Electro Optic Modulator (EOM). An EOM consists of a crystal and something that can generate an electrical field across it. Our crystal is lithium niobate ($LiNbO_3$) and due to Pockels Effect the birefringence of the crystal will depend linearly on the electrical field. If we send a sinusoidal field of frequency ω_m into the EOM, light propagating through the crystal will periodically speed-up and down. This is the same as saying that we periodically change the phase of the light, creating a phase modulation.

The effect of phase modulating light can be explained by representing the light as a phasor:

$$E = E_0 \cdot e^{i(\omega t+\theta)} \tag{14}$$

where E_0 is the amplitude, ω is the frequency, t is time and θ is the phase. Let us make a coordinate system that rotates with the same frequency as this phasor. The phasor will then just be a still-standing line in this coordinate system and its angle will given by θ. Modulating the phase is analogous to rotating the line back and forth with the modulation frequency of the EOM (see figure 3). Mathematically this means that a sine is added to the exponent in the previous equation:

$$E = E_0 \cdot e^{i(\omega t+\beta sin(\omega t))} \tag{15}$$

We have left out the original phase since we could just choose it to be zero.

The new phasor in equation 15 can now be described as a superposition of different frequencies. As it turns out we need an infinite series of frequency components to completely describe the phasor, but we only need a couple of terms to make a good description. The first term is called the carrier and the next two are called the first order sidebands. Mathematically these three terms are described as:

$$E = E_0(J_0(\beta)e^{i\omega t} + J_1(\beta)e^{i(\omega+\omega_m)t} - J_1(\beta)e^{i(\omega-\omega_m)t}) \tag{16}$$

where E_0 is the amplitude of the unmodulated light, J_0 is the zero'th order Bessel function, J_1 is the first order Bessel function, ω_m is the modulation frequency and β is the same as in the previous equation and named the modulation depth. The carrier is just the field we had before modulation. The sidebands have different frequencies from the carrier so they rotate in the coordinate system mentioned earlier. Since one of them has a slightly higher frequency than

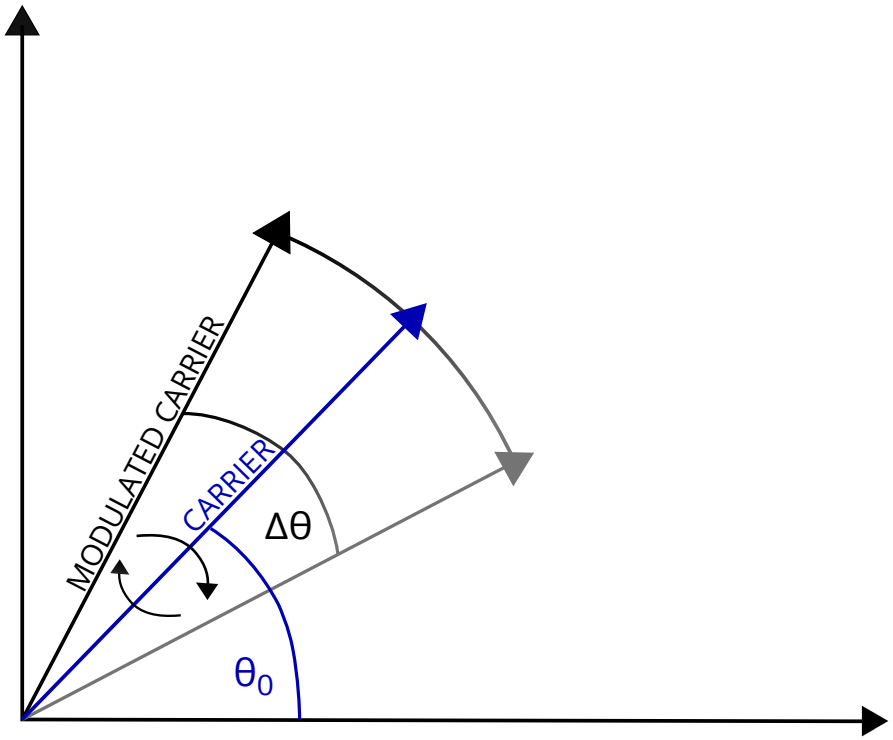

Figure 3: Before modulating the light, it is simply described as a phasor standing still as it has constant phase and the coordinate system rotates at the same frequency. We call this the carrier. When modulating, the phase changes sinusoidally between two values so the phasor rotates back and forth. In the figure θ_0 is the constant phase of the carrier and $\Delta\theta$ is the total deviation in phase in the modulated carrier.

the carrier and the other has a slightly lower, they rotate opposite ways. The situation has been drawn in figure 4. In the phasor representation it is easy to see that the sum of the carrier and the first sidebands are (to first order) equal to the modulated, rotating phasor, and so one can see visually how phase modulation is connected to frequency modulation. One can also see the error of the first order approximation: It changes amplitude.

Next the light is sent into the cavity and we look at the reflected part of it. As described in section 2.1 the reflected light's amplitude will be the original amplitude E_0 times a reflection coefficient r_C that depends on the frequency of the light (see eq. 11). This means that each part of the modulated light will be multiplied by this reflection coefficient after having been reflected off the cavity:

$$E_r = E_0(\, r_C(\omega)\, J_0(\beta)\, e^{i\omega t} + r_C(\omega+\omega_m)\, J_1(\beta)\, e^{i\omega+\omega_m t} - r_C(\omega-\omega_m)\, J_1(\beta)\, e^{i\omega-\omega_m t}) \tag{17}$$

Now, the entity measured by the photo detector is not the electrical field but

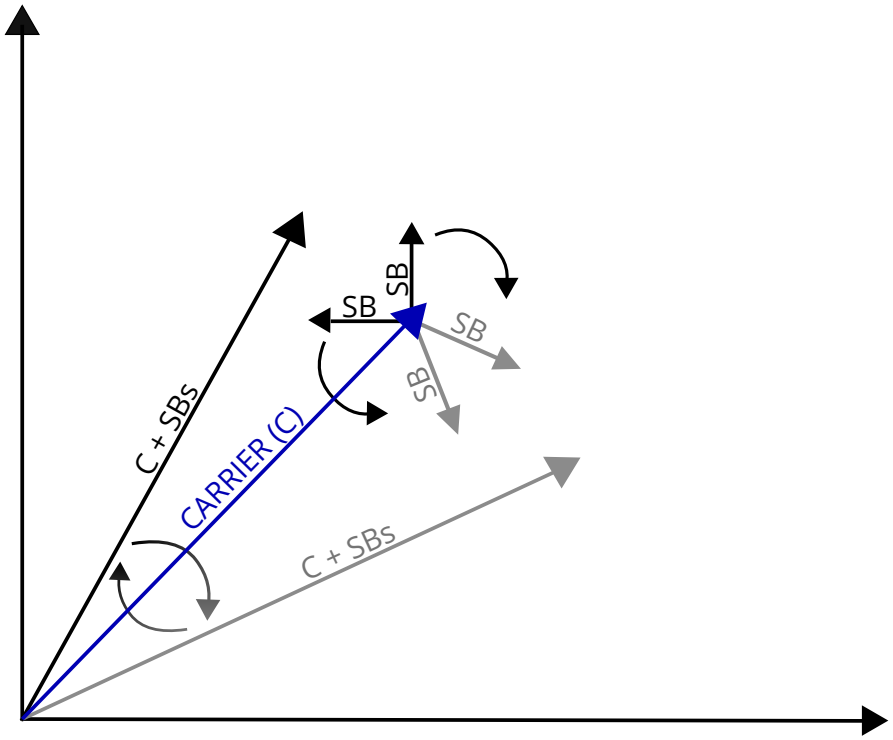

Figure 4: The phase-modulated phasor is now described by the sum of the carrier and two smaller vectors called sidebands that have slightly different frequencies than the original light. Since their frequencies do not match that of the coordinate system the sidebands rotate. The sum rotates back and forth and is almost equal to the modulated phasor of the last figure for small modulation frequencies. In the figure SB is short for sideband, and the grey sum and sidebands are at a later time than the black ones.

the power, so we use that $P \propto |E|^2$:

$$\begin{aligned} P_r = P_c|r_C(\omega)|^2 + P_s(|r_C(\omega+\omega_m)|^2 + |r_C(\omega-\omega_m)|^2) \\ +2\sqrt{P_cP_s}\{Re[r_C(\omega)r_C^*(\omega+\omega_m) - r_C^*(\omega)r_C(\omega-\omega_m)]cos(\omega_m t) \\ +Im[r_C(\omega)r_C^*(\omega+\omega_m) - r_C^*(\omega)r_C(\omega-\omega_m)]sin(\omega_m t)\} + (...) \end{aligned} \tag{18}$$

where $P_c \propto |E_0\ J_0(\beta)|^2$ and similarly $P_s \propto |E_0\ J_1(\beta)|^2$. The terms left out are the ones that depend on $2\omega_m$, and they have been left out because they do not matter for the technique. Really, the only terms that do matter are the ones multiplied by a sinusoidal with frequency ω_m, since these will allow us to make our discriminant when they are demodulated. To extract the desired terms the signal is multiplied by an oscillation of frequency ω_m to use the trigonometric identity:

$$\cos x \cos y = \frac{1}{2}(\cos(x-y) + \cos(x+y)) \tag{19}$$

which for x = y becomes

$$\cos x \cos x = \frac{1}{2}(1 + \cos(2x)) \tag{20}$$

One can show that at high modulation frequencies (compared to the cavity's optical response time) the imaginary terms will be dominant, and at low frequencies the real one will dominate. This just means that the phase of the oscillation we multiply with should be adjusted in phase. Once it has been adjusted the desired terms now exist in the signal as a DC signal: $r_C(\omega)\ r_C^*(\omega+\omega_m) - r_C^*(\omega)\ r_C(\omega-\omega_m)$, which can be filtered out by a low-pass filter. This is then sent to the servo system to be used as an error signal. A PDH signal with high modulation frequency is shown in figure 5.

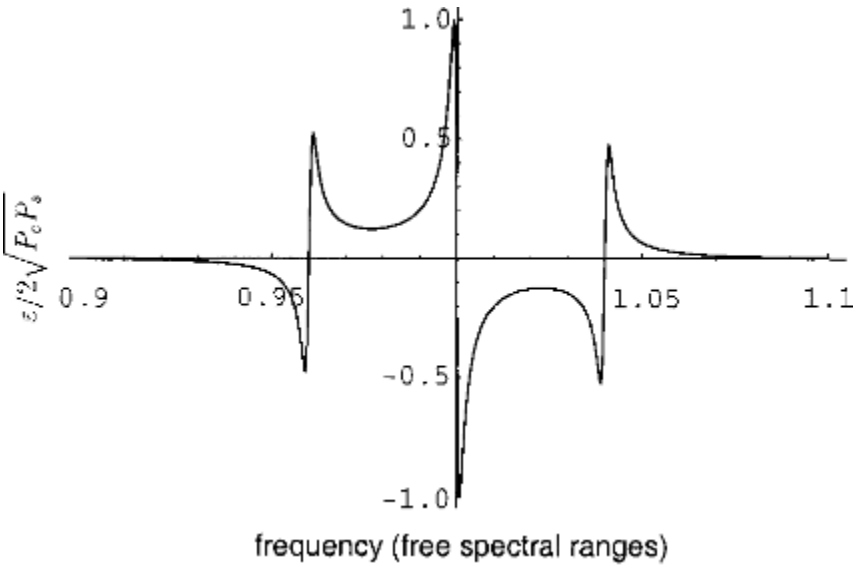

Figure 5: Pound-Drever-Hall signal with a high modulation frequency that is roughly 4 % of an FSR and a finesse of 500. ϵ is the imaginary term in equation 18. Figure is taken from [10] .

One can show that the optimal modulation depth is $\beta = 1.08$ since this makes the slope the steepest (which means that even a small deviation in frequency gives us a large error signal to act on) [10]. Furthermore higher finesse and higher intensity give a steeper slope.

2.4 Allan Deviation

When a frequency standard has been made, its stability must be tested, so that it can be compared to other frequency standards. To do so, one looks at either the phase deviation $\phi(t)$ or the frequency deviation $\Delta\nu(t)$. For comparison between frequency standards of different frequencies, these quantities are normalised by dividing them with the nominal frequency:

$$x(t) = \frac{\phi(t)}{2\pi\nu_0} \tag{21}$$

$$y(t) = \frac{\Delta\nu(t)}{\nu_0} = \frac{dx(t)}{dt} \tag{22}$$

where ν_0 is the nominal frequency. For many types of measurements the standard deviation can be used to describe how much the data points vary by looking at each data points' deviation from the mean of all the data. This is problematic to do with measurements of frequency standards, since the mean is not always constant, and so the standard deviation is not always convergent. The Allan deviation instead compares adjacent sets of points so that one can see how much the data changes over a specific amount of time.

The standard deviation of a sample can be estimated by the formula:

$$s = \sqrt{\frac{1}{K-1}\sum_{i=1}^{K}(y_i - \frac{1}{K}\sum_{u=1}^{K} y_u)^2} \tag{23}$$

where K is the amount of data points and y_i and y_u are data points. Now to get to the estimate for the Allan deviation instead, let us set K = 2, so that we get:

$$s_2 = \frac{1}{2-1}\sum_{i=1}^{2}(y_i - \frac{1}{2}\sum_{u=1}^{2} y_u)^2 = \frac{1}{2}(y_2 - y_1)^2 \tag{24}$$

This means that instead of comparing all values to the mean we simply compare each point to their neighbour. To get a good estimate we do this for all points:

$$\sigma_y^2 = \frac{1}{2(M-1)}\sum_{i=1}^{M-1}(y_{i+1} - y_i) \tag{25}$$

where M is the amount of data points, which for the Allan deviation will typically be frequency deviation measurements and σ_y^2 is the Allan deviation. Now all these data points are measured over a time interval, so each point is basically an average taken over that measurement time, which we call τ. If one changes the measurement time, one changes how large intervals of the measured entity the Allan deviation is comparing, which will naturally change the value of the Allan deviation. What is usually done, is to let the measurement time τ change from a small value to a large one to depict the Allan deviations dependency on it. Instead of actually changing the sample time one can simply take one series of measurements and then take increasingly larger averages of groups of data points, so that one effectively gets an increasing τ.

To get the most of one's data, and thus get higher statistical confidence, one can calculate the Allan deviation of averaged data with different offsets, so that you start comparing e.g. points 1, 2 and 3 to points 4, 5 and 6 and then compare points 2, 3, 4 to points 5, 6 and 7 and so on. This is called the

overlapping Allan deviation and is given by (see [6]):

$$\sigma_y^2 = \frac{1}{2m^2(M-2m+1)} \sum_{i=1}^{M-2m+1} \left(\sum_{i=j}^{j+m-1} (y_{i+m} - y_i) \right)^2 \tag{26}$$

Now if you have to do two nested sums for each τ this calculation will get quite tedious. Luckily if we use the phase deviation instead of the frequency deviation we lose a sum (see [6]):

$$\sigma_y^2 = \frac{1}{(N-2m)\tau^2} \sum_{i=1}^{N-2m} (x_{i+m} - 2x_{i+m} + x_i)^2 \tag{27}$$

where we've defined $N = M + 1$, because if we use frequency data to find the phase data we'll get an extra point, since we're integrating.

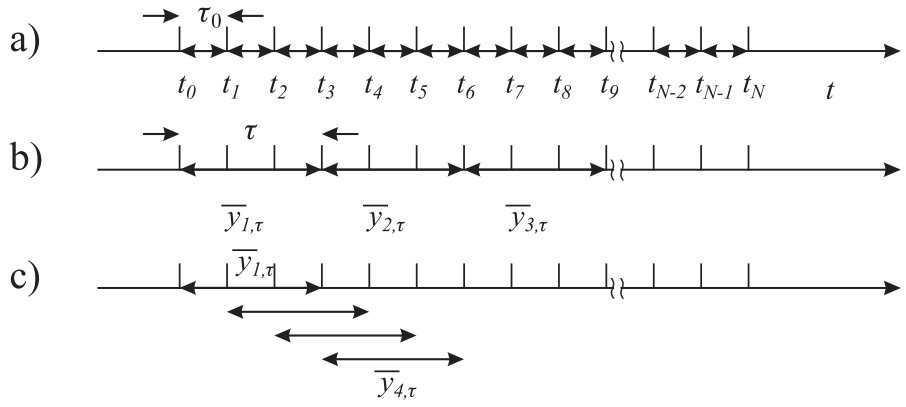

Figure 6: a) Measurements with measurement time τ_0. b) Averages over three data points which gives an effective measurement time of $\tau = 3\tau_0$. c) Overlapping averages that allows higher statistical confidence of the Allan deviation. Figure from [1] page 52.

Again it should be noted that the formulas we have written up so far concern how one would estimate the Allan deviation from a set of data. The true Allan deviation is given by the expectation value of this squared difference (see [1]):

$$\sigma_y^2 = \frac{1}{2}\langle (y_2 - y_1)^2 \rangle \tag{28}$$

The good thing about knowing this, is that now we can calculate how we would expect a specific frequency modulation to influence the Allan deviation. For example the very likely case of a harmonic modulation:

$$y(t) = \frac{\Delta\nu}{\nu_0} sin(2\pi f_m t) \tag{29}$$

Which gives the following Allan deviation (see [7]):

$$\sigma_y(t) = \frac{\Delta\nu}{\nu_0} \frac{sin^2(2\pi f_m \tau)}{\pi f_m \tau} \tag{30}$$

More specifically this can be done by following the procedure used in [7], where the modulation is described in the frequency domain. This is nice in case of more complicated modulations, because they might be simple in the frequency domain but quite hard to describe in the time domain. Now the last thing we really need to know about the Allan deviation is how to find the error of our estimate. The next section follows [11] page 13-15. The phase deviation y is in many cases distributed according to a Gaussian distribution. Since the Allan variance is the square of the difference between two y's it is distributed according to a χ^2 distribution. To be more specific it is distributed according to:

$$\chi^2 = \frac{df \cdot \sigma_{ys}^2}{\sigma_{yt}^2} \tag{31}$$

where df is degrees of freedom, σ_{ys}^2 is the Allan variance estimated from the sample and σ_{yt}^2 is the true Allan variance. Now to get a minimum and maximum value of the true Allan (at a given confidence interval) we can simply look at the distribution of the χ^2:

$$p(\chi^2) = \frac{(\chi^2)^{\left(\frac{df}{2}-1\right)}}{2\ df\ \Gamma\left(\frac{df}{2}\right)} e^{-\frac{\chi^2}{2}} \tag{32}$$

where Γ is the gamma function. The noteworthy thing about this distribution is that it only depends on degrees of freedom, so given those, we can choose an interval and find the minimal and maximal χ^2 value. When we've done that we can use formula 31 to determine a minimal and maximal Allan variance. The degrees of freedom are somewhat complicated for the case of the overlapping Allan deviation, since the overlapping estimates are naturally correlated to each other. Besides from that the different noise types correlate the data in different ways. Fortunately the degrees of freedom for overlapping Allan deviation estimates have been approximated for different noise types, see the table in [6] p. 40. This means that one has to know which type of noise is in the system to compute the error of the Allan deviation.

2.5 Saturated Absorption Spectroscopy

Eventually, when the cavity is locked to the laser, and the project can move on to locking the laser to an iodine transition, something called Saturated Absorption Spectroscopy (SAS) will be utilized. In short, the technique allows us to get around the Doppler broadening of the transitions by making a little peak at each hyperfine transition with a Doppler-free linewidth. If this peak is detected while scanning the iodine transition with the cavity lock engaged, the digital controller is good enough for the project.

Doppler broadening happens because the molecules inside the cell in the cavity have a velocity component in the direction of the light propagation. This shifts the observed frequency of the light for the molecules so it can be absorbed even when it's not at the molecules' rest resonance frequency. When the light interacts with the molecules while going one way through the cell, it will see a Doppler shift of one sign, but when it goes back, it will see the opposite. This means that it will only interact with these shifted molecules while going in one direction. The molecules that have a velocity perpendicular to the propagation of the light will, however, not be Doppler shifted. This means that the light will interact with them while going both ways. If the light is exactly on resonance it will only be able to excite the molecules with a perpendicular velocity, and since these are interacted with twice as much, the transition is saturated twice as much, which means that a transition peak with Doppler-free linewidth will appear. The peak is broadened by various effects, which is why it's not exactly the natural linewidth.

3 Experimental Set-up

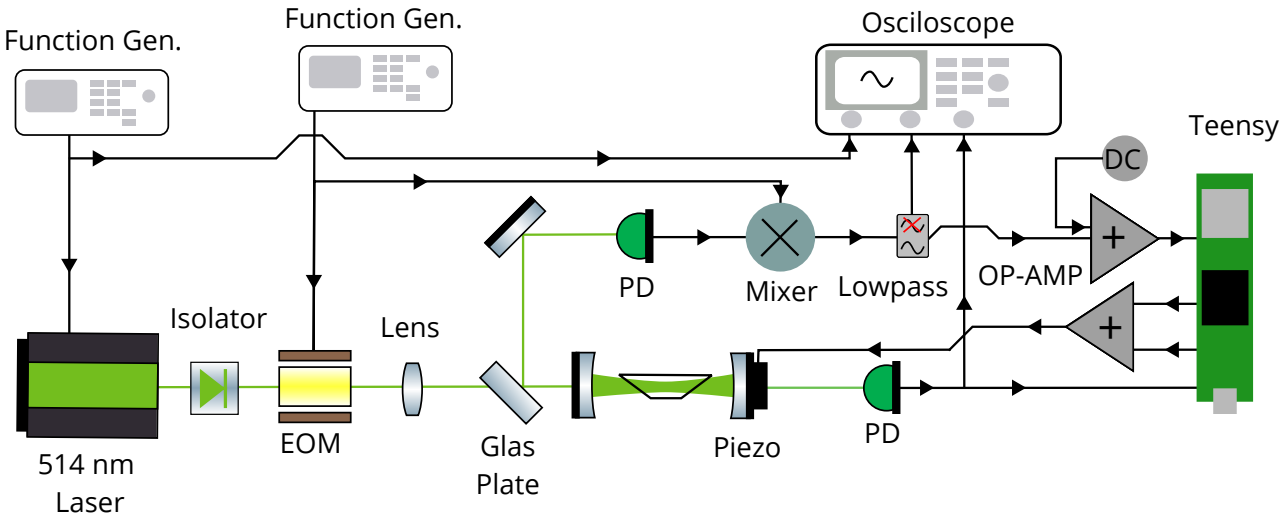

Figure 7: A simplified illustration of the experimental set-up. Phase modulated light is incident on an optical cavity. Transmission and PDH signal is fed to the Teensy, which controls the length of the cavity though a piezo.

The experimental set-up (fig 7) consists of a narrow linewidth green laser passed though an optical isolator and an EOM. It is then incident on the optical cavity with an inter-cavity iodine cell. The reflection is guided to a PhotoDetector (PD). The detector signal is then amplified, AC-coupled and fed into a mixer. The mixer is connected to the same frequency generator as the EOM, allowing us to produce/adjust the PDH signal and relative phase and amplitude

between the EOM and mixer. The PDH signal is then fed though a low-pass filter into a summing amplifier giving it a positive offset. The now always positive error signal is fed into the Teensy microcontroller which produces a control signal. The control signal is amplified and fed into the piezo mounted on a cavity mirror closing the loop. The cavity transmission is detected on a PD and sent to the Teensy directly. This signal is used to check if the cavity is really on resonance while running the PID code by checking that the transmission is above a certain threshold. This is necessary because the noise outside the PDH error signal has the same value as the error signal itself when exactly on resonance.

4 Methods

4.1 Setting up the Cavity

We first focused on mode matching the cavity. To help us we used the programs "Gaussian Beam" and "reZonator" to simulate our laser beam. ReZonator simulates the cavity to find the waist, based on the radius of curvature (ROC) of the cavity mirrors and wavelength of the laser. For a $\lambda = 514.67$ nm beam and ROC = 9 m mirrors, the waist is 434,5 μm. By requiring Gaussian Beam to produce this exact waist at the center of the cavity, it produced guidelines for a reasonable arrangements of lenses of various focal length to achieve mode match.

We also found the linewidth and finesse of the cavity before and after placing the iodine cell inside the cavity. This yielded a linewidth of $\sim$1.3 MHz and a finesse of $\sim$400 without the iodine cell, and a linewidth of $\sim$5.5 MHz and a finesse of $\sim$90 with the iodine cell.

4.2 Producing the Pound-Drever-Hall signal

When producing the PDH signal we align the cavity such that the transmission of the carrier frequency is as big as possible. Experimentally we varied the modulation frequency from 1-25 MHz. As a tendency we saw that higher modulation frequencies produced a steeper PDH slope, so we set the modulation frequency to 25 MHz, which was the largest we could produce. The iodine cell was then cooled to increase the finesse of the cavity, which made the slope of the PDH signal steeper in accordance with what we mentioned in section 2.3. We also made sure that the laser frequency was not on resonance with the iodine, which would increase loss in the cavity and reduce the finesse. Then we maximized the amplitude of the PDH signal by adjusting the phase of the signal into the mixer. This procedure ensures that we have a steep PDH signal,

which will be optimal for locking the cavity. Of course the PDH signal can not always be optimized (for example because we have to be on resonance with the iodine, when we in the future lock the laser to a hyperfine transition), but it was done at this point of the experiment to examine the PDH signal properly and optimize the servo system more easily.

4.3 Locking the cavity

We use a Teensy 3.6 microcontroller to lock the cavity to the laser. We have programmed the Teensy to work as a digital PID circuit, as described in section 2.2. A difference to an analogue circuit is that the Teensy cannot take negative voltages and it has an input and output range of 0 - 3.3 V. We therefore offset the signal by 1.65 V so it oscillates around the middle of the range. The PDH signal is also amplified, such that the amplitude of the signal matches the Teensy's operating range, which optimizes the resolution of the error signal. The cavity transmission signal - used to determine if the Teensy is on resonance - is fed directly into the Teensy. Once the Teensy is initialized, it is programmed to send a ramp signal to the piezo, linearly changing the cavity length to find the length that matches the laser wavelength. While scanning it finds the offset of the PDH signal by taking averages of the signal. The offset is then set as the setpoint for the error signal, since it should correspond to the middle of the PDH slope, i.e. resonance.

The Teensy is given a value of the expected height of the transmission peak, and when the signal from the transmission rises above a certain percentage of this, it initiates the PID lock.

Since the Teensy is running digital PID control code, we are limited by how fast and precisely the Teensy is capable of reading, processing and producing signals. We tested the Teensy by itself, by having it run a simple P-lock on a chirp signal, and found that it had a corner frequency of $\sim$35 kHz. The fastest PID-lock we were able to make work with the entire servo system, was only able to control disturbances up to 20 kHz (see section 5.2.2), which might be because of the piezo's frequency response. The PD we use has a bandwidth of 50 kHz, and is therefore the maximum speed we would be able to regulate. The low-pass filter right after the mixer has been chosen to be 50 kHz for this reason. The low-pass filter cuts out some of the noise from the mixer.

4.4 Searching for Lamb Dips

Once the cavity has been locked, and the lock has been optimized, we can begin to look for Lamb dips. We do this by linearly scanning the laser frequency ($\Delta\nu \approx 50$ MHz), while the servo system makes the cavity stay on resonance with

the laser. We set the laser's center frequency on resonance with the iodine, but this reduces the finesse, which changes the PDH signal, which changes what the optimal PID constants are, and therefore makes the servo system work poorly or disengage entirely. So every time one changes the laser's center frequency while on resonance with the iodine one has to readjust the PID constants. The lock can also disengage, if the frequency modulation becomes too large, which is why we only modulate the frequency 50 MHz. Through trial and error we found that the optimal constants approximately scaled with a common factor, so we implemented a scaling factor to the PID constant to make the work easier. There is a way to get around manually changing the scaling factor each time the center frequency is changed, and it is described in section 6. Once the lock has been re-engaged we are effectively scanning out a small section of the transition. We also change the laser intensity to different values because high intensity broadens the Lamb dips, so they may not be visible at all intensities. This also changes the scale of the PID constants, by increasing the amplitude of the PDH signal.

If a Lamb dip is found, it would be useful to estimate its optical linewidth. To do so the laser frequency is modulated in the same way as when the Lamp dip is measured, but with the servo system deactivated so the PDH signal will be measured instead of the error signal. The laser modulation frequency has to be turned up ten times, since the cavity is no longer locked and is thus subject to a lot of noise. The modulation frequency (not amplitude!) used to create sidebands for the PDH signal is turned down until the sidebands are both within scanning range. Then the known distance in frequency between the sidebands can be used to convert the Lamb dip's width in time to a width in frequency.

5 Results & Discussion

5.1 The Pound-Drever-Hall signal

Looking at the PDH signal and transmission in fig. 8, we see that the PDH signal looks much like predicted in section 2.3. The dashed lines in the figure indicates points of interest in the plot. The green line at -25 MHz indicates where one of our side bands is placed. It is easy to see on the PDH plot but completely invisible on the transmission. This is simply because the modulation depth β is small. The x-axis of the plot was originally in seconds, but using that the frequency difference between the two side bands is 50 MHz we converted the x-axis into frequency, making the side bands lie exactly at ± 25 MHz.

The cyan line is placed at the center of the transmission peak. Notice that this is not the center of the PDH signal. We do not currently have a good explanation for why this is. It is not a big problem since, we are simply able to

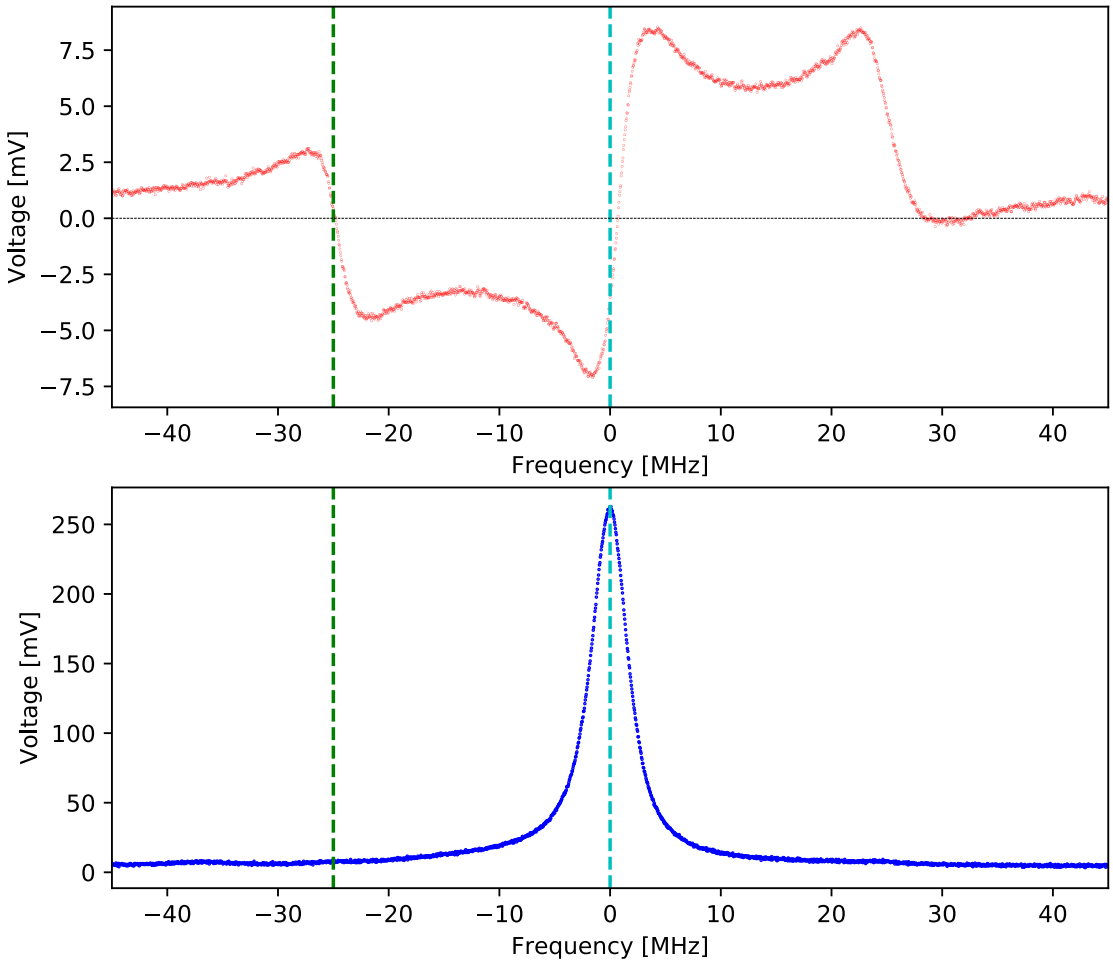

Figure 8: The Pound-Drever-Hall signal(red) and transmission(blue).The green dashed line marks one of the two sidebands, the x-axis has been converted from time to frequency using the difference between the sidebands as reference. The cyan line is the middle of the transmission peak. It is not at the center of the PDH slope. This is not of no concern since we can adjust our set-point in the Teensy. The offset from zero comes after the mixer and is dependent of the phase and amplitude relation between the EOM and mixer.

digitally adjust the PID set-point, until we maximise the cavity transmission. However, if the maximal transmission and middle of the PDH slope are too far apart, e.g. such that the resonance is a the bottom of the PDH slope, the lock would be more unstable since the effective locking range is reduced, or even worse the lock might disengage entirely .

The PDH signal has an offset of $\sim$1.25 mV, which would move the center of the slope a bit to the right. This offset appears right after the mixer. Before reaching the Teensy the offset is added to the offset we ourselves apply to the PDH signal. So it is of little concern, since the Teensy already compensates for an offset. But, we have seen the offset scale with the phase and amplitude of the demodulation signal, and we have also seen it drift significantly within minutes. This is a problem that can be solved by programming the Teensy to vary the set point to optimize the transmission signal.

5.2 Cavity & lock stability

5.2.1 Allan deviation

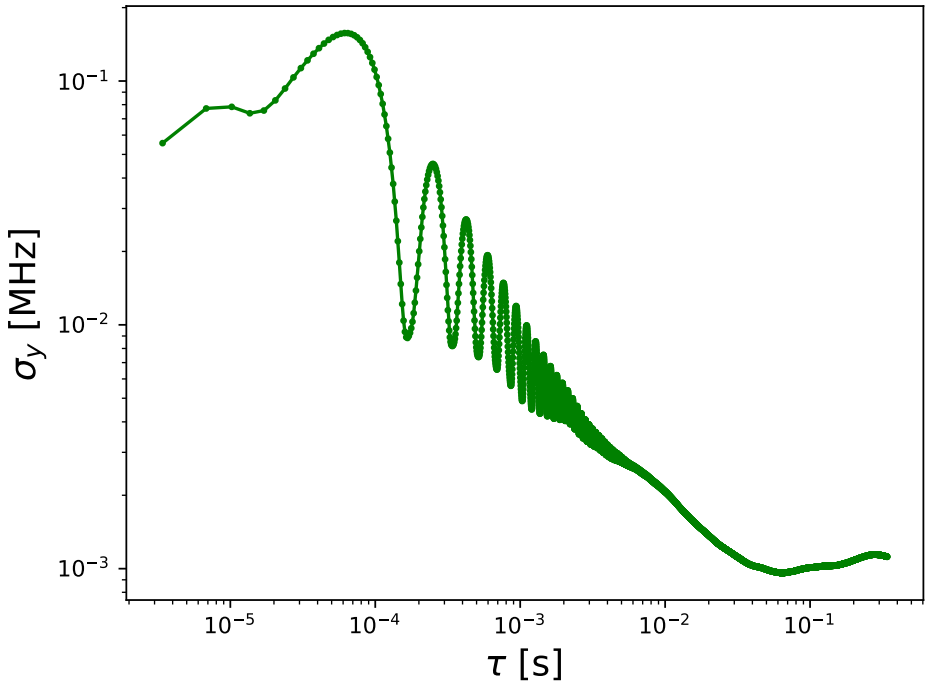

Figure 9: The Allan deviation of our system in log-log scale, showing a massive oscillation around $\tau = 10^{-4}$ s. We can also see something that might be the servo bump around $\tau = 6 \cdot 10^{-5}$ s. At high τ the Allan deviation rises because of long-term effects like rise in temperature. The Allan deviation is in MHz, because a cavity length deviation is analog to a frequency deviation. The errors are plotted but are so small that they cannot be seen.

The Allan deviation calculated from the in-loop measurement of the cavity length deviation has been plotted in figure 9. In the figure there is a bump around $\tau = 6 \cdot 10^{-5}$ s, which could very well be the servo bump. There is a large dampened oscillation appearing around $\tau = 10^{-4}$ s. According to our theory about Allan deviation (see section 2.4) this could be due to the length of the cavity being harmonically modulated (because this is completely analog to the frequency modulations mentioned in that section). To find the frequency of the modulation and see if it follows equation 30 we looked closer at the oscillation and attempted to compare it with the model (see fig 10).

The data was fitted with the theoretical model in equation 30 and afterwards a few insignificant coefficients were changed to make it fit the first oscillation, since it makes it easier to see the problem with the model: That it dampens the oscillation too slowly. That the model doesn't quite fit is actually not so surprising, since it holds true only for a modulation with a single frequency, that is, a modulation described as a delta function in the frequency domain. Practically this means that the model will only hold as long as the width of

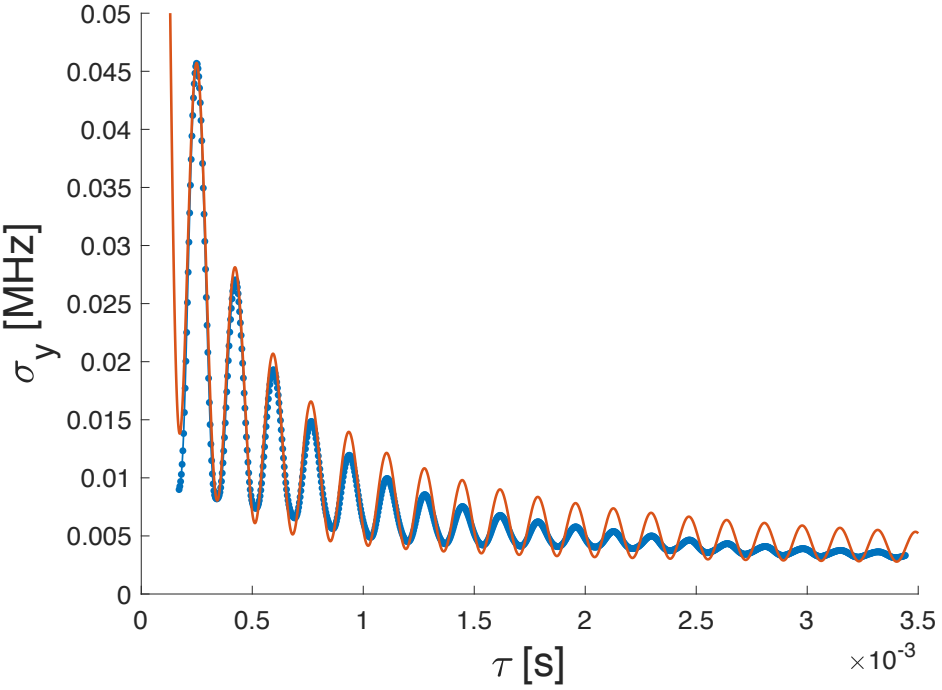

Figure 10: The theoretical model in equation 30 (orange) compared to the data (blue). The data clearly dampens much faster, so the model is not completely accurate. $\frac{a}{\tau}$ (where a is some constant) has been added to the model because there are many other frequencies present that all follow the model, that is, they all lose influence as $\frac{a}{\tau}$. Their oscillations are just too small to be seen.

the modulation in the frequency domain is very narrow compared to the range of frequencies present in the error signal. It could also be that our time axis is wrong. In our calculation of the Allan deviation we assume that the sample time is equal to the measurement time, that is, we assume that there is no dead time between each measurement point. If there should be a dead time, the measurement time would be lower than our estimate (which is simply the sample time), and the time axis would "shrink". This would effectively squeeze the data further together, but as we can see on the figure this would make the model even worse. So unless our measurement time has somehow been estimated too low (rather than too high) this is not an explanation. One could find the correct model with the procedure mentioned in section 2.4, but we didn't have the time. The frequency of the harmonic modulation was estimated by Matlab's fitting algorithm to be 5860 ± 3 Hz.

5.2.2 Fourier transform and piezo resonances

The error signal was Fourier transformed (see figure 11) to examine the modulation mentioned in the previous subsection and see if the system generally behaves as expected. From around 20 kHz, the amplitude of the modulations starts rising at roughly 27 dB/decade. This is probably just the lock getting

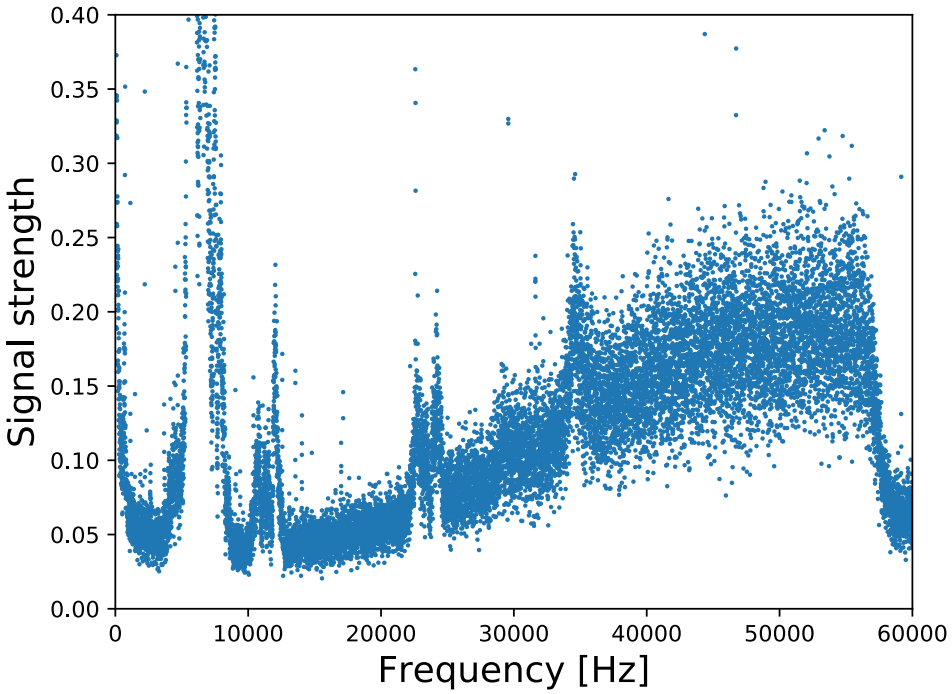

Figure 11: A Fourier transform of the error signal. Around 20 kHz the roll-off of the servo system starts and it is of roughly 27 dB/decade.

weaker and the disturbances getting stronger as a result, which means the corner frequency of the servo system is around 20 kHz. At around 57 kHz the signal suddenly disappears. This is due to our photo detector not being able to measure higher frequencies than 50 kHz and the 50 kHz low-pass filter, and this means that the range that contributes to the measured Allan deviation is approximately from 0 to 57 kHz.

A rough estimate puts the FWHM of the modulation at ~ 180 Hz and the FWTM (full-width-at-tenth-maximum) at ~ 570 Hz. However, as can be seen on figure 12 there are other peaks next to the big one and together they span several kHz (from around 5 kHz to around 8 kHz). So overall it makes sense that the Allan deviation did not completely follow the expected model, since the big peak spans $\sim 1\%$ of the whole range and the smaller peaks close to it span $\sim 5\%$ of the whole range. The middle of the top of the peak is located at ~ 5860 Hz, which matches the frequency found in the last subsection.
We believe that the peak is due to a mechanical resonance of either the piezo, or the system of mirror and post it is attached too, or most likely a combination of the two. The influence of the post could be tested by removing it. If the piezo was dismantled from the mirror and post one could furthermore measure its intrinsic frequency, and it would possibly be a simpler resonance peak, since the system would be much less complicated. One would also know for certain

whether the piezo was producing the peak or not, and the knowledge of the piezo's intrinsic resonance could allow one to design a future set-up in a better way. It would have been optimal to design the system so that the resonance had been well over the range that the lock was assumed to work in, and in future systems this should be implemented to improve system stability. The problem with the peak is, that the gain in the region around 6 kHz is much higher than we would like it to be, and that the phase just around the peak of the resonance might be the opposite of what we want, so that the lock will amplify the disturbances instead of attenuating them. Piezos make a 180 degree phase shift around resonance, so if we assume that our lock has perfect phase before it we would in fact start amplifying all frequencies after the resonance. This is not visible in our Fourier transform however, and it is probably because the phase is not perfect before. If the phase was 90^{o} before resonance it would be 180^{o} at the very top of the resonance and then at -90^{o} afterwards, which is effectively the same as 90^{o}.

If one imagines the system without the resonance, one could imagine building an optimal lock that would attenuate the noise as much as possible in the largest range of frequencies possible. If we applied this lock to our system, the gain would be unreasonably high at resonance, so our lock would oscillate the system violently and thus bring the cavity length outside of the scope of the PDH slope. In other words, the lock would be incredibly unstable and probably not function at all. The reason that our lock works is that there are two ways to modify the "perfect lock" so that it functions despite of the resonance. One can either turn down the overall gain so that even at resonance it doesn't become so high that the system starts oscillating and becomes unstable. Or one can set the corner frequency so low that the gain is small enough at resonance. It is clear from our Fourier transform that the lock is active until around 20 kHz, so the solution that we have found (randomly through trial and error while optimizing the PID constants) must be the first one. This makes sense as it's probably not possible to make a functioning lock that has a corner frequency at a couple of kHz. This means that if the effect of the resonance was cancelled it would be possible to turn the gain up, and thus attenuate the noise in the system much more.

5.3 Lamb Dip

On figure 13 a presumed Lamp dip is plotted. The signal to noise ratio is very low, so just from looking at the figure it can not be concluded, confidently, that this is indeed a Lamp dip. However, when taking this data it seemed clear to us that it indeed is a Lamb dip. This is partly because we could change the laser frequency, and see that the peak moved with it. This would suggest that the

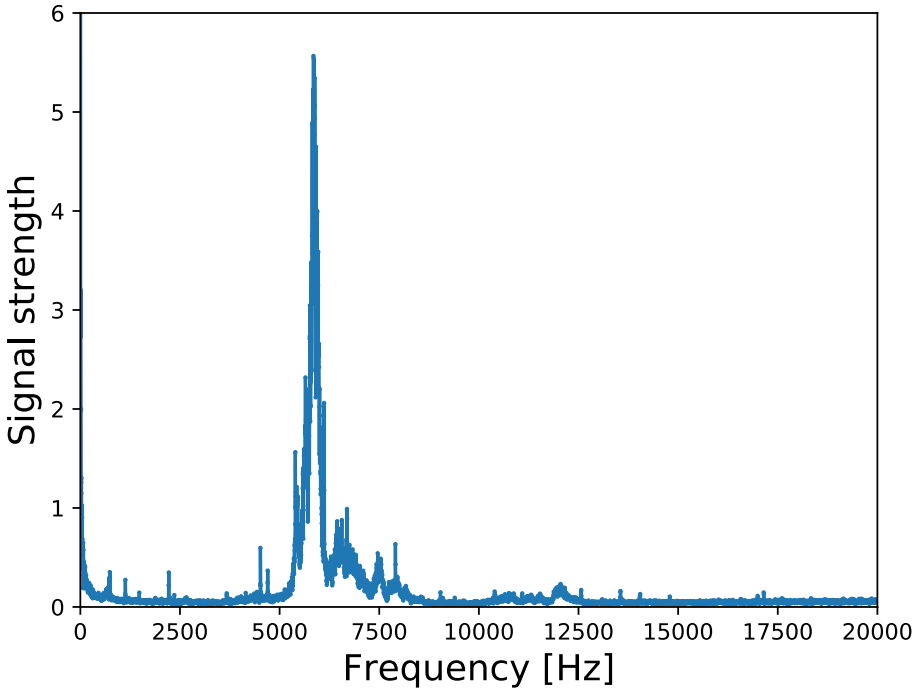

Figure 12: The Fourier transform of the error signal now with the x-axis going from 0 to 20 kHz. Here the width and height of the peak around 6 kHz can be seen clearly, and it is obvious that it would stand out in the Allan deviation, since it is much higher in amplitude than anything else in the servo system's range.

peak is positioned at a specific frequency. Furthermore we saw several other peaks like this one with similar behaviour. Unfortunately, it was not possible to scan wide enough to see two Lamb dips in the same data set, which would have enabled us to see if the frequency difference was as expected. This is, as mentioned earlier, due to the absorption changing our error signal and thus changing the optimal PID constants. A solution to this will be discussed in section 6. If there had been more time while performing this measurement, the intensity and temperature could have been optimized more, which may have made the Lamb dips more clear. System specification from this measurement can be found in the appendix, section 9.3.

With the set-up at the measurement it was not possible to determine the frequency at which the Lamp dip was located. It was however possible to estimate the width of the Lamp dip to 1.46 ± 0.02 MHz by following the procedure mentioned in section 4.4. The linewidth is consistent with what the Lamb dips in the transition should be and what has been found before [5].

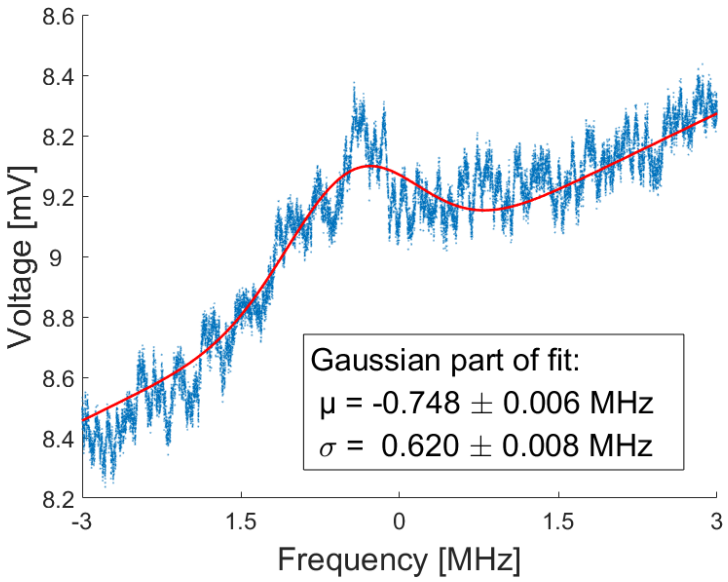

Figure 13: Plot of the found Lamb dip. The fit is a combination of a Gaussian and a linear fit. By making a similar scan of the PDH signal, we were able to calculate the FWHM of the Lamb dip as 1.46 ± 0.02 MHz. Unfortunately we do not know its position in the transition well, so we can not compare to one of the 21 transitions. [5]

6 Future Prospects

This experiment is due to see a lot of work in the future for it to become a frequency reference. This thesis has focused on locking the cavity itself and only paid little attention to the Lamb dip, however at the end of the day the Lamb dips are what we want to optimize and lock our laser to. To improve the signal to noise ratio of the Lamb dips, the cavity should be rebuild into a large waist cavity, which would increase the amount of iodine molecules we interact with, making the lamb dips taller and narrower. To lock the laser to a Lamb dip, a method called NICE-OHMS can be used to give the Lamb a PDH like profile, allowing one lock the laser to the top of the Lamb dip using control theory. Once the cavity is locked to the laser, and the laser is locked to the hyper-fine transition, the laser frequency can be beat against another (more precise) frequency reference (like strontium) to determine how good an atomic reference it has become.

What can be done right now to improve this set-up is listed below.

6.1 Attenuating the Resonance

Now that we know the piezo-mirror has a resonance spike in our effective range, we need a way to deal with it. As stated in sec. 5.2.2 one way to do it might be to simply mount the piezo on a shorter post. Otherwise a notch/band-stop filter could be implemented. Notch filters reduce the gain of a specific bandwidth, and at the center frequency of the band it flips the phase. This is the exact

opposite of what a resonance does to the system gain and phase. There are two ways to implement such a filter; digitally or analogously. There are advantages and disadvantages for each choice:

An analog notch filter is somewhat easier to build and configure (see more here: [8]). Making a single analog notch filter would attenuate the first (and most significant) resonance. If there is a need for improvement beyond this (i.e. attenuating more resonances), the analog method becomes undesirable since it would be bothersome to make a filter for every resonance.

A digital notch filter would be easily be able to attenuate many resonances, since one could simply repeat the code. The problem with this solution is that it may slow down the microcontroller excessively. One could simply implement it in the microcontroller code and see if the corner frequency changed. If the Teensy was upgraded (see next section) this would be a smaller issue.

6.2 Upgrading the Teensy

In section 4.3 we mentioned that the current resolution of the Teensy's output is around 13-bit. If it turns out that a higher resolution is needed, a new DAC could be bought and connected to the Teensy. It would be important to find one that doesn't operate so slowly that it brings down the corner frequency of the servo system. The current resolution of the ADC's is 16-bit, so if the resolution needs to be higher than that, they would need to be exchanged too. If a larger frequency range should be required for the lock, the ADC's of the teensy could be exchanged for faster ones, since they seem to be a bottle neck. We have tried changing the resolution of them, and found that it greatly influences the corner frequency.

6.3 Automatic PID adjustment

As mentioned in section 4.4 it is bothersome to scan the resonance with the laser frequency, since the optimal PID constants change when doing so. Therefore it would be helpful to implement a method that would automatically change the PID constants. A possibility is to add a sine to the Teensy's output which is far away from any piezo resonance in frequency and small in amplitude. One could then measure how much error signal this sine generates by looking at the Teensy's input. This gives us a relative measure of how much frequency deviation a certain amount of error signal signifies. Thus it tells us how and how much we should change the PID constants. This feature would allow wider scans of the iodine resonances and it would save time. It can be done very simply and probably without making the system significantly slower.

7 Conclusion

The cavity was successfully locked to the laser by setting up and programming a Teensy 3.6 microcontroller. We found that in the current set-up of the servo system there is a huge resonance around 6 kHz. This limits the quality of the lock, and it is highly recommendable to either implement a notch filter or to alter the physical set-up of the piezo in future work on the project. Despite of this resonance it was possible to find a Lamb dip, which proves that the digital controller we have made is usable in the project.

8 References

[1] Fritz Riehle - "Frequency Standards, Basics and Applications"

[2] `https://en.wikipedia.org/wiki/Clock#History`

[3] Milonni & Eberly - "Laser Physics"

[4] R.J. Barlow - Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences

[5] A. Knorr & I. Stoustrup - Laser Spectroscopy of the Hyperfine Splitting of Iodine

[6] W.J. Riley - "Handbook of Frequency Stability Analysis"

[7] Jacques Rutman - "Characterization of Frequency Stability In Precision Frequency Sources"

[8] `http://www.electronics-tutorials.ws/filter/band-stop-filter.html`

[9] `http://openaudio.blogspot.dk/2016/10/benchmarking-teensy-36-is-fast.html`

[10] Eric D. Black - "An introduction to Pound–Drever–Hall laser frequency stabilization"

[11] D. A. Howe, D. W. Allan and J. A. Barnes - Properties of Signal Sources and Measurement Methods

9 Appendix

9.1 Transmission and finesse

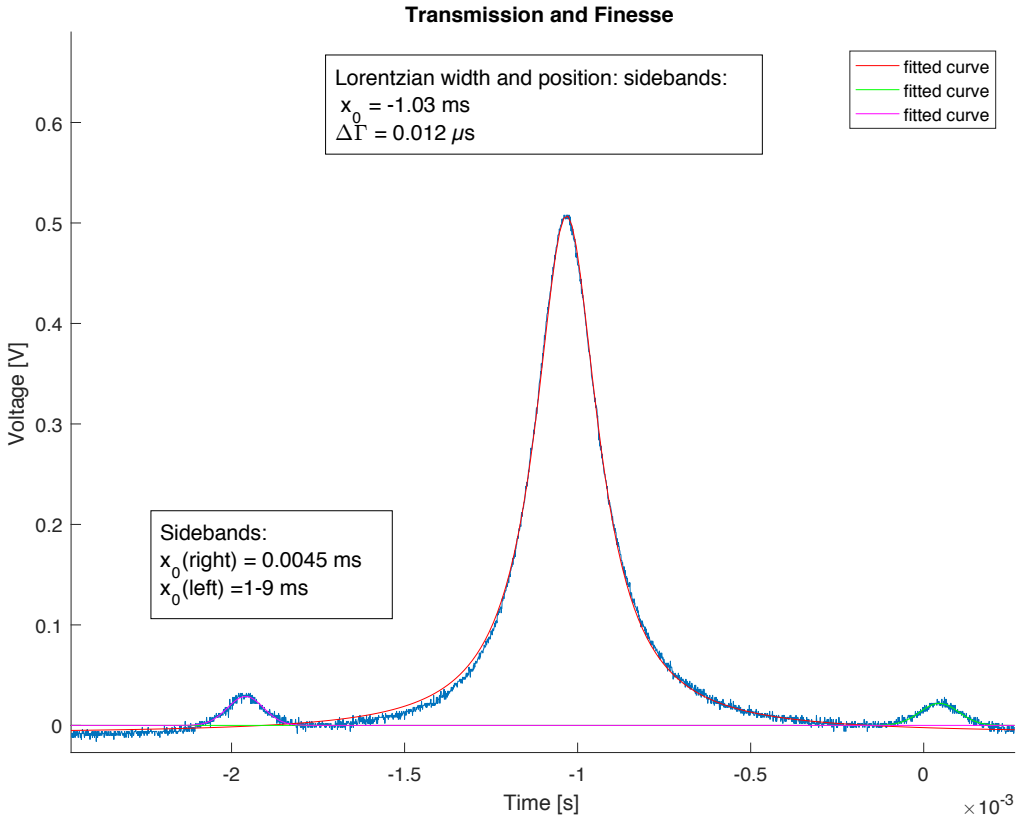

Figure 14: Transmission data used to approximate the Finesse and linewidth of the cavity. Data is taken when off resonance with the iodine, to have a high finesse.

9.2 Teensy Resolution

The Teensy's output is produced by a 12-bit Digital to Analogue Converter (DAC). 12-bit resolution means that it can produce $2^{12} = 4096$ different discrete voltages between 0 and 3.3 V. This is enough to produce a lock, but it's not one good enough to find Lamb dip. So to effectively increase the resolution we used both of the Teensy's DACs.

We attenuated the output one of the DACs by 1/43.5. This reduces its effective range but increases its resolution. By summing the output of two DACs we are able to coarsely set the correct voltage with one DAC and then use the other one to finely adjust the last bit. The resolution of this double DAC is theoretically $\log_2(43.5) \simeq 5$ bit $\rightarrow 5 + 12 = 17$ bit. But in reality the digital to analogue converters of the Teensy are not perfect. The Teensy does not always output exactly the same voltage for a given bit-value, there is some uncertainty, meaning that in reality the lock is not quite 17-bit quality. How

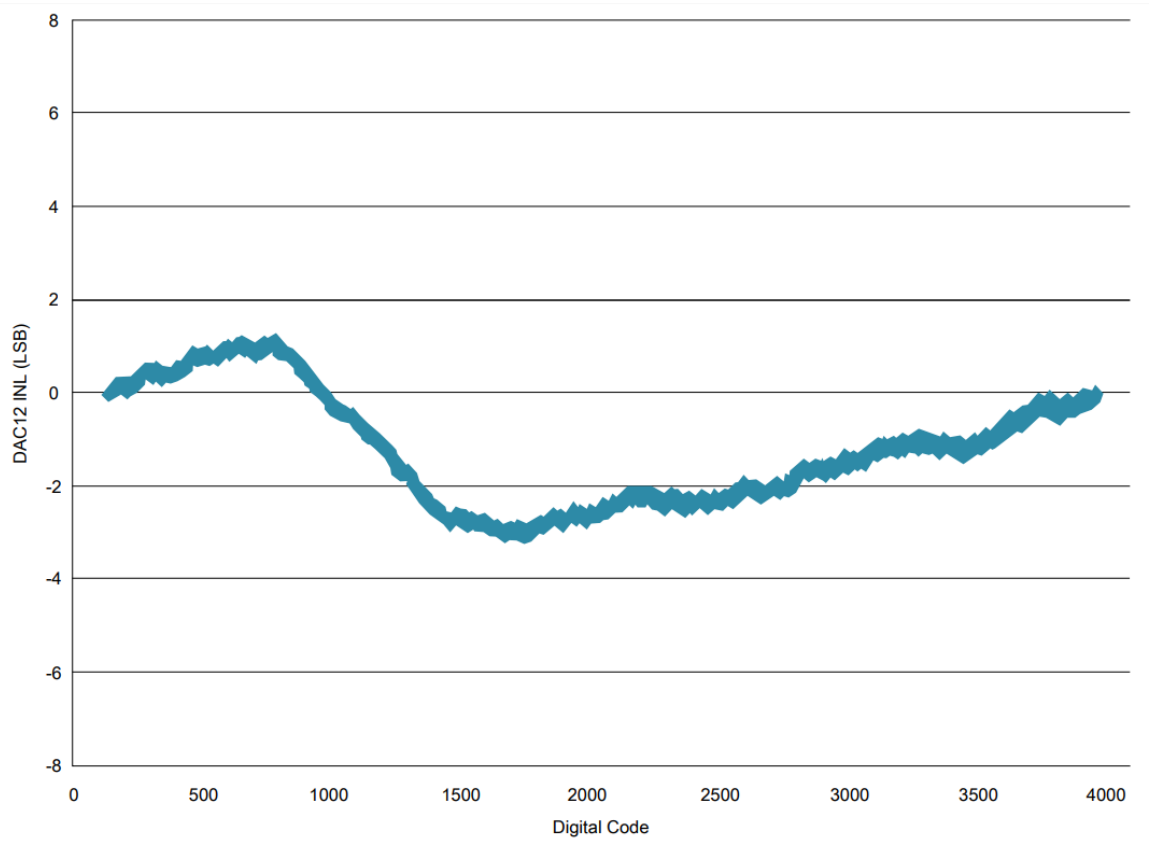

Figure 15: Bit Noise and error graph of the Digital-Analog Converter(DAC) taken from: https://www.pjrc.com/teensy/K66P144M180SF5V2.pdf On the x-axis are the 4095 different bit values one can have the Teensy output since it has a 12-bit DAC. On the y-axis is the error on that bit value. So at e.g. digital code = 2000 the Teensy will actually set the bit value to what equates to 1998 instead. The "noise" on this is by eye about 1/4 the an entire bit value. If one were to improve the DAC by using two DACs (as we do) to an effective 13-bit DAC, the noise would go to half a bit value and further upgrading this DAC to 14-bit would mean the noise is an entire bit value. Once the noise extends an entire bit value the corrections the second DAC makes are drowned in noise. Therefore this 12-bit DAC can not be improved beyond 13-bit.

far off the Teensy is from the real value is plotted in fig. 15 in the appendix. From the figure we can by eye read off the "uncertainty/noise" of the offset which is approximately 1/4 of a bit-value with 12-bit resolution. If one were to improve the DAC to 13-bit, there would be double as many bit-values as in the 12-bit DAC. This would mean that equivalently the uncertainty is doubled in size. On the 13-bit DAC the uncertainty would then be 1/2 a bit-value. Further improving the DAC to 14-bit would mean that the uncertainty is 1 entire bit-value. Now the uncertainty is as large as the precision of the correction. This is simply not noticeable, it is drowned in noise. Therefore the actual improvement of the DAC is about 1 bit.

To clarify: By theoretically improving the Teensy to 17-bit, it actually improves to 13-bit, but it does not degrade the lock that we are trying to improve it beyond 13-bit, some of the improvement is just drowned in noise. Once the

lock is engaged it can be optimized by digitally changing set-point and PID constants to reduce the amplitude of the error signal. Once the lock is optimized, we take data of the error signal to be used to calculate the Allan deviation. It should be noted that this of course is a closed loop measurement, meaning that if e.g. the piezo drifted but our electronics drifted at approximately the same speed, we would not be able to see the length drift in our closed loop system. The transmission is an out of loop measurement because it is directly affected by the length of the cavity and not by internal drift in the locking system, so we could have used it to calculate the Allan deviation. However we did not do this.

9.3 Lamb Dip, system specifications

λ_{laser}	1029.344 nm
Output Power	25.2 mW (24.5 mW actual)
Power pre cavity	0.66 μW
T_{iodine}	-10.2 ℃
Piezo Amp	45.8 V
P	30,000
I	4,000
D	10,000
Scale	100%
Laser Step freq	13 Hz
Laser Step Amp	50 mV_{pp}
Laser Step DC	2.120 V
Osc. Scale V	1 mV
Osc. Scale s	5 ms