Automatic Bias Controller for an Electro-Optic Modulator

A Thesis Presented

by

Ian Schwartz

to

The Graduate School

in Partial Fulfillment of the Requirements

for the Degree of

Master of Science

in

Physics

Stony Brook University

May 2021

Stony Brook University

The Graduate School

Ian Schwartz

We, the thesis committee for the above candidate for the Master of Science degree, hereby recommend acceptance of this thesis.

Harold Metcalf – Thesis Advisor
Distinguished Teaching Professor, Department of Physics and Astronomy

Thomas C. Weinact
Professor, Department of Physics and Astronomy

John D. Hobbs
Professor, Department of Physics and Astronomy

Laszlo Mihaly
Professor, Department of Physics and Astronomy

This thesis is accepted by the Graduate School.

Eric Wertheimer
Dean of the Graduate School

Abstract of the Thesis

Automatic Bias Controller for an Electro-Optic Modulator

by

Ian Schwartz

Master of Science

in

Physics

Stony Brook University

2021

Adiabatic rapid passage is an efficient method of exerting a large force on two-level atoms via chirped pulses of light. To create these pulses from CW light, one can utilize an amplitude modulating electro-optic modulator (EOM), based on a lithium niobate ($LiNbO_3$) Mach-Zender interferometer. Unfortunately, the crystals in EOMs are susceptible to thermal drifts that can negatively affect the output pulse quality. This thesis describes a method to measure these drifts by computing the ratio of the harmonics of an added dither signal. This device will control the bias voltage across the EOM and correct for the drifts using the measured ratio

as a feedback parameter.

Contents

List of Figures

Acknowledgements

My first time in the physics building, my first class at Stony Brook, my first optics project, and the MSI program itself all have one thing in common: Hal Metcalf. I suppose it is no surprise that I was so excited to join the Metcalf research group. Hal, I cannot thank you enough for so many fantastic opportunities and for making me feel so welcome. I am grateful for all of the lessons, advice, and things I have learned because of you. It has been an absolute pleasure working in your lab.

I cannot overstate my thanks to everyone I have had the privilege of working with. Marty Cohen, thank you for all of your guidance on my thesis as well as every project and presentation I have worked on. Eric Jones, thank you for always being willing to lend a helping hand, for pointing me in the right direction, and for all of the friendly conversations. Yifan Fang, Brian Arnold, Eli Rafkin, and Xiaoyang Liu, thank you for your warm welcome and patience with me when I joined the lab. Thank you for everything you taught me and the help along the way. Chiu Yin Liu and Edoardo Buonocore, thank you so much for all of the help on my project, either in getting it off of the ground or in keeping it running. I wish you all the best.

Of course, I also owe a huge thanks to everyone outside of the group that

made my thesis project possible. Eugene Shafto, James Eksi, and Anthony Olivo, without you none of this would have happened. Thank you for answering my endless electronics questions and for all of the assistance in making everything work the way I'd hoped.

I imagine my time at Stony Brook would have been very different if it weren't for my time in the LTC during my first summer. Michael Dapolito, Samet Damerican, Chiu Yin Liu, and Richard Joyce, thank you for the great discussions and even greater card games over lunch.

Despite being in a windowless basement most of my time at Stony Brook, I will always look back at this time fondly because of all of the friends I made throughout department. Michael Stewart, Alfonso Lanuza, Joonhyuk Kwon, Youngshin Kim, Brian Kaufman, Yusong Liu, Chuan Cheng, Samuel McClung, Myles Silfies, Jay Rutledge, Anthony Catanese, Guodong Cui, Fernando Araiza Gonzalez, and so many more than I cannot possibly name everyone. Thank you all for all of the great times. I've had a lot of fun with you all and I will miss grabbing meals with you.

Finally, I want to thank my family and friends for all of the love and support along the way. I could not have done any of this without you, and I cannot express how grateful I am to have you all in my life.

Chapter 1

Introduction

Adiabatic rapid passage (ARP) was first discovered within the context of nuclear magnetic resonance but it was soon shown to be applicable in two-level optical systems as well [1, 2]. By applying oppositely directed, alternating laser pulses to move an atom between its ground and excited state, it can be used to apply a significant momentum change in a short amount of time. Under proper conditions, the subsequent force is much larger than the radiative force and is the natural evolution from using just optical absorption to apply an optical force. In this chapter, I will discuss optical forces on an atom and provide an introduction to the ARP force.

1.1 The Radiative Force

The process behind the radiative force is relatively straightforward and is displayed in Figure 1.1. A stationary atom in a resonant light field will absorb light and move from the ground state to the excited state. Since the light

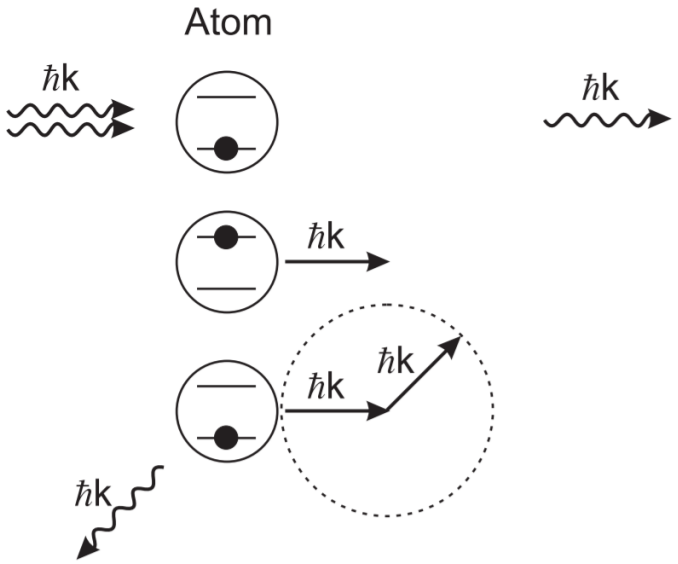

Figure 1.1: Diagram of the radiative force sequence. The atom absorbs resonant light and experiences a momentum shift of $\hbar k$ in the direction that the light was travelling. After some time, the atom spontaneously emits light in a random direction, experiencing a momentum kick in the opposite random direction. By repeating this an arbitrary number of times, the momentum resulting from the spontaneous emission will average to 0.

is carrying energy $\hbar\omega$ it must also have momentum $\hbar\omega/c = \hbar k$, which the atom also absorbs. Because of this momentum change, the atom experiences a velocity change of $\hbar k/m$ in the direction of the light, where m is the mass of the atom. In most cases, this velocity change is on the scale of a few cm/s [3].

After some time (typically $\sim \tau$, the excited state lifetime) the atom will spontaneously emit light in a random direction and consequently experience another $\hbar k$ momentum shift in the opposite direction of this emitted light. By repeating this process an arbitrary number of times the momentum shifts due to spontaneous emission will average to zero but those due to absorption will

add in the direction of the light field. If the repetition rate is fast enough, the resulting force can reach a maximum of

$$\overrightarrow{F}_{Rad} = \frac{\hbar \overrightarrow{k} \gamma}{2}$$

where $\gamma \equiv 1/\tau$. Of course, this force is limited by the lifetime of the atom. This limit can be avoided, though, by stimulating emission.

1.2 The ARP Force

As mentioned above, a resonant light pulse can excite an atom and generate a momentum change of $\hbar \overrightarrow{k}$. A second light pulse in the same direction can stimulate emission, returning the atom to the ground state. In this case, the atom will emit light in the direction of the pulse and experience a momentum shift of $-\hbar \overrightarrow{k}$. The result is a net momentum change of zero over the full cycle. However, by sending alternating light pulses in opposite directions (Fig. 1.2), these momentum shifts would be in the same direction, creating a net momentum shift of $2\hbar \overrightarrow{k}$ per cycle.

In our experiment, we focus on the $2^3S_1 \longleftrightarrow 2^3P_2$ transition of the helium, which has a resonant wavelength of 1083.33nm. For each set of pulse pairs that excite and stimulate emission, the atom experiences a net velocity change of ~ 18 cm/s.

There is no guarantee that single-frequency pulses excite or stimulate emission in an atom, even if they are on resonance. If the frequency of the applied field is swept across resonance, though, there is a very high probability that

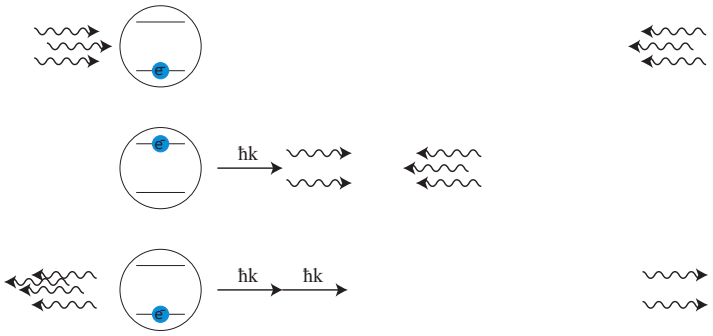

Figure 1.2: Diagram of the ARP force sequence. Light travelling in one direction excites the atom, creating a momentum change. Immediately following this, light travelling in the opposite direction stimulates emission in the atom, creating a momentum change in the opposite direction of this light. In a single cycle, the atom experiences a momentum change of $2\hbar\overrightarrow{k}$.

the light will move the atom from the ground to the excited state (and vice versa). If the frequency is swept slowly enough, one can actually expect to excite (or stimulate emission in) the atom with very high efficiency. Hence the *adiabatic* in ARP [4]. In our case we use a sweep rate of $\omega_m = 2\pi \times 160$ MHz $\approx 100\gamma$.

Since it does not rely on spontaneous emission, the ARP sequence is not limited by the atom's lifetime; in fact, the ARP sequence needs to be faster than the decay rate. Should an atom be excited by the first light pulse and decay before the second pulse occurs, the second pulse will now excite the atom, reversing the process. To avoid this, these pulse pairs must occur much faster than the decay rate and hence the *rapid* in ARP. The pulses in our experiment each last for $\pi/\omega_m = 3.125$ ns, which is sufficiently small when compared to the excited state lifetime of $\tau \approx 98$ ns.

Unfortunately, there is still a chance that the atom decays before the second pulse arrives, reversing the process and causing the net force to approach 0.

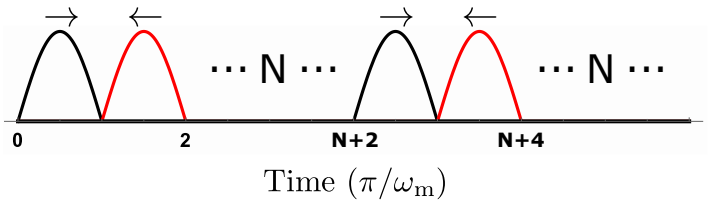

Figure 1.3: The ARP pulse sequence over time, with dead time of N pulses between each pulse pair.

To account for this, one can introduce a dead time between each pulse pair (Fig. 1.3). In the case that the pulse pairs successfully excite and stimulate emission, this dead time will lower the net force felt by the atom. On the other hand, if spontaneous emission occurs, the increased time in the excited state will increase the probability that the atom will decay again before the next pulse pair, correcting the direction of the force. When taking spontaneous emission into account, this dead time increases the net force for large number of pulse pairs. For our experiment we use a dead time of $t_d = 2\pi/\omega_m$ (the length of 2 pulses) because it was calculated to sustain the largest average force after many pulse pairs.

Considering all of this, in a single cycle an atom gains $2\hbar k$ momentum in a time frame of $4\pi/\omega_m$. This results in an ideal force of

$$F_{ARP,ideal} = \frac{2\hbar k}{2\pi/\omega_m} \approx \frac{100\hbar k\gamma}{2\pi} \approx 32F_{Rad}.$$

Taking the possibility of spontaneous emission into account, this ideal force is reduced [5, 6] to

$$F_{ARP,spont} \approx 19F_{Rad}.$$

For a more detailed description of this process see Ref. [7].

Chapter 2

Experimental Setup

In this chapter I will discuss the equipment used to observe the ARP force in helium atoms. I will begin with the lasers used for the experiment and the instrumentation used to create suitable ARP pulses. Following that, I will describe the vacuum chamber, the method for applying the force to metastable helium atoms, and the mechanism used to measure this force.

2.1 ARP Light Production

2.1.1 DL 100 Lasers

The current optical system for our experiment exists largely due to the work of T. Inaki [8]. It begins with three DL 100 laser systems from Toptica Photonics, Inc. tuned to 1083.33 nm. The linearly polarized, collimated outputs of their extended cavities are coupled into Panda type PM fibers via Toptica FiberDocks^{TM}. The wavelength of each laser can be altered in one of three ways: the angle of the internal diffraction grating, the diode current, and the

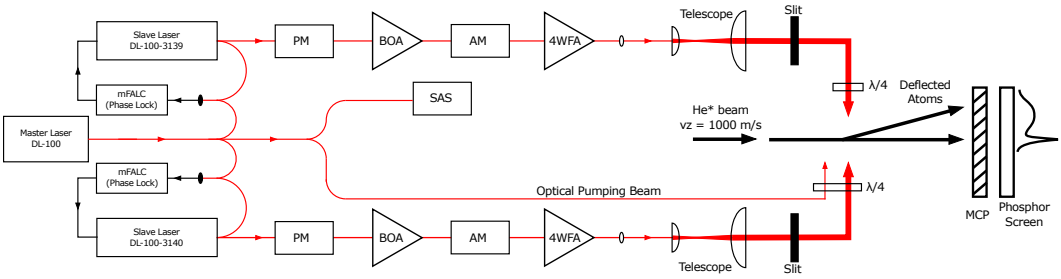

Figure 2.1: A block diagram of the major components for our experiment, including phase modulator (PM), booster optical amplifier (BOA), amplitude modulator (AM), 4 Watt fiber amplifier (4WFA), mixed fast analog linewidth control (mFALC), saturated absorption spectroscopy (SAS), and multi-channel plate (MCP).

diode temperature. The diffraction grating has coarse manual control via a screw and fine electronic control via a piezo-electric transducer (PZT).

2.1.2 Frequency and Phase Locking

The master laser output is split into four directions via fiber splitters, as shown in Fig. 2.1. While one of these paths is used for optically pumping the helium in the experiment, the remaining three paths are used to ensure frequency stability and coherence among the three lasers.

To stabilize the laser frequencies, we use a portion of the master laser light and a helium discharge cell for the common technique of saturated absorption spectroscopy (SAS) locking, shown in Fig. 2.2. When subtracted from the unaltered Doppler-broadened spectrum, we can feed the resulting spectrum into a Toptica DigiLock 110 to generate an error signal. The DigiLock 110 then uses two PID controllers to modulate the current of the laser diode and the voltage to the PZT, which respectively correct for high and low frequency fluctuations.

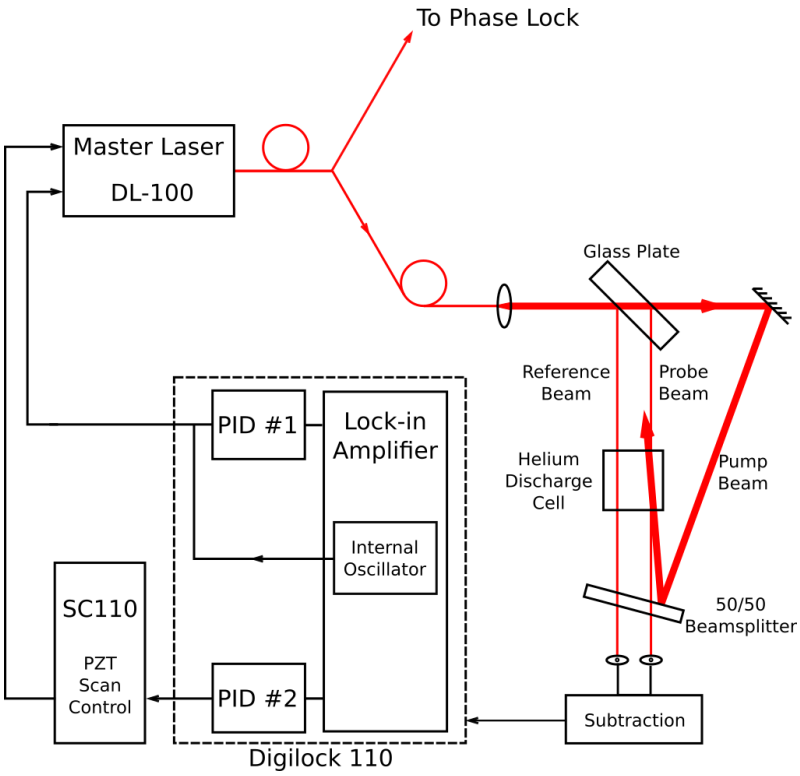

Figure 2.2: Diagram of the saturated absorption spectroscopy setup for frequency locking the master laser.

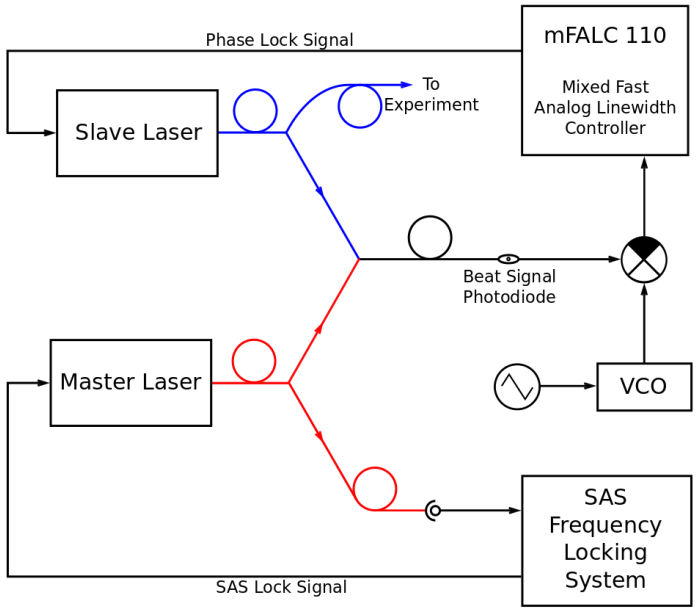

Figure 2.3: Diagram of the phase locking setup for each of the slave lasers, developed by T. Inaki and outlined in [8].

The remaining two paths of the master laser light are directed to separate but identical phase locking setups, one for each slave laser. The phase locking system is shown in Fig. 2.3. Each slave laser's output is split via a 50/50 fiber splitter, with one path continuing toward to the experiment and the other joining the master laser light in the phase lock. In the phase locking setup, the slave and master lasers' outputs are combined into one fiber, where the overlapped beams create a beat note that is equal to the frequency difference of the two lasers. A photodiode then converts this optical beat signal into an electrical one and mixes it with a voltage controlled oscillator (VCO) in a Toptica Mixed Fast Analog Linewidth Calculator (mFALC). This device

provides feedback to the slave laser diode's current, fixing the frequency and phase difference between the two lasers. More on the phase locking process can be seen in Ref. [8].

2.1.3 Electro-Optic Modulators

The electro-optic effect is a phenomenon that occurs in certain materials that describes the change in the index of refraction in the presence of an external electric field. In crystals that lack inversion symmetry, like lithium niobate ($LiNbO_3$), the change in refractive index is linearly proportional to the applied electric field and therefore directionally dependent. Our experiment makes use of two types of electro-optic modulators (EOMs) that utilize this effect to transform the CW light of the slave lasers into the pulsed light with a controlled chirp that the ARP force requires.

Phase Modulators

A phase modulating EOM consists of a parallel plate capacitor with a lithium niobate crystal set inside of it, as shown in Fig. 2.4.a. As light travels trough a waveguide perpendicular to the capacitor's electric field it will experience a phase shift proportional to the voltage across the capacitor. Our phase modulators are responsible for sweeping the frequency of the ARP pulses. The introduced chirps consist of half of a cosine wave symmetrically sweeping through the laser frequency at a modulation frequency, ω_m.

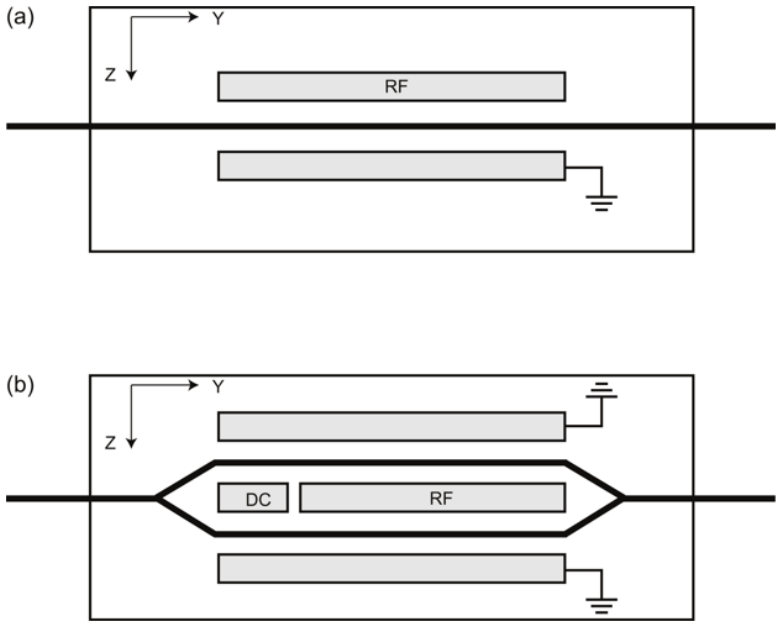

Figure 2.4: Diagrams of two types of lithium niobate electro-optic modulators. (a) The NIR-MPX-LNO3 phase modulator. (b) The NIR-PX-LNO3 amplitude modulator.

Amplitude Modulators

An amplitude modulating EOM consists of the same components as its phase modulating counterpart, but is instead configured as a Mach-Zender interferometer, as shown in Fig. 2.4.b. Incoming light is split into two paths and recombined. Outside of the paths are two grounding electrodes and between the paths are two powered electrodes: one connected to a DC supply and the other to a signal generator. By applying a voltage to the center electrodes, each path will experience an electric field of equal magnitudes but pointed in opposite directions. Due to the inversion asymmetry of lithium niobate, this means that each path will experience opposite phase shifts, $\pm\phi$. This design allows for control of the destructive interference at the recombination point to

determine the output light power.

To create light pulses, we tune the DC voltage to create total destructive interference. The signal generator then generates an RF signal that intermittently creates total constructive interference. The signal generator operates at 320 MHz with 1/4 pulse rate and 17% duty cycle. As a result, the output light from the amplitude modulator is a nearly triangular pulse with a a base width of 3.125 ns and a repetition rate of 80 MHz. The measured pulse train can be seen in Fig. 2.5.

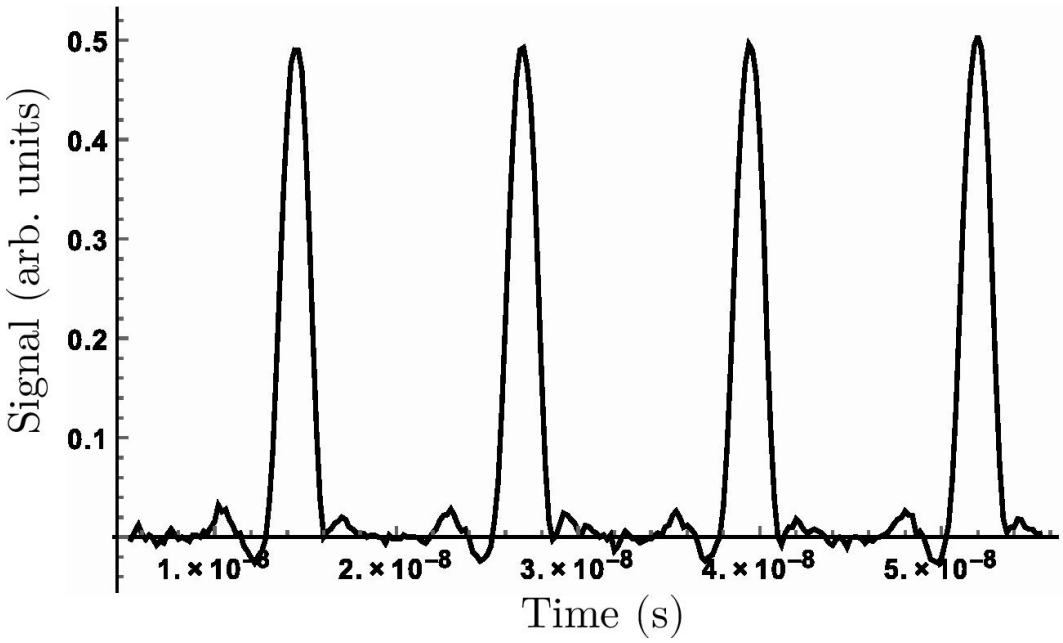

Figure 2.5: Optical pulses used in the ARP experiment as measured by a Thorlabs, Inc. DET08CFC photodiode.

2.1.4 Amplifiers

Between the light source and experimental chamber, the slave lasers undergo a great deal of attenuation. Immediately following the phase modulators, we

send the beams through a pair of Thorlabs BOA1137P booster optical amplifiers (BOAs). The design of these BOAs is similar to a semiconductor laser, but with anti-reflective-coatings to prevent laser operation. These amplifiers are designed to amplify low power signals ($\sim \mu W$) to something easier to work with ($\sim$ mW).

The BOAs provide the small signal gain we need but with a maximum output of $\sim$ 15 mW, we require another boost to the power. To accomplish this we employ a pair of Optocom OI-30-dBm-Yb fiber amplifiers. These utilize Ytterbium doped fibers that are pumped via a diode to amplify light passing through and increase the power of the beam to $\sim$ 70 mW.

Finally, after the amplitude modulators are a pair of Keopsys KPS-BT2-YFA 4 W fiber amplifiers. Similarly to the Optocom fiber amplifiers, these make use of Ytterbium doped fibers for a large power gain. Through these high power amplifiers we can achieve pulses $\sim$ 1 W for the experiment.

2.2 Helium Vacuum System

A full diagram of our vacuum system can be seen in Fig. 2.6. Helium enters the source chamber, where it is excited to the metastable state and then goes to the interaction chamber. From there, the helium beam is collimated and light pulses deflect the atoms. Finally, the deflected beam enters the detection chamber, where its pattern is displayed on a phosphor screen and recorded via a CCD camera.

2.2.1 Source Chamber

Since our experiment focuses on the $2^3S_1 \longleftrightarrow 2^3P_2$ transition of helium, our first step is to excite the atoms to the metastable 2^3S_1 state. Unfortunately, the optical transition $1^2S_0 \rightarrow 2^3S_1$ in helium is doubly forbidden by selection rules, so we resort to exciting the atoms through collision. We accomplish this via a reverse flow discharge modeled after the designs seen in [9] and [10]. The source configuration can be seen in Fig. 2.7.

The helium source itself consists of a 1 cm diameter glass tube mounted in the center of a 3 cm stainless steel jacket cooled by liquid N_2. A tungsten needle, serving as the cathode of the source, is placed in the center of the glass tube. The anode is a grounded aluminum plate with a 0.5 mm aperture placed at the front end of the steel jacket, called the nozzle plate.

When opened, helium gas flows between the jacket and the glass tube towards the tip of the tapered glass tube, cooling as it travels. A portion of the gas sustains a glow discharge between the needle and the nozzle plate as the cathode is set to a high voltage of ~ -2100 V via an HP 6525A. The plasma

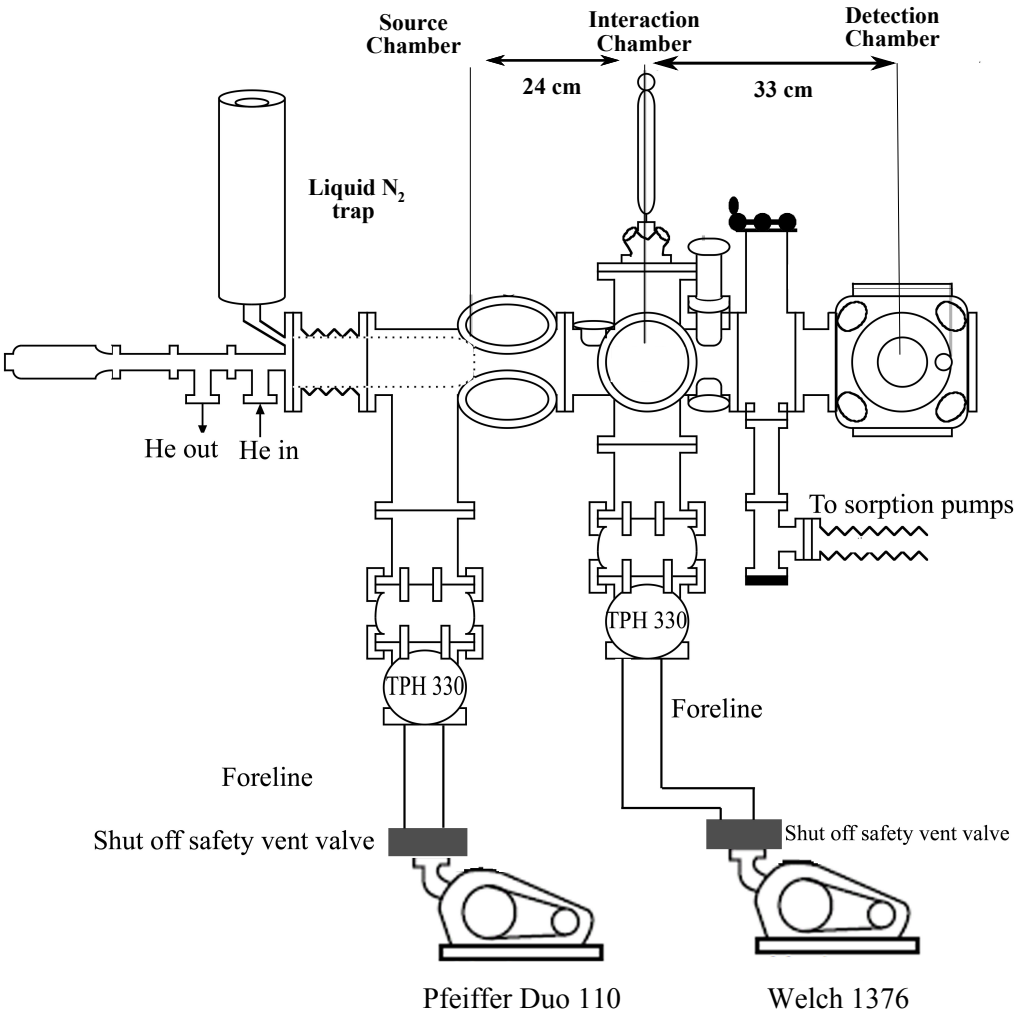

Figure 2.6: Schematic of the vacuum system used in our experiment.

formed creates the desired metastable helium (He*) via electron impact or ion-electron recombination. The produced He* atoms travel through the nozzle plate and skimmer holes, creating an atomic beam with a roughly Maxwellian distribution of longitudinal velocities centered at $\sim$ 1000 m/s with a FWHM

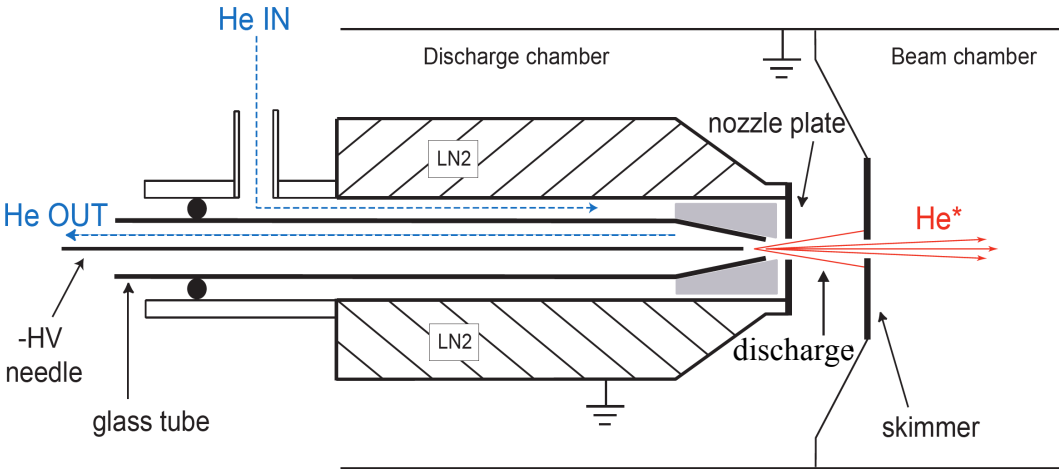

Figure 2.7: A schematic showing production of metastable helium via reverse back flow discharge. Helium gas is pumped into the system and cooled by liquid N_2 as it flows to a high voltage cathode needle. A portion of the gas is excited and accelerated toward the interaction chamber (He*).

of $\sim$ 400 m/s [11]. The source output flux is $\sim 10^{14}$ He* atoms/s·sr [12].

The source chamber has a dedicated Pfeiffer Duo 110 mechanical pump which lowers its pressure to $\sim$ mTorr. This chamber is also equipped with a Pfeiffer Balzers TPH 330 turbo pump, which keeps the pressure in the system $\sim$ µTorr. Since most of the helium is not excited to the metastable state, the back of the glass tube is attached to a Welch 1376 pump. This relieves some of the pumping burden from the TPH 330 and defines the helium flow within the source.

2.2.2 Interaction Chamber

After the skimmer plate, the He* beam travels 24 cm to a 1 cm tall, 250 µm wide slit meant to constrain the transverse velocity of the atoms to $\pm$ 1 m/s. This slit is mounted to the front of a pair of Helmholtz coils, which

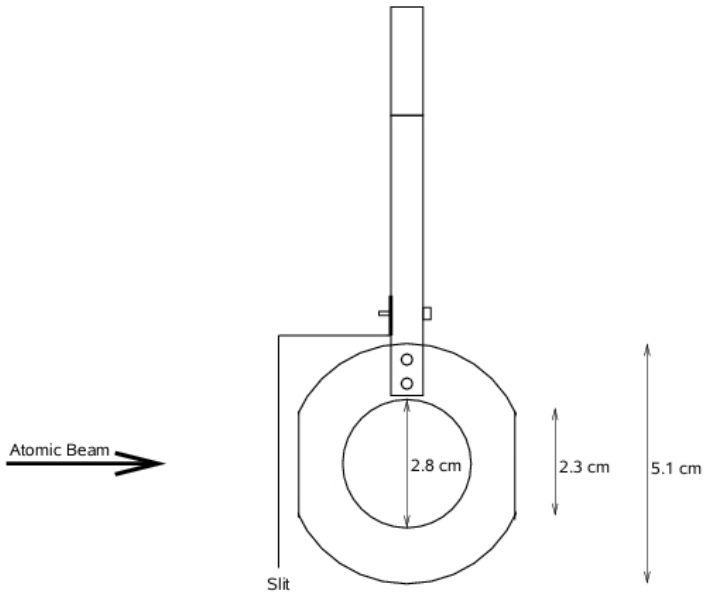

Figure 2.8: A side view of the Helmholtz coil, with slit attached to the front.

provide a uniform magnetic field transverse to the atomic beam and break the Zeeman degeneracy of the 2^3S_1 state. These coils only need to create a field of a few Gauss, which is easily achieved with a few amperes of current. These Helmholtz coils and slit can be seen in Fig. 2.8.

The two slave laser beams enter the interaction chamber in opposite directions through the center of the Helmholtz coils. As they cross the He* beam, they exert a force transverse to the motion of the atoms. This deflects the atoms at an angle proportional to the exerted force.

Similarly to the source chamber, the interaction chamber is pumped by a Pfeiffer Balzers TPH 330 turbo pump, backed by a Welch 1376. This lowers the pressure in the chamber to $\sim$ µTorr.

2.2.3 Detection Chamber

Finally, the He* beam enters the detection chamber, which is displayed in Fig. 2.9. First in the detection chamber is a microchannel plate (MCP), which is a thin, 1" diameter disc comprised of many parallel 10 µm diameter channels. These lead-coated glass channels are arranged in a hexagonal lattice with spacings of 12 µm.

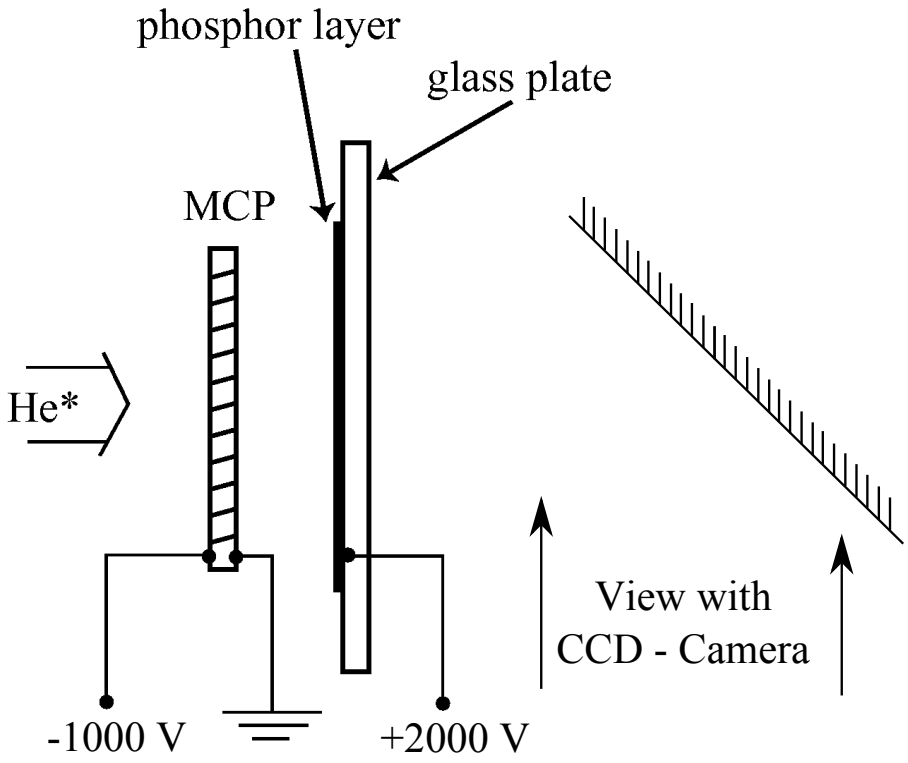

Figure 2.9: A diagram of the detection setup in our experiment, featuring a multi-channel plate (MCP), phosphor screen (PS), and CCD camera (not pictured). When struck, the MCP creates a cascade of electrons directed toward the PS, causing the phosphor to fluoresce. A CCD camera outside of the vacuum system captures this fluorescence pattern.

The He* atoms carry 19.82 eV, which is more than enough to eject an electron as it collides with the metal surface of the MCP. By applying a large

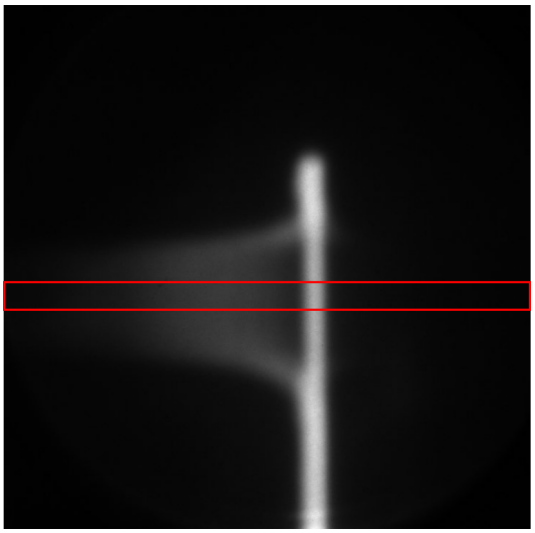

Figure 2.10: A CCD image of the phosphor screen during the experiment. The red box shows the area of interest.

bias of $\sim$ -1000 V across the MCP, the ejected electrons will accelerate and collide within the channel walls, causing an electron avalanche.

Placed directly behind the MCP is a phosphor screen (PS), which consists of a 1.16" diameter glass plate with a thin layer of indium tin oxide and aluminum for conductivity, underneath a layer of P43 phosphor. By applying $\sim$ +2000 V to the screen, we accelerate electrons from the output of the MCP toward the PS. As the electrons collide with the PS, the phosphor layer will fluoresce at the point of impact with a brightness that is related to the electron flux. An angled mirror behind the PS allows us to view this fluorescence pattern outside of the chamber through a window, and we capture it via a CCD camera. A CCD image from a recent experiment can be seen in Fig. 2.10.

This detector is not only sensitive to He*, but also the UV light that is emitted from the helium source chamber, which passes through the slits and

follows the same path that an undeflected He* beam would. By comparing this to the pattern of He* atoms that are deflected by the laser pulses, we can determine the magnitude of deflection and thus measure the magnitude of the ARP force on the beam.

The detection chamber is connected to the interaction chamber by a gate valve and is kept at a low pressure through this opening. The valve can be sealed, allowing maintenance to the detector without opening the interaction and source chambers to atmospheric pressure. When this is the case, we use two sorption pumps to lower the detection chamber pressure to $\sim$ mTorr before slowly opening the gate valve again.

Chapter 3

EOM Drifts

As discussed in 2.1.3, our experiment utilizes two EOMs to transform our CW light into pulses and their precision is critical. Unfortunately, there are potential drawbacks that come along with EOMs (temperature drifts, aging effects, etc.) and our devices are not immune to these issues. In this chapter, I will discuss further the operation of an amplitude modulating EOM, the drifts that can occur during operation, and the proposed solution to efficiently correct for these drifts.

3.1 Operating Points

Amplitude modulating EOMs are typically kept at one of four voltage dependent operating points: MAX, MIN, increasing quadrature (QUAD+), and decreasing quadrature (QUAD−). The MAX point refers to complete constructive interference of the two paths, while the MIN point refers to a π phase difference between the branches (i.e. $\pm\frac{\pi}{2}$ phase shift in each branch)

and thus complete destructive interference when the paths recombine. The QUAD+ and QUAD− points refer to the inflection points halfway between the MAX and MIN, either on the increasing or decreasing side. These points are shown with respect to the DC bias voltage in Fig. 3.1.

In the case of our experiment, the DC bias puts the amplitude modulators into the MIN point, creating total destructive interference. The RF signal then applies a voltage of V_π to intermittently switch to the MAX point. This creates optical pulses.

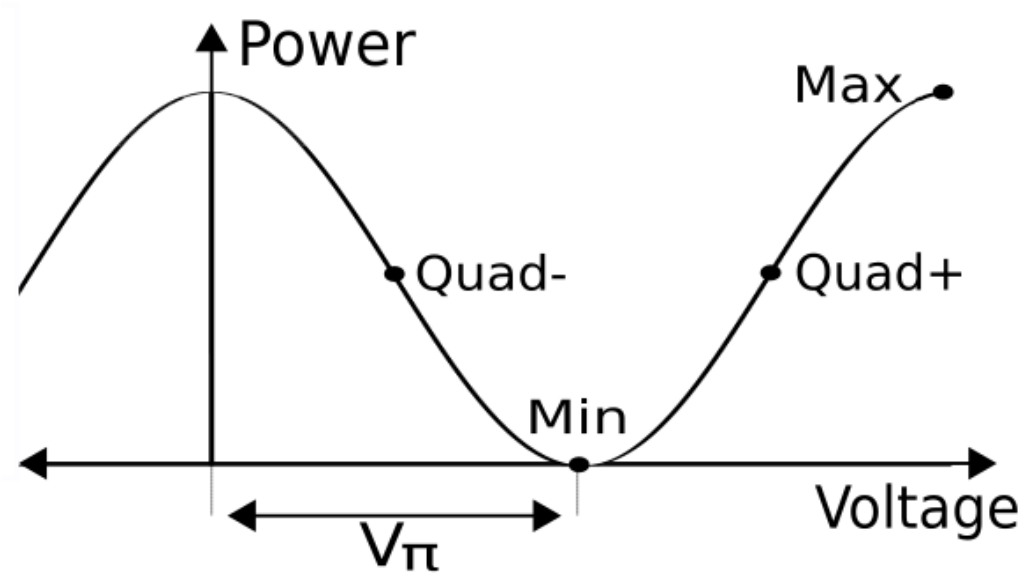

Figure 3.1: Output power from an amplitude modulator for a given voltage, and the 4 common operating points for the device. V_π is the voltage required for a 180° phase shift.

3.2 Drifts

As in a traditional Mach-Zender interferometer, small changes in the path length of an EOM (e.g. due to thermal expansion/contraction) can drastically affect the end interference. While these path lengths change, the voltage required to achieve a certain type of interference does, too.

In our lab, this means that static DC bias will no longer keep the modulator at the MIN point and that the modulator's operating point will drift along the curve shown in Fig. 3.1. As the RF signal applies a V_π voltage, it will no longer put the modulator at the MAX point, but some distance away proportional to the drift. For pulse production, these drifts will lower the power of the pulses and increase the amount of light during the dead time. These drifts are detrimental to our experiment and cause noticeable effects in just minutes. Should they go uncorrected for long enough, the modulation can have no effect and the EOM will output CW light.

The simple solution to such a drift is to modify the DC bias voltage to compensate. In the past we did this by manually controlling the DC power supply in response to changes in the optical pulses. This process was slow, required near constant attention, and often meant that data needed to be collected more than once. By automating the DC bias correction process, these issues can be eliminated.

3.3 Theoretical Solution

Automating the DC bias control is, in theory, a straightforward process. One can tap a small portion of the EOM output and use it as an error signal for a PID algorithm. When the light power changes, the algorithm can correct the input voltage. However, when operating at the MIN point, as in our case, an issue presents itself: whether the DC voltage is too high or too low, the output power will increase. This is a consequence of the fact that the photodiode measures intensity, which is proportional to the square of the electric field in the modulator. The solution to this problem is to add an AC dither tone to the DC bias voltage to help pinpoint the operating point. A method of ratioing dither tone harmonics such that the error signal accomplishes this while remaining independent of input optical power is outlined in [13].

The output optical power from an EOM, P_{out}, can be expressed as

$$P_{out} = \frac{T_D P_{in}}{2}\big(1 + cos(\phi_0 + \Delta\phi)\big) \tag{3.1}$$

where T_D is the transmission coefficient when biased for maximum transmission, P_{in} is the input optical power, ϕ_0 is the DC phase bias, and $\Delta\phi$ is the phase shift resulting from an AC signal. For a sinusoidal dither signal at angular frequency ω,

$$\Delta\phi = \frac{\pi V}{V_\pi(\omega)} sin(\omega t) \tag{3.2}$$

where V is the AC signal amplitude and $V_\pi(\omega)$ is the half-wave voltage as a

function of frequency, as shown in Fig. 3.1. Substituting (3.2) into (3.1) yields

$$P_{out} = \frac{T_D P_{in}}{2}\left(1 + cos\big(\phi_0 + \alpha sin(\omega t)\big)\right) \tag{3.3}$$

where $\alpha = \pi V / V_\pi(\omega)$. For small enough dither, where $\alpha \ll 1$, (3.3) can be expanded to

$$P_{out} = \frac{T_D P_{in}}{2}\Big(1 + cos\phi_0 \cdot cos\big(\alpha sin(\omega t)\big) - sin\phi_0 \cdot sin\big(\alpha sin(\omega t)\big)\Big).$$

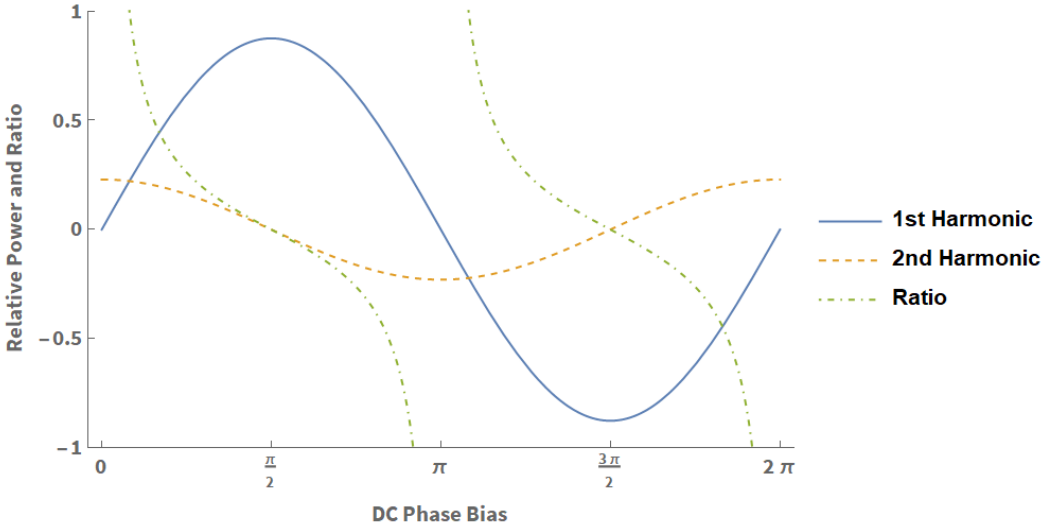

Figure 3.2: The simulated signals and their ratio for varying DC phase bias. Negative and positive represent 180° phase changes relative to the input dither signal.

Applying a fourth order Taylor expansion to this expression yields

$$P_{out} = \frac{T_D P_{in}}{2}\left(1 + cos\phi_0\left(1 - \frac{\alpha^2}{4} + \frac{\alpha^4}{64} + \left(\frac{\alpha^2}{4} + \frac{\alpha^4}{48}\right)cos(2\omega t) + \frac{\alpha^4}{92}cos(4\omega t)\right)\right.$$

$$-sin\phi_0 \left(\left(\alpha - \frac{\alpha^3}{8} \right) sin(\omega t) \cdot \frac{\alpha^3}{24} sin(3\omega t) \right) \Bigg).$$

The optical signal will be measured via a photodiode, and from the above expression the corresponding photocurrents at first and second order will be given by

$$I_1 = \eta \frac{T_D P_{in}}{2} sin\phi_0 \left(\frac{\alpha^3}{8} - \alpha \right)$$

$$I_2 = \eta \frac{T_D P_{in}}{2} cos\phi_0 \left(\frac{\alpha^2}{4} - \frac{\alpha^4}{48} \right),$$

where η denotes the photodiode responsivity. Though much simpler than the previous expression, these equations depend on input light power, which is subject to change. Taking the ratio of the first and second order photocurrents, on the other hand, yields

$$R = \frac{I_2}{I_1} = cot\phi_0 \frac{\left(\frac{\alpha^2}{4} - \frac{\alpha^4}{48} \right)}{\left(\frac{\alpha^3}{8} - \alpha \right)}, \tag{3.4}$$

which depends only on the characteristics of the dither signal, half-wave voltage, and phase angle. Additionally, due to the asymptotic nature of the cotangent function at 0° and 180°, R is highly sensitive to deviations from the MIN bias point, making it a very reliable feedback parameter. The individual signals and their ratio is simulated in Fig. 3.2.

Therefore, by applying an AC dither to the DC bias of an EOM and isolating the harmonics in the output light, we can pinpoint the EOM's operating point. This information can be used as the input to a PID algorithm to efficiently correct for any drifts in the EOM from the desired locking point.

Chapter 4

Modulator Bias Controller

The device used to control the drifts in our EOM, called a modulator bias controller (MBC), can be seen in 4.1 and a diagram of the entire locking system can be found in 4.2. For signal processing purposes, a local oscillator generates an AC signal, which acts as the dither. Its output is combined with a DC signal and input to the EOM's DC bias port. When light exits the modulator, a small portion of it is tapped off and measured via a photodetector. The resulting electronic signal is amplified and split between two band pass filters, which select either the fundamental or second harmonic of the dither signal. These signals are then converted to digital ones and input to a Raspberry Pi 4 (RPi), where a PID algorithm corrects for any drifts in the EOM and corrects the DC bias. The circuit diagram of the electronics used in the controller can be found in Appendix A.

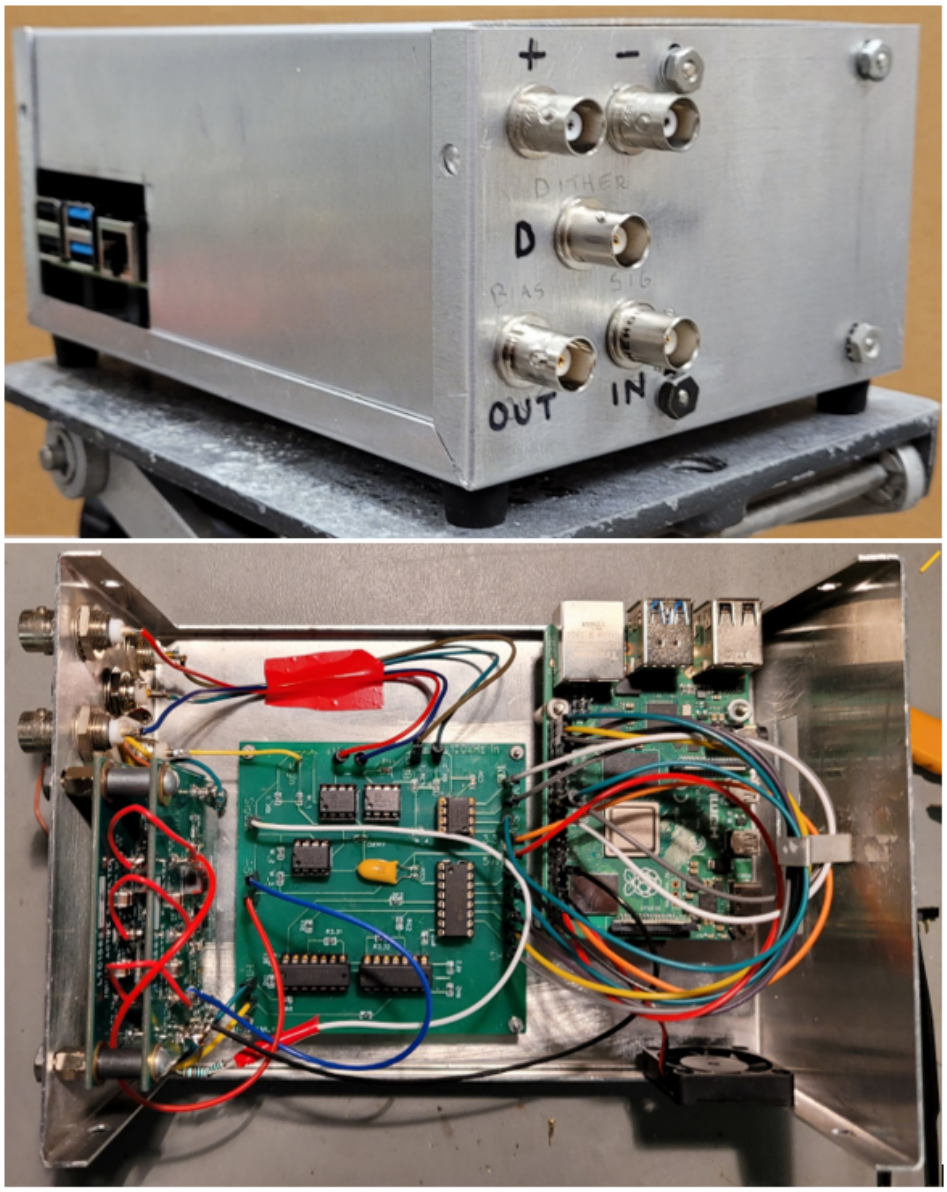

Figure 4.1: Photographs of the MBC. (Top) The +/- ports are for the DC supply to power the device, the D port is for the 100 kHz dither signal input, Out is for the output bias signal to be input to the EOM, and In is for the detected optical signal from the EOM as measured by the photodetector. (Bottom) is the MBC with the top of the case removed. On the right is the RPi; in the middle is the circuit containing the amplifiers, ADC, and DAC; and on the left wall are the band pass filters.

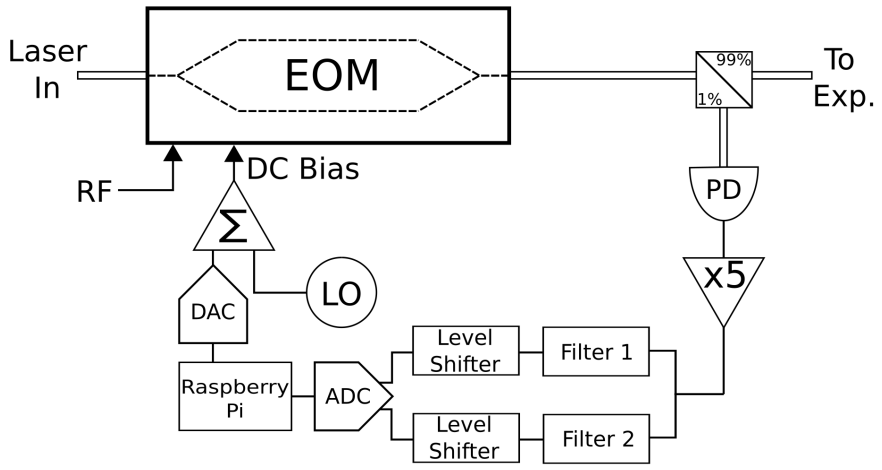

Figure 4.2: Block diagram of the modulator controller including photodetector (PD), 5 times gain amplifier (denoted by x5), analog-digital converter (ADC), digital-analog converter (DAC), local oscillator (LO), and summing amplifier (Σ).

4.1 EOM Input Voltage

In the past, the bias voltage sent to the EOM was generated with an analog DC power supply, but for the purposes of automation the DC bias is now generated digitally by a Raspberry Pi 4B (RPi). While microcontrollers and field-programmable gate arrays (FPGAs) can offer many technical advantages, the RPi boasts approachable programming, affordability, and terrific documentation that helped it stand apart from other options.

Since the RPi is capable of outputting only digital signals (i.e. signals of either 0V or 3.3V), its output must first be converted to an analog signal. This is accomplished via the MCP4911, a 10-bit digital-analog converter. This IC allows us to create a DC voltage in the 0-3.3V range. An external function generator creates a 100 kHz dither signal, which is combined with the DC bias

signal via a summing op-amp. This amplifier is configured with a gain of 4.7x, creating an AC signal of 500 mV_{pp} with a DC offset of up to 15.5 V with ∼15 mV increments. This dithered bias signal is then input into the EOM DC port.

4.2 EOM Output Signal

The output of the EOM is pulsed light modulated by the relatively small dither signal. Immediately following the modulator is a 99/1 fiber splitter. The 1% branch of this splitter leads to a Thorlabs DET08CFC photodetector, which converts the light signal into an electrical current. Since we are already looking at such a small portion of the EOM output, the direct output of the photodetector can be too small to effectively work with. To compensate for this, the detector's output leads directly to an amplifier with a gain of ∼ 5x.

After this amplification, the signal is split and sent to two LTC1562-2 fourth order band pass filters mounted on DC266B-B demonstration boards. In its default configuration the DC266B-B has a center frequency of 200 kHz, allowing it to select the second harmonic. By replacing the resistors on the second board, the other filter has been tuned to a center frequency of 100 kHz allowing it to select only the fundamental dither signal. The schematic for the default board and its transfer function can be seen in Fig. 4.3.

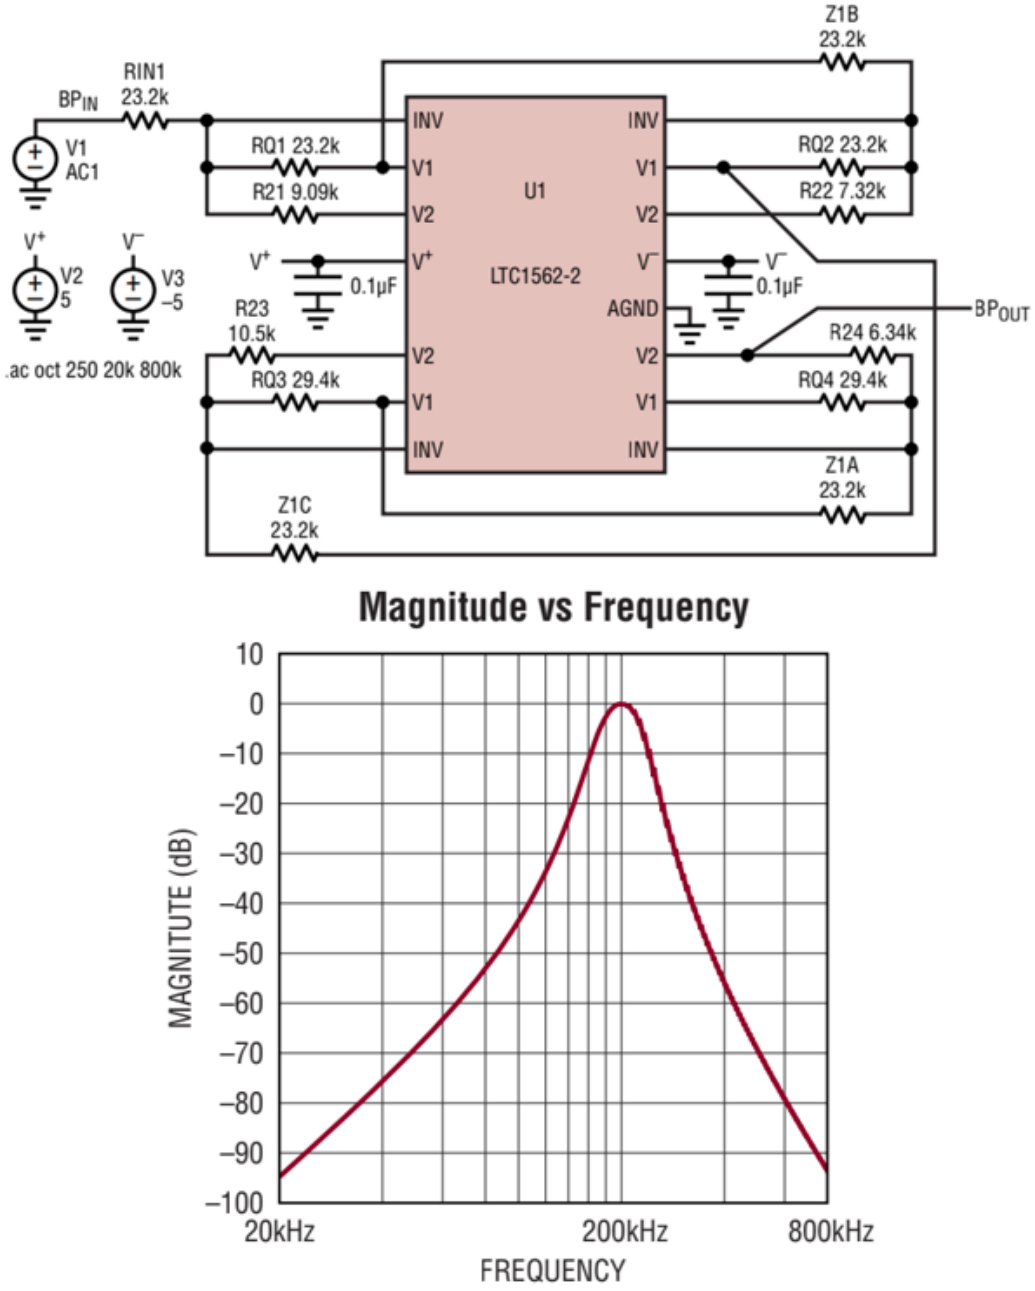

Figure 4.3: The schematic for the default DC266B-B demo board used as the 200 kHz band pass filter (Top) and its transfer function (Bottom). By replacing each of the resistors we created a 100 kHz bandpass filter using the same board.

4.3 Reading the Signals

Unfortunately, the simplicity of the RPi does come at several costs. The first of these is the voltage limit: its general purpose input/output (GPIO) can only support positive voltages smaller than 3.3 V. To ensure that the signal is contained within these bounds, the filter outputs are scaled using voltage level shifters before continuing to an analog-to-digital converter (ADC).

The ADC used is an MCP3008. In converting back to a digital signal, one can see the second major drawback of the RPi: when powered by the RPi's 3.3 V, the sample rate of the MCP3008 is limited to 75 ks/s. By the Nyquist-Shannon sampling theorem, this is not nearly enough to record a 100 kHz signal, let alone a 200 kHz signal. Fortunately, the method described in Sec. 3.3 requires only the amplitude of the signals, which can be estimated within ∼14% even with the slow sample rate.

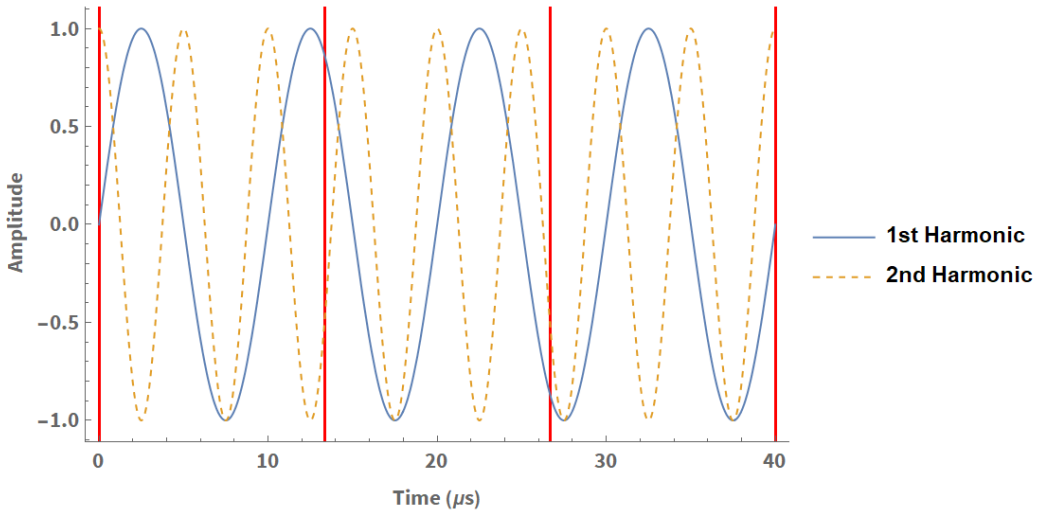

Figure 4.4: The sampling points (red vertical lines) as it compares with the two harmonic signals.

Since the ADC samples at a rate of 75 kHz, it will ideally measure the same point in the harmonic every 4th sample (See Fig. 4.4). Considering this and assuming that one of the samples occurs at one of the extrema of a harmonic, then one needs to sample 3 times and pick the largest value to measure the amplitude of the signal. It is worth noting that should the sampling be perfectly synced with the maxima of one of the harmonics that it will measure the minima with nearly 50% error, and vice versa. However, one only needs to measure one of the extrema to determine the amplitude of the signal and by selecting the value with the largest magnitude one can find the amplitude.

Unfortunately, the sample rate is likely out of sync with one (or both) of the signals by some phase, φ, especially since the two harmonics are 90° out of phase from one another by default. As the phase difference between the sample rate and signal changes so, too, does the accuracy of the measurement.

There is a limit to this inaccuracy, though: as one sample moves away from one extrema, another sample approaches another extrema (Fig. 4.5). For the fundamental frequency, this limit is reached at $\varphi = n \cdot \frac{\pi}{3}$ for $n \in \mathbb{Z}$, where measurement of the amplitude has an error of 13.4% (See Fig. 4.6). The 2nd Harmonic has the same maximum error, but for a phase difference of $\frac{\pi}{12}(m+1)$ for $m \in \mathbb{Z}$. Therefore by sampling each signal in groups of 3 and selecting the measurement of maximum magnitude, one can find the amplitude of the signal with an error of $< 14\%$.

To compensate for this inaccuracy and any noise, the RPi collects the amplitude in this way 100 times and averages over the values before passing the measurement onto the PID algorithm. Of course, this method provides infor-

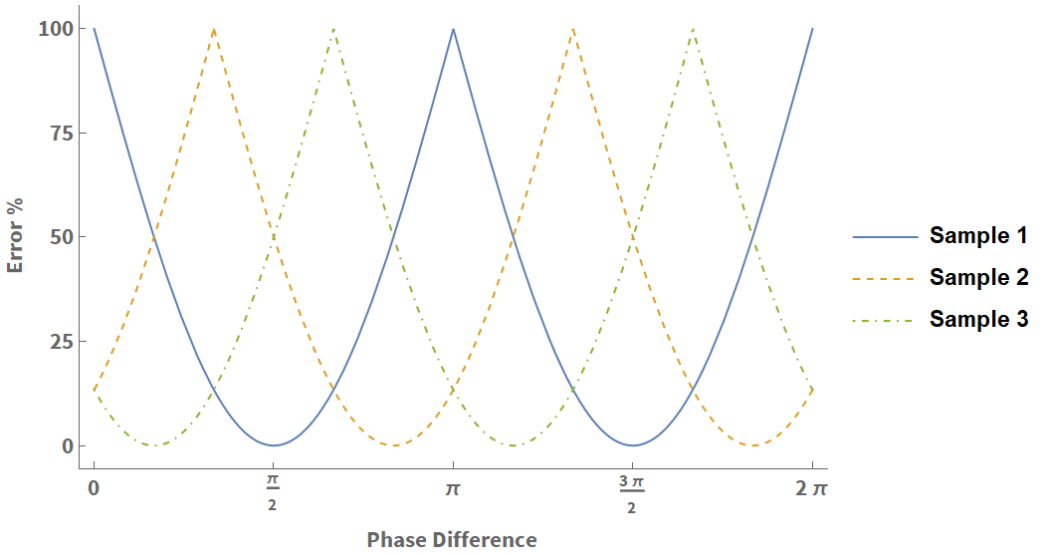

Figure 4.5: The distance between each of the three samples and its nearest extrema for varying phase difference between the sample rate and the 100kHz signal.

mation on the amplitudes of the harmonics without any phase information. This means that despite the generality of the theory discussed in Sec. 3.3, the device is unable to measure the EOM's exact operating point in a single sample in its current state. Luckily even with the increased computation time, the drifts occur on a much longer time scale and the device can still fulfill its purpose.

4.4 PID Program

Though the RPi is equipped with a 1.5 GHz CPU, its serial peripheral interface is limited to a fraction of that [14]. Further, the limited speed of the ADC and DAC mean that the improved speed of a compiled language like C is wasted when engaging with the GPIO. This means that utilizing a script language

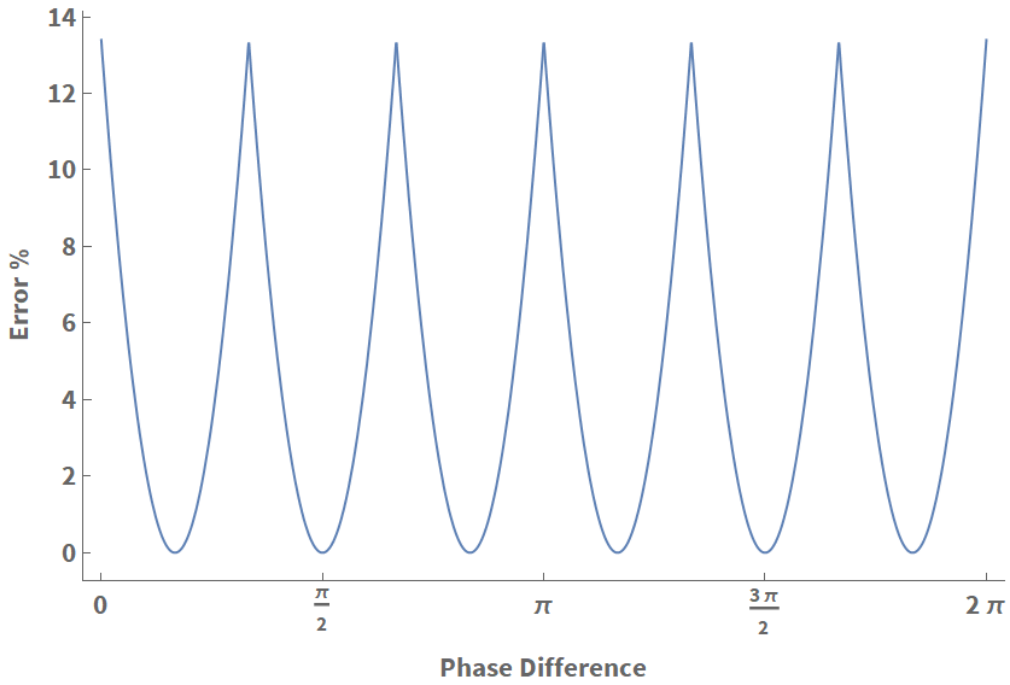

Figure 4.6: The inaccuracy of each group of samples for varying phase difference between the sample rate and the 100 kHz signal. Error reaches a maximum of 13.4% when the phase is a multiple of $\pi/3$.

like Python comes at little cost for the increased simplicity. Considering this, the program written on the RPi to interact with the ADC and DAC as well as execute the PID algorithm is written in Python. The program in full can be seen in Appendix B.

When started, the program tells the DAC to output a starting voltage of 6 V before beginning the PID loop. In the loop, the program first reads in the values of the harmonics from the MCP3008 using the method described in Sec. 4.3. It then calculates the ratio of these signals, converts it from a bit value into a voltage, and using this voltage and Eq. 3.4 finds the phase deviation from the desired operating point. Measurements of these signals and their ratios over a range of bias voltages can be seen in Fig. 4.7. Since

we are using the EOM at the MIN operating point, these voltage and phase measurements can then be used directly as the error value for a traditional PID algorithm. The program then calculates a correction to the DC bias, alters the DAC output correspondingly, and repeats the process.

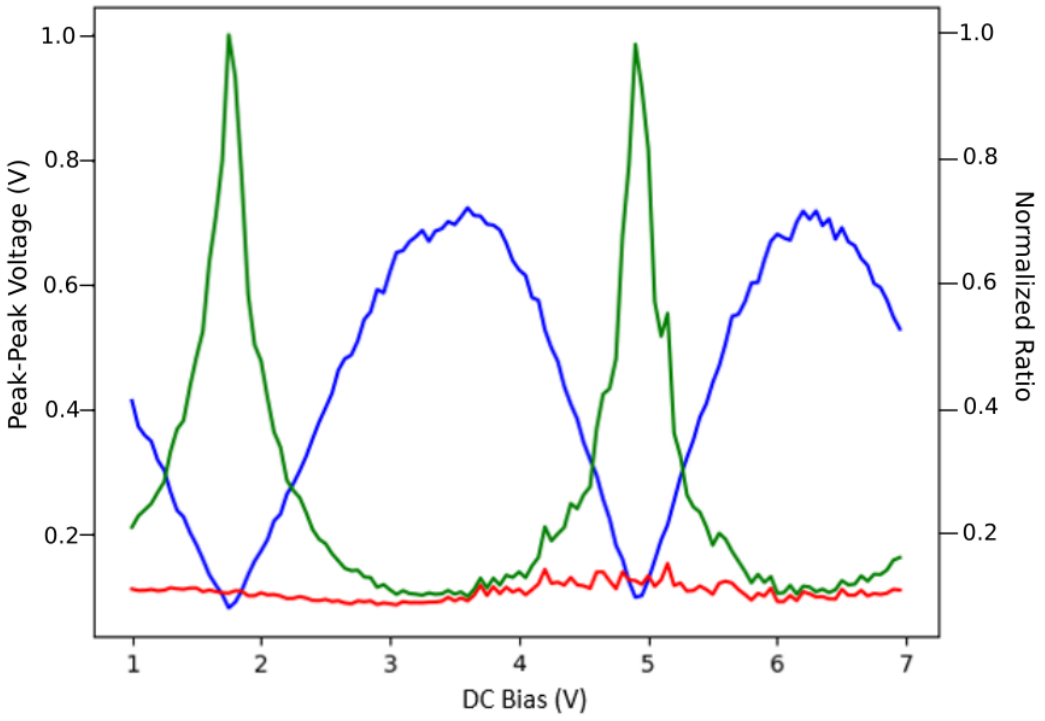

Figure 4.7: Measurements of the 1st harmonic (Blue), 2nd harmonic (Red), and their normalized ratio (Green).

Due to the $\sim$ 14% error window when reading the signals, the program would often continue to make larger than necessary corrections to the DC bias even when very close to the desired locking point. To avoid this, we introduced a tolerance value. This tolerance would represent an allowable error in the locking point and was selected to be large enough that the PID could easily reach the value while remaining small enough that uncorrected drifts would not negatively impact the experiment. We chose this tolerance to be 60 mV.

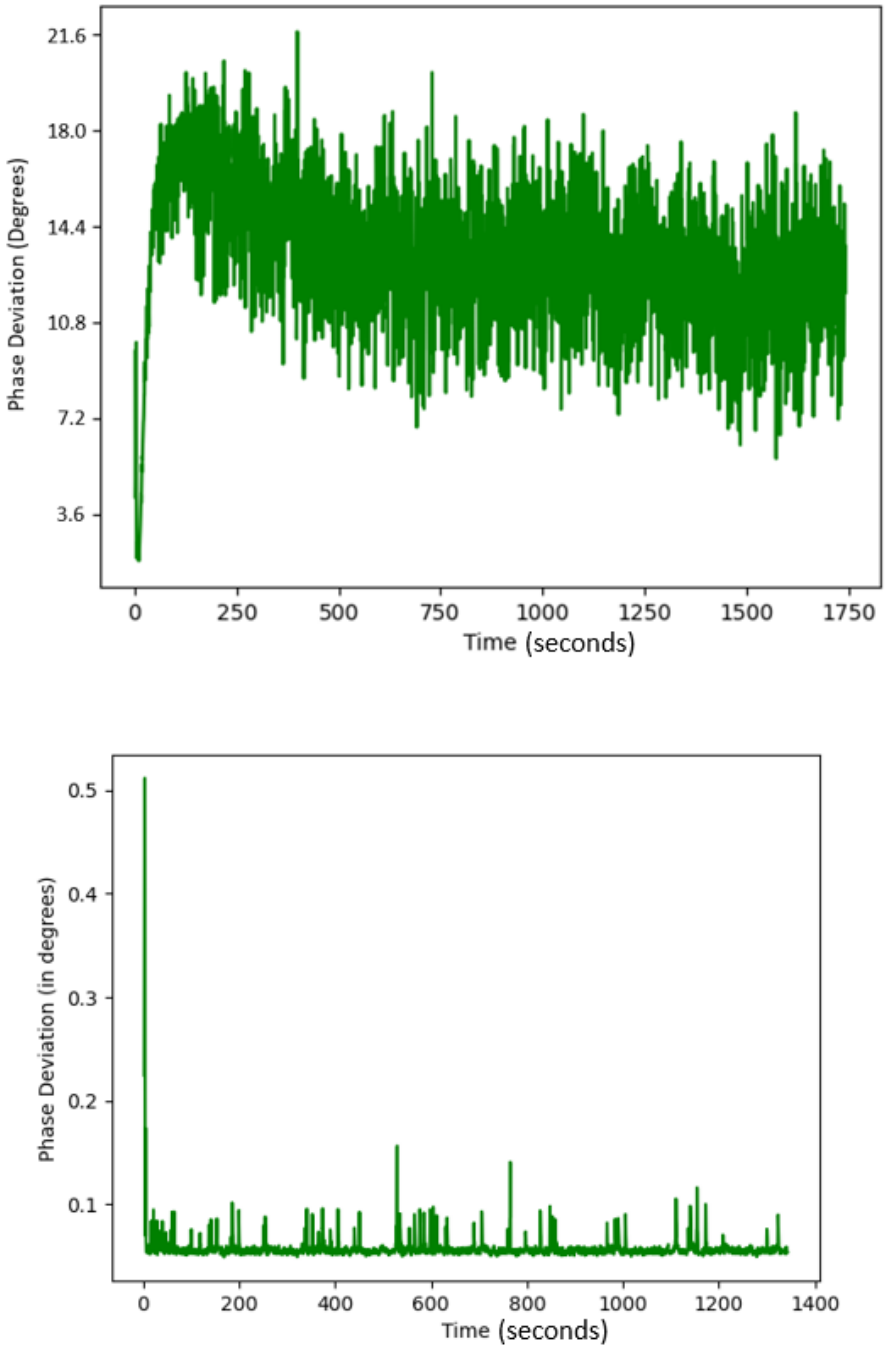

Figure 4.8: The distance from the desired point locking point with the PID deactivated (Top) and with the PID activated (Bottom).

With the controller activated, we are able to find the desired locking point in under a minute and remain within 0.15° of the phase point for arbitrarily long. The effects of the PID can be seen in Fig. 4.8. Since the previous solution required near constant human attention and could affect the results of the experiment, the controller is a vast improvement.

Chapter 5

Outlook

While the modulator bias controller described in this thesis is a notable improvement upon the previous way of doing things, there are still several ways to upgrade the experiment and the MBC itself. In this chapter, I will discuss potential hardware upgrades to improve things.

5.1 Methods of Improvement

As in any feedback system, there are two factors that can always be improved: speed and accuracy. For this MBC, both of these factors are affected by the current sampling rate. One of the most obvious ways to improve the sampling rate would be to replace the MCP3008 with a faster ADC, but this would present its own challenges. Because of the way the RPi interacts with its GPIO, even ADCs with higher sampling rates would not be able to consistently guarantee faster speeds. Without an experienced programmer to modify the timing within the RPi, this direction will not be straightforward.

The next logical idea would be to replace the RPi itself, due to its many limitations. The RPi was selected because of its versatility as this project evolved but now that the system is well established, replacing it with a faster device would be advantageous. Many other devices even include integrated ADCs and DACs, on top of increased speed. In doing this, though, one would need to replace the circuitry of the MBC and rewrite the PID program as they are designed around the RPi.

Rather than speeding up the sampling rate, another option would be to simply lower the frequency of the dither signal. The current 100 kHz dither frequency was selected after we found the LTC1562-2 filters. By replacing these with a pair of lower frequency band pass filters, the dither frequency could be lowered and the sampling issues could be avoided. In fact, if the second harmonic frequency can be less than 37.5 kHz (i.e. $1/2 \times 75$ kHz), then the system could digitally recreate the harmonic waveform without any further modifications.

There is one promising approach that can be made without any hardware modifications that was suggested by Edoardo Buonocore, a current MSI student. It is mentioned in 4.3 that sampling more than 3 times is redundant because the sampling rate and dither signal are periodic and the ADC's fourth sample will have the same result as its first. In other words, because 75 and 100 share the common multiple of 300, the sample and dither will sync back up every fourth sample. By lowering the sampling rate to a number that has a larger lowest common multiple with the dither frequency, one can sample many more points on the waveform. For example, with a sample rate of 61 kHz, the ADC could sample 99 times before sampling the same point on the

signal twice. In this case, one of the samples would measure the amplitude of the signal with $< 1\%$ inaccuracy. Though a single measurement would take more time to sample, this increased accuracy would eliminate the need for averaging and ultimately take the same amount of processing time. Edoardo is currently testing this method.

5.2 Project Expansion

This modulator bias controller was used only on one of four EOMs used in our experiment, but the drifts are not unique to this modulator. The natural next step is to expand this method to the other amplitude modulator. Now that this method is tested, creating a second controller should be much more straightforward and many of the shortcomings of this MBC can be avoided.

Appendices

Appendix A

Circuit Board Schematic

Below are the schematics of the circuit used for teh modulator bias controller. This circuit acts contains all amplifiers in the device and connects the Raspberry Pi 4B, function generator, EOM, and band pass filters. H1 and H2 refer to the first and second harmonic signals output from the band pass filters, respectively. Dither In refers to the function generator output. Bias Out is the combined DC bias and dither signal as input to the EOM. Signal In refers to the EOM output and Signal Out is the amplified EOM output. CLK, Dout, Din, SHDN/CS, SDI, CS, and SCK refer to the digital signals from the Raspberry Pi SPI. Decoupling capacitors are used near DC inputs and near each op-amp to filter high frequency noise.

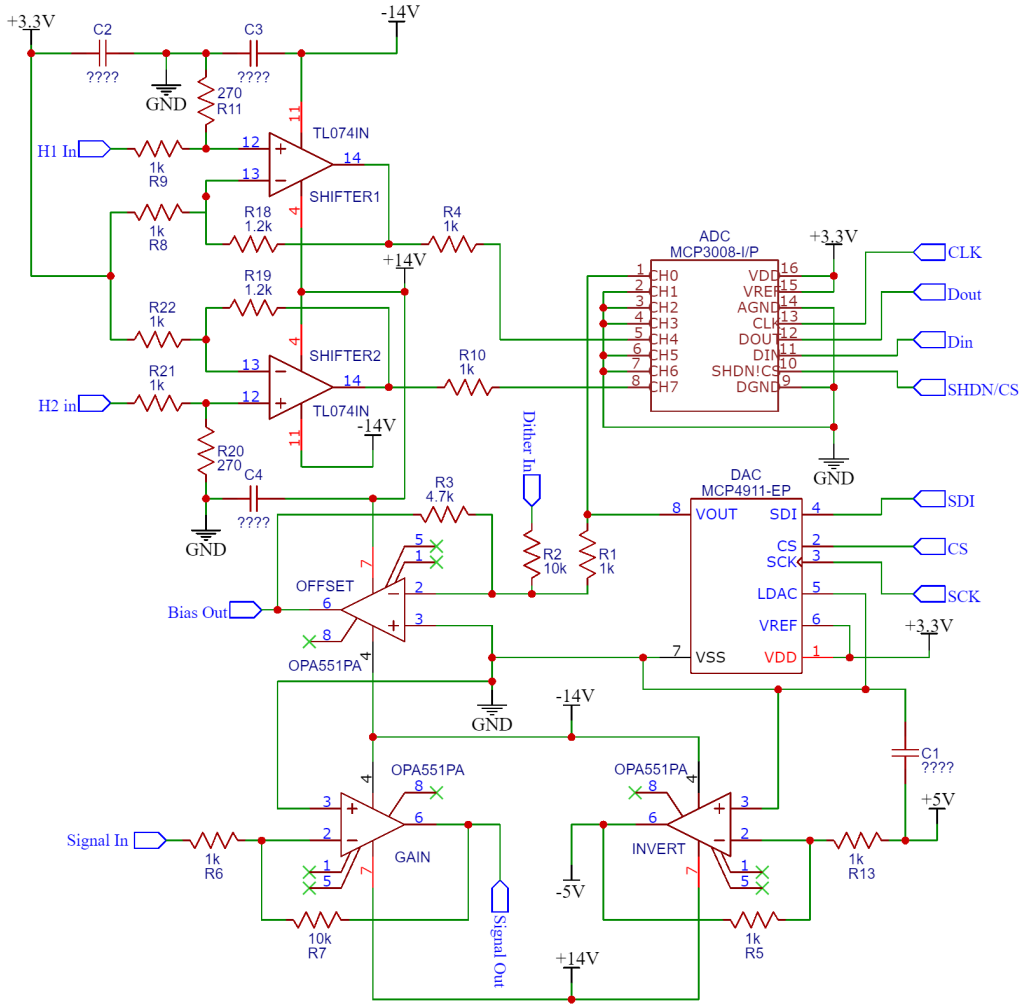

Figure A.1: A circuit diagram for the electronics in the modulator bias controller.

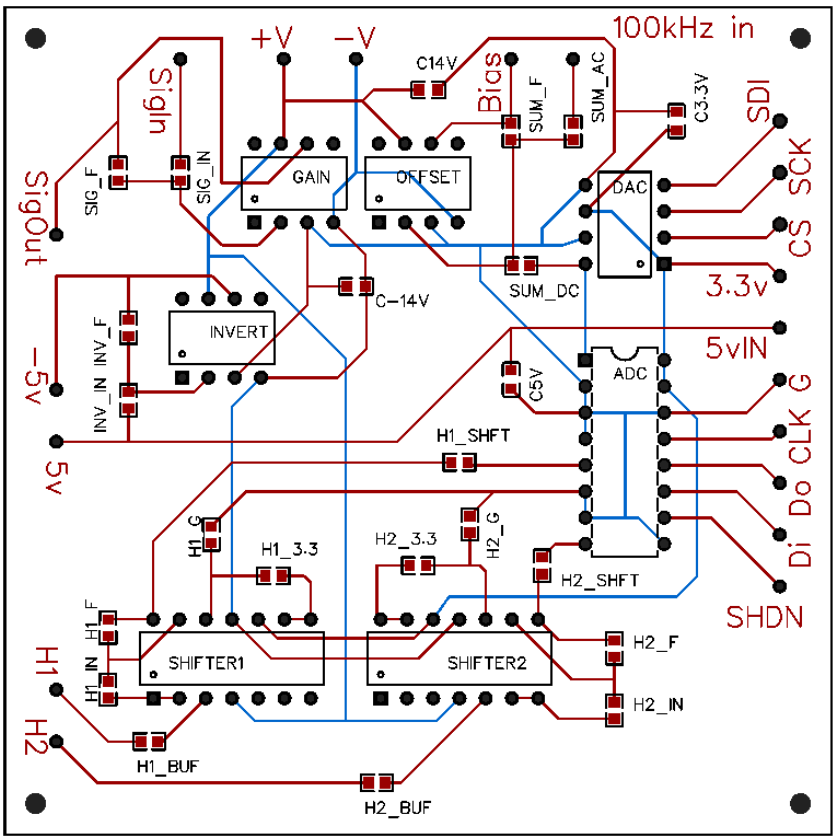

Figure A.2: The PCB schematic for circuit shown in A.1. The top layer is in red while the bottom layer is blue.

Appendix B

Python Code

Included here is the code used on the Raspberry Pi to interact with the peripheral electronics and execute the PID algorithm. While this program is my own creation, the imported libraries are not and are under The MIT License (MIT), which is quoted here:

The MIT License (MIT)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

```
import time
import sys
import RPi.GPIO as GPIO
import math
import spidev
import busio
import digitalio
import board
import adafruit_mcp3xxx.mcp3008 as MCP
from adafruit_mcp3xxx.analog_in import AnalogIn
from matplotlib import pyplot as plt
import numpy as np
import mpmath as mp
from DAC.MCP4922 import MCP4922
from pynput import keyboard
import time

#ADC -- MCP3008 stuff
spi1 = busio.SPI(clock=21, MISO=19, MOSI=20)
cs = digitalio.DigitalInOut(board.D16)
adc = MCP.MCP3008(spi1, cs)
#Record from the DAC and each filter
DACout = AnalogIn(adc, MCP.P0)
h1in = AnalogIn(adc, MCP.P4)
h2in = AnalogIn(adc, MCP.P7)

#DAC -- MCP4911 stuff
GPIO.setmode(GPIO.BCM)              # use the Broadcom pin
```

```
                                numbering
GPIO.setwarnings(False)                 # disable warnings
dac = MCP4922(spibus=0, spidevice=0, cs=8)

#SPI -- GPIO clock stuff
spi0 = spidev.SpiDev()
spi0.open(0,0)
spi0.max_speed_hz = 125000000

#Skip to the next min bias point when user presses spacebar.
                                Resets locking point
def on_press(key):
    global V
    global lastError
    global I
    if key == keyboard.Key.space:
        if V >= 9:
            V = 1
        else:
            V += 3.5
    lastError = 0.0
    I = 0.0

#Track key presses
listener = keyboard.Listener(on_press=on_press)
listener.start()

#Measures the amplitudes of the signals
def read(h1in, h2in,avg_num):
    Vpp1 = np.array([])
```

```
    Vpp2 = np.array([])
    for a in range(avg_num):
        h1 = []
        h2 = []
        #Measure the voltage 4 times
        for b in range(4):
            h1.append(h1in.voltage)
            h2.append(h2in.voltage)
        #Estimate the amplitude
        Vpp1 = np.append(Vpp1, max([max(h1) - 1.54, 1.54 - min(
                                        h1)]))
        Vpp2 = np.append(Vpp2, max([max(h2) - 1.29, 1.29 - min(
                                        h2)]))
    #Average over avg_num measurements and find the ratio
    return np.mean(Vpp2)/np.mean(Vpp1)

#Initializations
#Bias voltage
V = 6
#PID Stuff
lastError = 0.0
lastde = 0
lastTime = 0.0
P = 0.0
Kp = 2.0
I = 0.0
Ki = 0.05
D = 0.0
Kd = 0.1
windupGuard = 1.0
```

```
errorp = []
timep = []
try:
    while True:
        #Find the error ratio
        ratio = read(h1in, h2in, 100)
        #Convert to a phase angle
        phase = mp.acot( ratio / 0.0788668)
        #Convert to volts
        error = (phase)*(5 / mp.pi)

        # THE PID ALGORITHM #

        if lastError == 0.0:
            de = 0.0
        else:
            de = error - lastError
        lastError = error

        currentTime = time.process_time()
        dt = currentTime - lastTime
        lastTime = currentTime

        errorp.append(error)
        timep.append(currentTime)
        if error < 0.13:
            if Kp > 0:
                Kp = 1.0
            else:
                Kp = -1.0
```

```
elif Kp > 0:
    Kp = 2.0
else:
    Kp = -2.0

#Only correct if the error is larger than 60 mV
if error > 0.06:
    P = error
    #A work-around since error is always positive
    #If it got further from the set point, then work in
                                        the other
                                       direction
    if de > 0:
        I -= error*dt
        #If it got further twice in a row, then flip P
        if lastde > 0:
            Kp = -Kp
    else:
        I += error * dt

    if (I < -windupGuard):
        I = -windupGuard
    elif (I > windupGuard):
        I = windupGuard

    if (dt > 0):
        D = de/dt
    else:
        D = 0
```

```
            V += Kp * P + Ki * I + Kd * D

        print("Error: " + str(error) + " | P " + str( Kp*P ) +
                                        " | I: " + str( Ki*I )
                                        + " | D: " + str ( Kd*D
                                         ) )
        if V > 12:
            V = 2
            print("Postive")
        elif V < 0:
            V = 10
            print("Negative")

        #Convert to something the DAC can put out (0-3.3V)
        x = -0.21595*(0.0474-V)
        dac.setVoltage(0,int(4095*(x/3.3)))
        #Wait for the EOM to fully respond
        #time.sleep(0.1)
        lastde = de

except KeyboardInterrupt:    # Press CTRL C to exit program
    dac.setVoltage(0, 0)
    plt.plot(timep,errorp, "g", label="error")
    plt.xlabel("Time")
    plt.ylabel("Phase Deviation (in degrees)")
    plt.show()
    dac.shutdown(0)
    GPIO.cleanup()
    sys.exit(0)
```

Bibliography

[1] E.B Treasy. *Adiabatic Inversion with Light Pulses*. Phys. Lett. A, 27:421, 1968.

[2] M.M.T. Loy. *Observation of Population Inversion by Optical Adiabatic Rapid Passage*. Phys. Rev. Lett., 32:814–817, 1974.

[3] P. van der Straten and H. Matcalf. *Atoms and Molecules Interacting with Light*. Cambridge University Press, Cambridge, 2016.

[4] V.S. Malinovsky and J.L. Krause. *General Theory of Population Transfer by Adiabatic Rapid Passage with Intense, Chirped Laser Pulses*. European Physical Journal D 14, 147–155 (2001)

[5] D. Stack, J. Elgin, P.M. Anisimov, and H. Metcalf. *Numerical studies of optical forces from adiabatic rapid passage*. Phys. Rev. A, 84:013420, 2011.

[6] X. Miao. *Optical Forces on Atoms with Periodic Adiabatic Rapid Passage Sequences*. PhD thesis, Stony Brook University, 2006.

[7] B. Arnold. *Velocity Dependence of the Adiabatic Rapid Passage Force in Metastable Helium*. PhD thesis, Stony Brook University, 2019.

[8] T. Inaki. *Optical phase coherence in the adiabatic rapid passage force.* Master's thesis, Stony Brook University, 2017.

[9] J. Kawanaka, M. Hagiuda, K. Shimizu, F. Shimizu, and H. Takuma. 56(1):21–24, 1993.

[10] H. Mastwijk. Cold Collisions of Metastable Helium Atoms. PhD thesis, Utrecht University, 1997.

[11] D. Stack. Optical Forces from Periodic Adiabatic Rapid Passage Sequences on Metastable Helium Atoms. PhD thesis, Stony Brook University, 2012.

[12] M. Cashen. Optical Forces on Atoms in Polychromatic Light Fields. PhD thesis, Stony Brook University, 2002.

[13] L. L. Wang, T. Kowalcyzk. *A Versatile Bias Control Technique for Any-Point Locking in Lithium Niobate Mach–Zehnder Modulators.* Journal of Lightwave Technology, Vol. 28, 2010.

[14] Raspberry Pi Foundation. *Raspberry Pi Documentation.* https://www.raspberrypi.org/documentation/, 2021.