Development of a Vector Magnetometer based on Electromagnetically Induced Transparency in
^{87}Rb Atomic Vapor

A thesis submitted in partial fulfillment of the requirement
for the degree of Bachelor of Science with Honors in
Physics from William & Mary

by

Alexander M. Toyryla

Accepted for _____ *Honors* _____

_____*Irina Novikova*_____
Irina Novikova, Advisor

_____
David Armstrong

_____
Jonathan Frey

Williamsburg, VA
May 4, 2022

Contents

Acknowledgments

List of Figures

List of Tables

Abstract

We present progress towards the development of an atomic magnetometer capable of accurate scalar and vector magnetic field measurements with high sensitivity and no need for external calibration. The proposed device will use the interaction between a bi-chromatic laser field and rubidium vapor to derive magnetic field magnitude and direction from measured amplitudes of Electromagnetically Induced Transparency (EIT) resonances. Since the proposed method requires precision control of light polarization, we observe the performance capabilities of a liquid crystal device to dynamically rotate the polarization of the laser field. Another goal in this project is to establish a polarization locking mechanism that tracks the magnetic field's azimuthal angle for streamlined measurement. Finally, we realize methods to derive the field's polar angle from comparisons of EIT resonance strengths. The work completed herein will inform the greater project's noise reduction, component manufacturing, and device refinement steps that aim to produce the first compact, accurate, and unobtrusive vector magnetometer.

Chapter 1

Introduction

1.1 Motivation for the Magnetometer

Humans have been interested in the precise measurements of the magnitude and direction of a magnetic field for hundreds of years, but the instruments and methods used have often been inherently imperfect. Hall probes and search coils use effects like Lorentz forces and magnetic induction to spatially dislocate electrons in metals, producing measurable voltages indicative of magnetic field strength. Magneto-resistive sensors use spin-dependent tunneling through magnetic thin films and are more sensitive than Hall probes. The most sensitive magnetometers we have today are the superconducting quantum interference device (SQUID) and the atomic magnetometer; both take advantage of quantum effects to translate magnetic fields into measurable quantities [1].

Today, atomic magnetometers are among the few we truly consider accurate as they rely on fundamental constants of nature like the atomic g-factor (measured with high accuracy for most atoms) and the Bohr magneton. Other instruments rely too heavily on uncontrolled variables or material properties. However, many types of atomic magnetometers are scalar devices, sensitive only to the total magnetic field value, and cannot determine its vector components, which is crucial information for

a plethora of magnetometer applications in aircraft, spacecraft, geological surveys, positioning systems, and submarine vision [2]. Many modifications to the atomic magnetometer have been implemented, but all attempts to date either degrade accuracy or apply an external field which can interfere with nearby sensors.

This thesis is a part of a longer-term project aimed to demonstrate a vector atomic magnetometer operating at Earth-like magnetic field and capable of measuring magnetic field strength with sub-picotesla precision. For these experiments we will use Rubidium-87 (^{87}Rb) atoms interacting with a bi-chromatic laser field whose frequency will be tuned to match two optical transitions in Rubidium atoms. Under these conditions physicists have observed several spectrally narrow resonances in the laser transmission [3]. This effect is called Electromagnetically Induced Transparency, or EIT, and is caused by quantum interference of the multiple excitation pathways of a valence electron in Rb atoms [4]. Knowing the position of each resonance allows us to precisely determine the magnitude of the external magnetic field, while the relative amplitudes of the peaks provide information regarding the orientation of this field with respect to the propagation direction and the polarization of the laser beam.

The unique advantage of this approach is that no external calibration is required; the magnetic field orientation is measured relative to the laser light polarization and propagation vector. Contrasting with existing devices, this project combines scalar accuracy, long-term stability, vector measurement capability, and the absence of interfering applied fields into a single sensing unit. In the future, such a unit could be manufactured as a compact, chip-scale device that would facilitate magnetic object measurement and location tracking.

1.2 The Goals of this Thesis

My role in this project is establishing, demonstrating, and refining methods to measure the precise direction of the magnetic field. Analysis of electromagnetically induced transparency (EIT) resonances and data-driven characterization allows for precise methods of determining the two angles required to define field direction [3, 5]. The end goal is to tune the setup such that variable factors have no effect on the ability to extract precise information from EIT resonances; the methods demonstrated solely depend on magnetic field properties. When achieved, this work ensures this magnetometer's stability and accuracy even in the presence of technical noise and component fluctuations.

I first evaluated the performance and capabilities of a voltage-driven polarization rotator essential to field direction measurement. An automatic measurement device cannot include a manually-rotated half-wave plate, so I examined the range, reliability, response time, input methods, and drawbacks of a liquid crystal device which acts as a half-wave plate using a markedly different polarization method.

Next, a vital step in measuring field angle was the construction of a feedback loop that finds and locks the polarization rotator to an angle with known angular distance to the field's azimuthal angle. This locking mechanism took preparation and good knowledge of the system to install, but it serves to mitigate the effects of outside factors by dynamically holding the polarization in plane with the field and laser vector. Automating a method to ensure predictable angles between the device's laser and the field created a functional, reliable field direction measurement protocol.

Finally, establishing this proof-of-concept device necessarily uncovered limitations that needed to be overcome. Since this device will be improved over the next few years, subsequent researchers require a complete understanding of the drawbacks

faced like unwanted polarization rotation and loss of angular sensitivity when the field is oriented in specific directions. I explored and mitigated many device drawbacks, however I showed one perceived issue of additional atomic polarization rotation provided a method to solve a directional symmetry problem that this development faces.

This work will inform the greater project's subsequent noise reduction, component manufacturing, and device refinement steps that aim to produce this vector magnetometer.

Chapter 2

Theory

2.1 Rubidium Vapor as the System

Rubidium is an alkali metal. This means it has a single valence electron in an S-orbital, much like hydrogen; the difference is that we can approximate the nucleus and all 36 other electrons in the rubidium atom as being the hydrogen-like nucleus. It thus resembles hydrogen in its interactions with light and magnetic fields. Another advantage of a hydrogen-like atom is the absence of multiple electrons in its outermost orbital, meaning interactions with light are predictable and free of complex inter-particle interactions. The single electron is easily excited.

Rubidium has several other desirable properties for the development of this device. Rubidium metal is easily vaporized at temperatures not far from room temperature, allowing easy production of vapor to send a laser through. The wavelength absorption range of rubidium is convenient as well, sitting in the close-to-visible infrared spectrum. Many of the advantages of rubidium are shared by caesium, and a similar device could be produced using caesium and a laser tuned to its relevant resonant frequency. The reason for choosing rubidium over caesium simply lies in the lab's experience and equipment tailored for work with rubidium.

2.2   Coherent Population Trapping

The central concept for this magnetometer's operation is electromagnetically-induced transparency (EIT) made possible through coherent population trapping (CPT) in rubidium vapor. The three-level system, or Λ-system, is a fantastic model for rubidium atoms and can be seen in Fig 2.1. An electromagnetic (EM) field, such as that produced by a laser, can excite atoms from a ground state to an excited state through the absorption of a photon of a frequency $\omega_{EM(1,2)}$ matching the difference between the two states, like ω_1 or ω_2 in the diagram. These are called resonance frequencies. An atom in an excited state can spontaneously emit a photon of the exact frequency difference between the excited state and some lower state, thus returning the atom to that lower state. The emitted photon could travel in any direction and thus has a negligible effect on laser transmission measurements through atomic vapor.
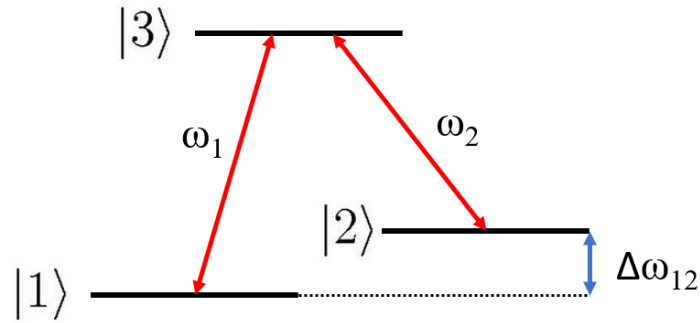
Figure 2.1:   A generic Λ-system. Energy levels 1 and 2, in our case, are states in $5S_{1/2}$ for rubidium, separated by a frequency difference $\Delta\omega_{12}$ (6.834 GHz in our case). Energy level 3 is an excited state.

The probability that an atom is in the excited state $|3\rangle$ can be manipulated, however, by placing an atom in a two-photon resonance, which involves applying two EM fields of different frequencies to the atom. If the two EM fields are close to the resonant frequencies of the two transitions, the difference between the EM field and resonant frequency is $\Delta_s = \omega_{EM(s)} - \omega_s$ where $s = (1, 2)$ is the specific ground state.

Given two incident fields, if the difference between these "detunings" is equal, the atom is placed in a quantum superposition of the two ground states called a "dark state" in which the atom does not interact with either EM field. This state does not follow the same visual intuition as the states in Fig. 2.1 and thus cannot be pictured. Mathematically, the condition

$$\Delta_1 - \Delta_2 = (\omega_{EM(1)} - \omega_{EM(2)}) - (\omega_1 - \omega_2) = 0 \tag{2.1}$$

must be satisfied [4]. Equation 2.1 shows it is equivalent to the difference between the EM fields being approximately equal to the difference between the resonant frequencies.

Since atoms in the dark state cannot interact with the EM field, they have been "trapped" and CPT causes electromagnetically-induced transparency. The EM field passes through the Λ-system with near zero chance of absorption, meaning there is nearly 100% transmission. Around this pair of frequencies, transmission increases sharply, forming what will be referred to as an EIT peak.

2.3   Zeeman Shifts and Scalar Field Measurement

Several magnetic energy sub-levels F exist in the ground state of the rubidium atom; this is called hyperfine splitting. The Zeeman effect serves to split these hyperfine levels into several extra levels, with an energy difference linearly proportional to the static magnetic field applied to the system, $\Delta E_m = m\hbar\gamma B$. Here m is the azimuthal quantum number, $\hbar$ is the reduced Planck's constant, B is the magnetic field strength, and γ is a constant based on other atomic effects. EIT peaks occur between Zeeman sublevels with $\Delta m = 0, \pm 1, \pm 2$, where Δm is the difference in m between two given levels.

[87]Rb has the energy level structure shown in Fig. 2.2 that allows seven EIT peaks
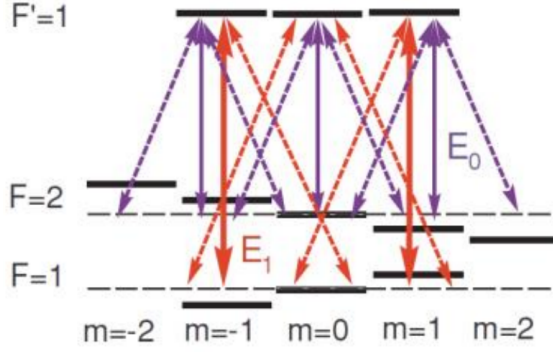
Figure 2.2: Atomic levels in the ^{87}Rb $5S_{1/2} \to 5P_{1/2}$ optical transition, interacting with two EM fields E_0 and E_1 with frequencies ω_m separated by ground-state hyperfine splitting. Arrows represent possible transitions between Zeeman sublevels. The horizontal dashed lines indicate the unshifted positions of Zeeman sublevels, $F = 1, 2$. Solid arrows represent transitions that do not change the m quantum number, whereas dashed arrows represent transitions with a ± 1 change in m. From reference [3].

caused by CPT. Optical transitions may only occur between Zeeman sublevels with $\Delta m = 0, \pm 1$, so the diagram tells us it is possible for a given Λ-system ground state energy level difference of up to $\Delta E_m = 3\hbar\gamma B$. These seven peaks are separated in frequency from their neighboring peaks by $\Delta\nu = \gamma B$ ($\hbar$ has been absorbed by γ since they are both constants). Therefore, we arrive at a method to measure the magnetic field strength. The frequency separation of EIT peaks does not depend on the direction of the field, so this is a reliable method of measuring the scalar magnitude of any external magnetic field. An example of these EIT peaks found after a laser radio-frequency (RF) sweep is in Fig. 2.3.

Field direction affects the coupling strength between magnetic sublevels and therefore changes the amplitude of EIT peaks. Previous studies [3], however, have demonstrated the existence of two permanent peaks corresponding to $a_{\pm 2}$ that do not disappear for any field orientation. When the laser polarization $\vec{E}$ is in-plane with the plane formed by the magnetic field $\vec{B}$ and the laser vector $\vec{k}$, quantum selection rules dictate that $\Delta m = 0$ transitions are allowed (the solid vertical arrows in Fig. 2.2); this
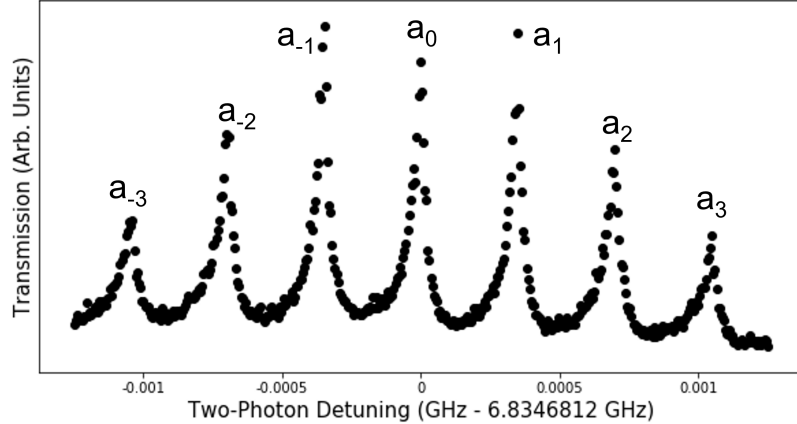
Figure 2.3: Arrangement of seven EIT resonances evenly spaced in frequency $\Delta\nu = \gamma B$. Each is labeled a_n where n indicates the peak's relative position to the central resonance, separated by $\Delta\nu = n\gamma B$ where $n = 0, \pm 1, \pm 2, \pm 3$. The frequency value in the x-label is the approximate frequency of the central resonance, 6.834 GHz.

restricts us to Λ-systems that are an even number of levels off of hyperfine splitting, the a_0 and the $a_{\pm 2}$ peaks. However, an optical transition between states with the same F number and $m = 0$ is forbidden by quantum selection rules, and thus only the $a_{\pm 2}$ peaks exist in this regime [3]. These peaks can be used for reliable scalar field measurements.

2.4   Field Direction Measurement

The end of Section 2.3 explains that EIT peaks change in amplitude based on the orientation of the magnetic field. This phenomenon presents a unique opportunity to measure the direction of the field without a need for external orientation calibration. The EIT resonance amplitudes depend on the relative orientations of three vectors: the laser wave-vector $\vec{k}$, the laser field polarization $\vec{E}$, and the magnetic field direction $\vec{B}$ [6, 7]. These three vectors are geometrically represented in Fig. 2.4.

Transitions between energy levels with the same m (solid arrows in Fig. 2.2) are solely enabled by the laser polarization component along the magnetic field. Tran-
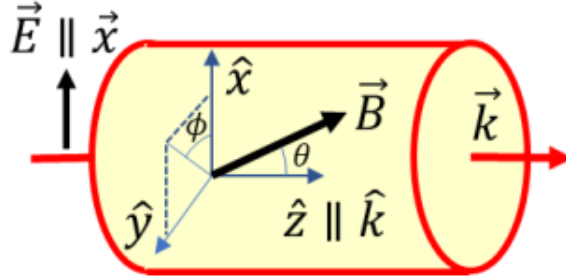
Figure 2.4: Field geometry in the vapor cell. $\vec{k}$ is the laser wave-vector, always oriented along the central axis of the vapor cell. $\vec{B}$ is the magnetic field vector, an angle θ off of the $\vec{k}$ vector. $\vec{E}$ is the laser's polarization, and ϕ is the angle between $\vec{E}$ and plane formed by $\vec{B}$ and $\vec{k}$. From reference [3].

sitions between energy levels with $\Delta m = \pm 1$ (dashed arrows in Fig. 2.2) are solely enabled by the remaining orthogonal component. Theoretically, if we examine the possible two-photon transitions and all their combinations, we can calculate the resulting amplitude of each EIT resonance as a function of the angles between laser propagation, polarization, and the magnetic field direction vectors.
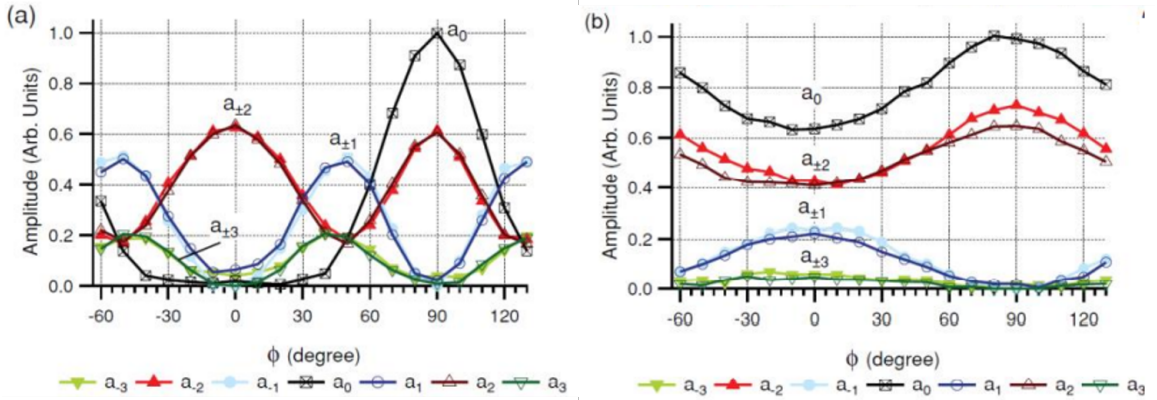
Figure 2.5: Experimental Dependence of EIT resonance amplitudes on polarization angle ϕ in a magnetic field at angle (a) $\theta = 90°$ and (b) $\theta = 15°$ with respect to the laser wave-vector $\vec{k}$. Curves are labelled to match their corresponding EIT peak(s) in Fig. 2.3. From reference [3].

Fig. 2.5(a) visualizes this dependence with experimental results [3]. It depicts a sweep of ϕ while the magnetic field is oriented at angle $\theta = 90°$. The amplitudes for

10

$a_{\pm 1}$ and $a_{\pm 3}$ are only nonzero when both polarization vector components parallel to and orthogonal to the $\vec{B}$ field component in plane with polarization (the xy-plane in Fig. 2.4) are nonzero. The central peak a_0 has zero amplitude when polarization is parallel to $\vec{B}$'s polarization plane component. The peaks $a_{\pm 2}$, as discussed at the end of Section 2.3, always have nonzero amplitude, making them vital for measuring the scalar magnitude of the field in any direction.

We see that EIT peak amplitudes exhibit universal extrema whenever the light polarization is parallel or perpendicular to the magnetic field's polarization plane component. This forms an important basis for measuring magnetic field direction that is free from dependence on other experimental parameters like laser power. Thus, knowledge of the polarization direction ϕ and laser propagation direction $\vec{k}$ can tell us about the azimuthal angle of the magnetic field if we compare the relative amplitudes of measured EIT resonance peaks.

In Fig. 2.5(b), the relative angle between the laser propagation $\vec{k}$ and magnetic field direction $\vec{B}$ has been reduced to $\theta = 15°$. The overall variation in each of these amplitudes has decreased significantly. Once the $\vec{B}$ field is parallel to the laser vector $\vec{k}$, each of these curves flattens entirely. Therefore, by measuring the variation in the amplitude of a single CPT resonance when ϕ changes, we can establish the angle θ between $\vec{k}$ and $\vec{B}$.

In summary, with the scalar magnitude of the field established through the separation of EIT peaks, the azimuthal angle ϕ established through the relative amplitudes of the EIT peaks, and the angle θ between the magnetic field and laser propagation direction established through a resonance's response to changing polarization, this method fully characterizes any constant magnetic field vector $\vec{B}$.

Chapter 3

Experimental Arrangement

Before a detailed explanation of each component of the magnetometer, here is a brief overview. Our laser is first split in two directions, the secondary direction being towards the dichronic atomic vapor laser lock (DAVLL) which acts as a feedback loop holding one laser frequency on-resonance while we sweep the other to observe EIT. The primary laser path goes through the Liquid Crystal Variable Retarder (LCVR) which polarizes the light based on a modulating voltage we send in with a function generator. The laser is split once again, where the secondary direction this time goes to a polarization detector. The primary arm continues through the main ^{87}Rb cell surrounded by magnetic-field-producing coils and housed in magnetic shielding. The remaining light goes to the post-cell photodetector, which can measure the full transmission and polarization rotation simultaneously. Data from both detectors is recorded by a computer, but the transmission signal is additionally fed into the lock-in amplifier and feedback controller which find where the amplitude of the current EIT peak is maximized and supplies an additional voltage to the LCVR to ensure this maximum. In-depth descriptions of each component follow.
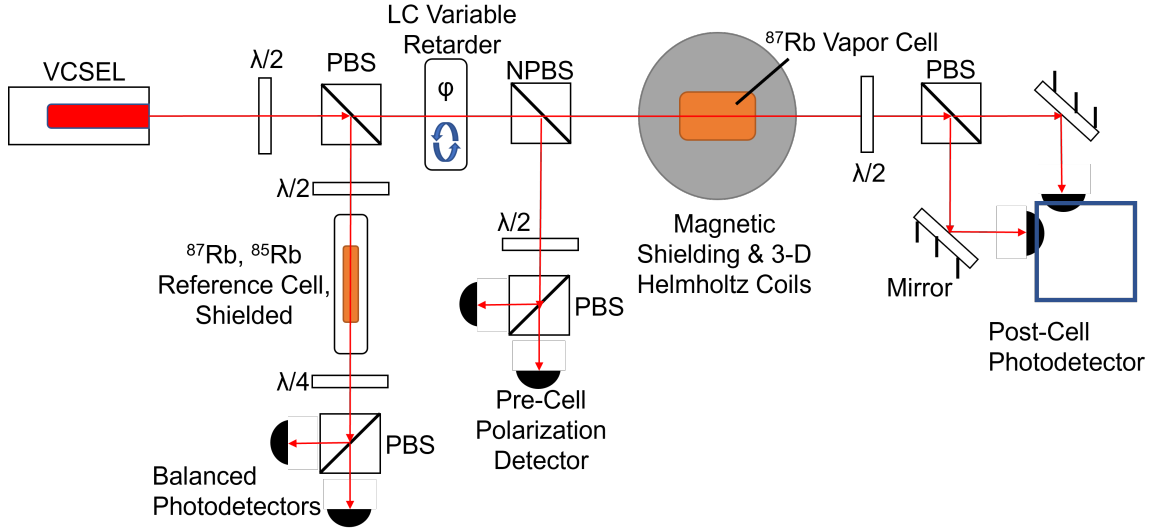
Figure 3.1: Optics Setup. VCSEL stands for vertical-cavity surface-emitting laser, $\lambda/2$ and $\lambda/4$ are half- and quarter-wave plates respectively, PBS stands for polarizing beam splitter, NPBS is a non-polarizing beam splitter, and the LC variable retarder (LCVR) is the device capable of effective polarization rotation.

3.1 Optical Components

3.1.1 Laser

The vertical-cavity surface-emitting laser (VCSEL), henceforth referred to as the laser, is a low power consumer. Its current can easily be modulated using RF signals. The laser emits light of wavelength 794.7 nm, in the near-infrared range [8]. Two factors influence the laser's output frequency: temperature and input current. We typically only change the current, so the temperature is held constant near 50°C using a standard temperature controller. For most measurements we lock the precise laser frequency to a ^{87}Rb atomic line, as described in Section 3.1.2. The laser has a maximum current of 2 mA, which means the current source has to have small output and add very little noise to the current. The light is not visible to the naked eye, but special viewing cards or optics are used to view the laser at any given point for alignment purposes. Fig. 3.2 is photo of the device we use.
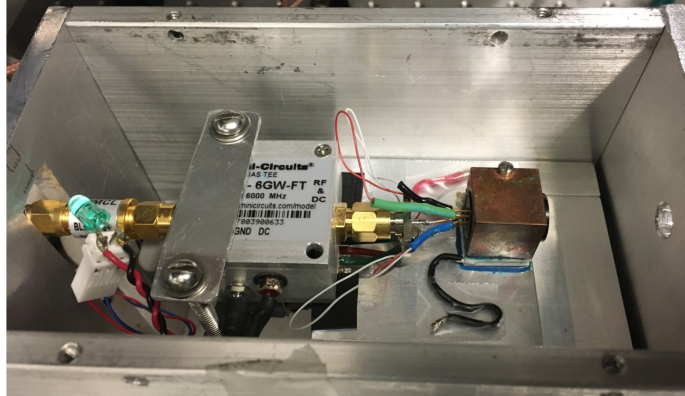
Figure 3.2: The VCSEL (laser) used in the setup. When in use, a cover is placed over the box and the beam exits through the right hole.

EIT requires two electromagnetic fields of differing frequency to achieve two-photon resonances, and EIT is sensitive to the relative phase between the two. There would be no coherence between two independent laser systems combined into one beam, as they would 'jump' around the set frequency in a random way and out of synchronization. To avoid this we modulate the current of the VCSEL at the desired frequency separation, effectively creating multiple frequency components within the same laser field. As a result, all random jumping applies to both frequency outputs coherently. An external generator sets the relative frequency between the two fields such that a carrier and sideband comb form at a one-to-one intensity ratio 6.834 GHz apart (the hyperfine splitting of ^{87}Rb) [9]. Credit goes to Nathan T. Belcher for the design and construction of this apparatus [8].

3.1.2 DAVLL

The dichronic atomic vapor laser lock, or DAVLL, is a system designed to lock the laser onto a specific rubidium resonance. It applies the raw differential signal, or error signal, from the optics later in this section to adjust the laser current so that it never strays from that which generates the correct frequency. Fig. 3.3 is a photo of the

beam splitter and photodetectors that provide a basis for the error signal.
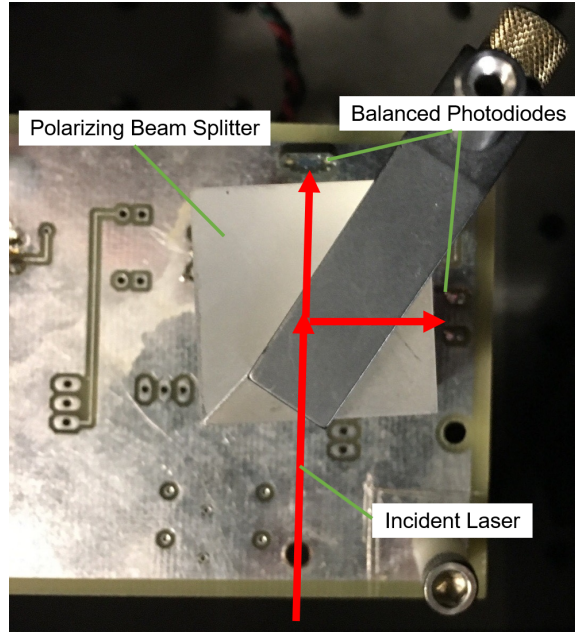
Figure 3.3: Balanced photodetectors on different sides of a beam splitter; this is enclosed in a black box when in use, as it is very sensitive to environmental light.

Observe Fig. 3.4, an example signal from the DAVLL near a resonance. Each of the Zeeman-shifted sublevel resonances is detected individually by one of the balanced photodetectors in Fig. 3.3, thus when the detector signals are subtracted, a curve like this occurs with two visible peaks (one negative due to subtraction). The quasi-linear portion in between that crosses zero is the exact location of the resonance we want to lock to, seeing as these visible levels were shifted exactly the same amount in opposite directions (Section 2.3 has more detail on how this occurs). After a human increases the laser's current and flips a switch on the DAVLL to lock at this zero-crossing frequency, the DAVLL will detect with a feedback loop when the error signal strays above or below zero; it will thus increase or decrease the current to the laser accordingly, fixing its frequency.

The entire lower arm of Fig. 3.1 is part of the DAVLL system and consists of an
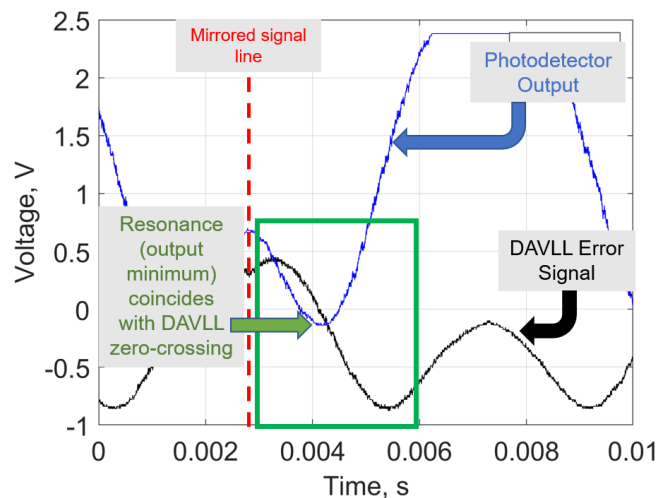
Figure 3.4: An example differentiated signal from the DAVLL. Each photodetector detects transmitted light from one of the separated resonances, so when the output signals are subtracted, we end up with a curve as such. The 'linear' zero-crossing is where we lock the laser and how the DAVLL detects if we have strayed off resonance. The zero-crossing lines up with the dip in the photodetector output signal where the resonance occurs.

additional rubidium cell. This cell contains a natural abundance mixture of ^{87}Rb and ^{85}Rb atoms (^{85}Rb is about three times more abundant in nature) and no buffer gas, as no EIT needs to occur here. The cell is encased in cylindrical shielding and subjected to a localized, strong (50 Gauss), homogeneous magnetic field. The incident light is linearly polarized and directed before the cell. The strong magnetic field causes a large Zeeman shift in two absorption lines for circularly-polarized components of the laser field that increase and decrease in energy respectively [8]. These components are uniquely reflected or transmitted through the subsequent beam splitter and the strength of each signal is measured by the photodetectors. Fig. 3.5 is a photo of this arm of the apparatus.
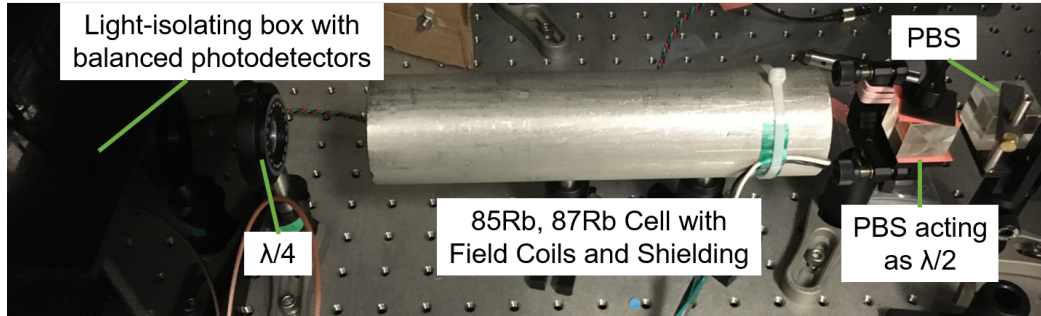
Figure 3.5: Optics used in the diverging DAVLL arm of the setup.

3.1.3 Liquid Crystal Variable Retarder

The liquid crystal variable retarder (LCVR) is a key device of interest because its function is to rotate the plane of polarization of incoming light. In terms of the geometry of Fig. 2.4, this device changes the direction of $\vec{E}$ and thus necessarily ϕ, a key variable in determining magnetic field direction. The device performs the same function as a half-wave plate ($\lambda/2$), however it changes the angle at which it polarizes incoming light dependent on the voltage applied to it. This presents the opportunity to change ϕ to specific angles, sweep ϕ over its whole or partial range of possible values, and vary ϕ with a desired frequency. Fig. 3.6 is the device we use.
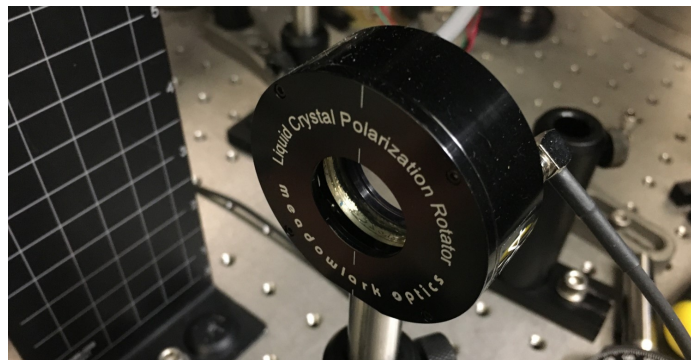
Figure 3.6: The LCVR used in the setup.

Figure 3.7 shows a molecular cross-section of the LCVR. Applying no voltage to the device means all anisotropic nematic liquid crystal molecules are parallel to

17
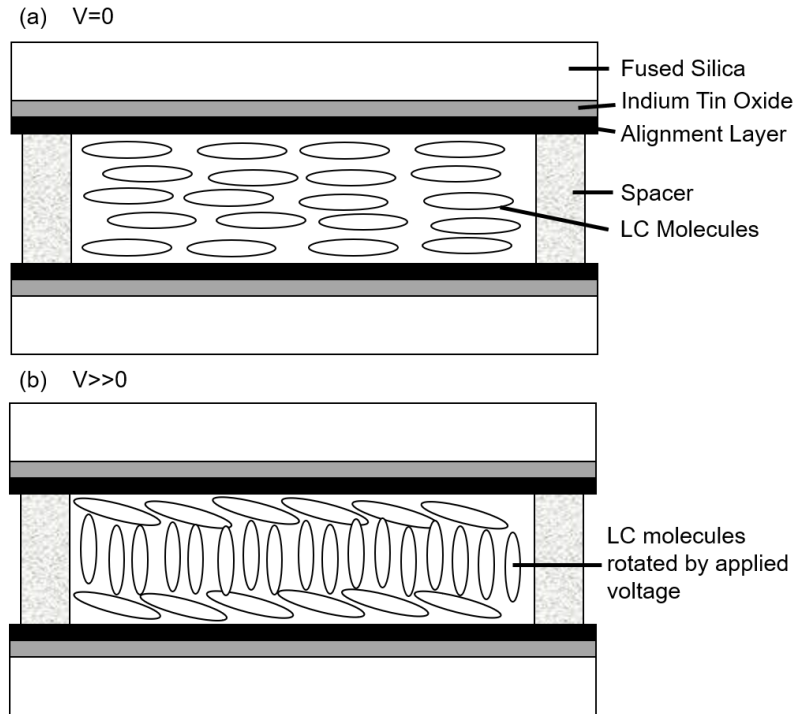
Figure 3.7: Schematic of LCVR operation. In (a), no voltage is applied and all molecules in the liquid crystal layer are relaxed and oriented parallel to the outer fused silica and alignment layers. In (b), a voltage $V >> 0$ is applied and all molecules not pinned to the alignment layer rotate towards a perpendicular orientation, where minimum retardance is achieved.

the fused silica and alignment layers. Maximum retardance is achieved in this state. An applied voltage will forcefully rotate and hold molecules in their rotated state so long as that constant voltage is applied. In the 0-10 V range the device is capable of taking in, higher voltages move the atoms towards being perpendicular to the alignment layer. The layer of molecules closest to each alignment layer are pinned, however, and are unable to rotate fully. Maximum and minimum retardance do not mean anything as metrics in this experiment, so the device will be calibrated in Section 4.1, finding a function mapping voltage to angle and vice versa.

While more detailed results from calibrating and investigating the performance capabilities of the LCVR reside in Section 4.1 and Section 4.2, the primary result was

that the conversion between input voltage and resulting angle is

$$\phi_{LCVR} = \frac{a}{1 + (V_{in}/b)^c} + d, \tag{3.1}$$

where $a = -239.1 \pm 0.4$, $b = 1.868 \pm 0.001$, $c = 4.12 \pm 0.01$, and $d = 249.3 \pm 0.1$ are fit parameters, ϕ_{LCVR} is the rotation angle relative to the x-axis in Fig. 2.4, and V_{in} is the input voltage using computer input (as opposed to function generator input, which needs to be halved for this device). We also resolved to modulate the input voltage at 18 Hz with 30 mV peak-to-peak amplitude, as we require a modulating signal for the lock-in amplifier described in Section 3.2.1.

3.1.4 Rubidium Cell

The last uncommon component in the optical setup is the primary rubidium vapor cell, the top-rightmost cell in Fig. 3.1. A shielded cylindrical container holds the cell to avoid interference from non-homogeneous residual fields in the environment. Within this cylinder, three sets of Helmholtz coils (aligned longitudinal to the laser z, transverse horizontal y, and transverse vertical x) form a three-dimensional apparatus to create a magnetic field in any direction by adding the axis-aligned magnetic vectors each set of coils produces. Fig. 3.8 shows the cell enclosure and final photodetector, with a simple lens in between to focus incoming light.

The cell itself holds only the ^{87}Rb isotope of rubidium and neon buffer gas. Rubidium atoms in the dark state in an EIT resonance tend to lose this superposition state when they collide with the walls of the cell. The atoms have high velocity and can move in any direction, thus many collisions happen per unit time; we do not want our specially-prepared atoms to decohere rapidly after achieving EIT. The addition of buffer gas is an established trick to alleviate the issue. Rubidium atoms will collide with neon atoms more often and the walls much less, but neon collisions preserve the spin state of rubidium atoms [10].
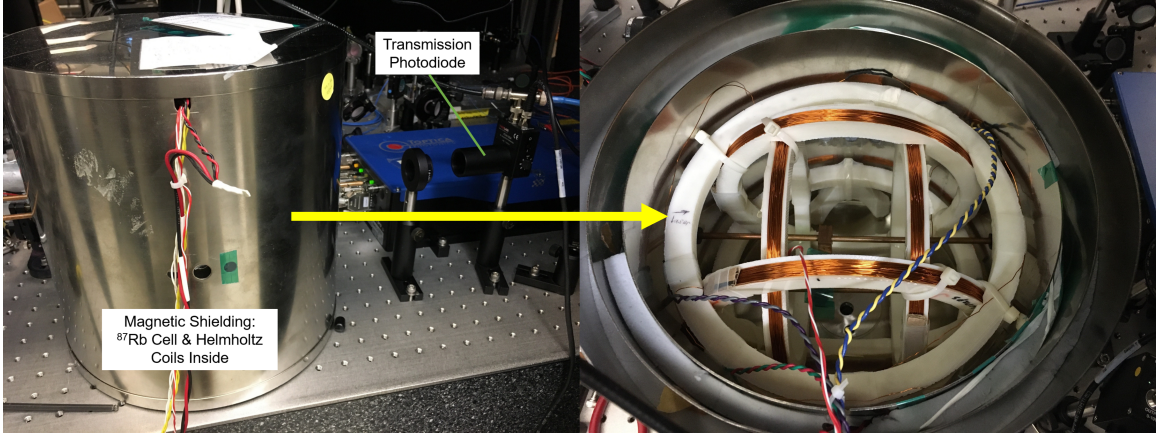
Figure 3.8: The shielded ^{87}Rb cell chamber and simple photodetector at the end of the setup. The right image shows the 3D-printed three-dimensional Helmholtz coil setup within.

A part of the collaboration's exploration is to determine the optimal cell parameters. The pressure of buffer gas must be calibrated because too little allows more Rb wall collisions, but too much dulls the effects of EIT transmitted through the cell. Cell temperature is also important and under investigation, because higher temperatures lead to a higher pressure of vaporized ^{87}Rb atoms but less transmitted light. We used 56.0°C.

3.2 Electronics

Several electronic systems are necessary for control and measurement of this system. Depending on the task, electronics may be connected in a variety of ways. This section will give an overview of each system and how it is typically used in this study. Chapter 4 will note any vital changed connections during experimentation.

3.2.1 Lock-in Amplifier

A lock-in amplifier is a device that can ensure a signal lock, like on a polarization setting maximizing transimission, by means of finding the first derivative of an input
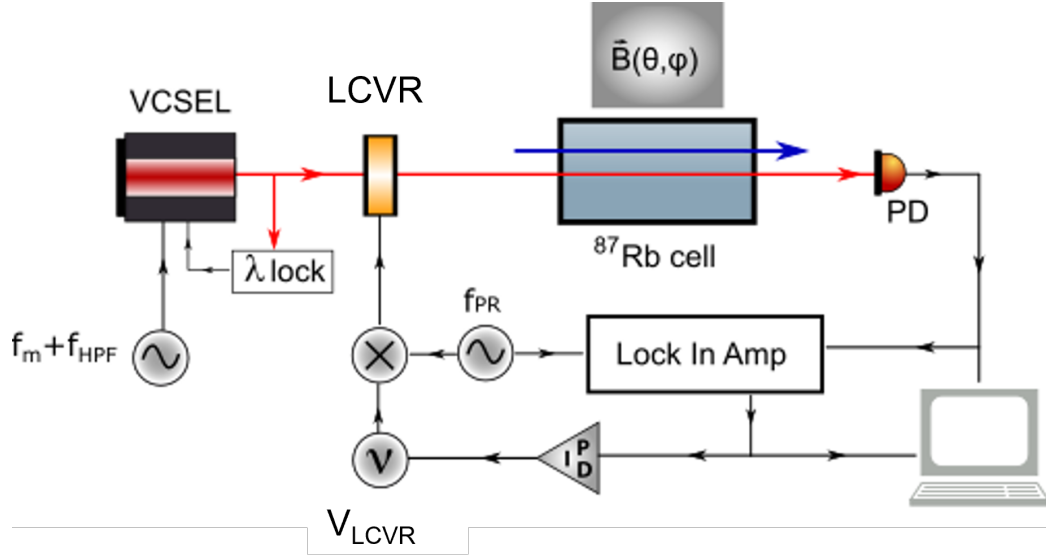
Figure 3.9: Key electronics and connections in the magnetometer. The λ lock is the DAVLL, V_{LCVR} is the DC voltage supplied to the LCVR, f_{PR} is the added modulation frequency added to the LCVR input, and PID stands for proportional-integral-derivative controller.

signal. This device can resolve small signals within a great deal of noise so long as it is given the carrier wave of the target signal. It takes in the final photodetector's output signal, a reference signal carrier wave of the same frequency as the output we want, and it outputs what is desired: either an amplified target signal without the noise or a higher derivative of this signal [11].

In our setup, we want the capability to hold the LCVR polarization at maximum transmission as seen on the curves in Fig. 2.5. This requires effective dynamic locking of the LCVR input voltage, meaning we need a feedback loop. The derivative of the output signal is a fantastic basis for a feedback loop, since the derivative at any rounded local maximum or minimum is zero. In a sweep around the minimum or maximum, the derivative will look much like the DAVLL signal in Fig. 3.4. Thus, ensuring the lock-in amplifier's signal is zero ensure we are situated at a transmission peak: this can be used as an error signal for the proportional-integral-derivative (PID)

controller, locking the LCVR at a polarization with peak transmission.

Another application of the lock-in amplifier (Section 4.3) is to determine the frequency of an EIT peak. We can use it to look for the frequency of the zero-crossing of the derivative of laser transmission when we sweep and modulate the RF frequency rather than the LCVR polarization, indicating the peak's maximum resides at the same frequency.

3.2.2 Photodetectors
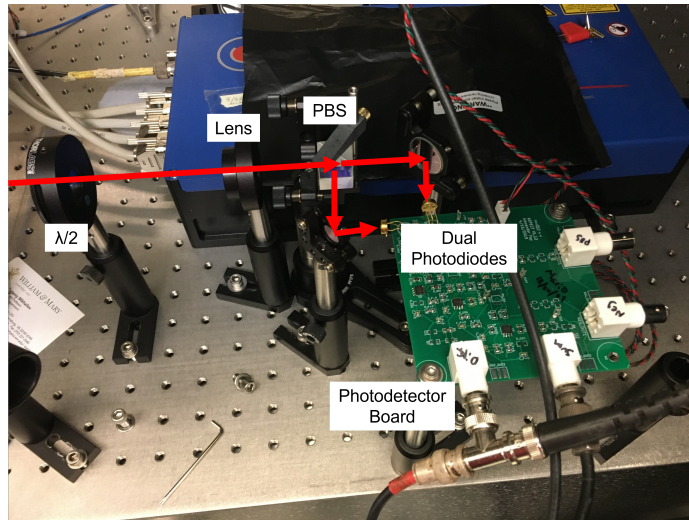
Figure 3.10: The differential photodiode board. The incident laser is split by a polarizing beam splitter and each resultant beam is guided to separate photodiodes.

The level of transmitted laser light passing through the cell is recorded as the voltage of the signal output by the post-cell photodetector receiving the light. For many investigations, a simple photodiode measuring the total transmission is enough, pictured in Fig. 3.8. However, if we desire a measurement of how the laser's polarization has changed, a more complicated two-channel "differential" photodiode board is required as pictured in Fig. 3.10. This device takes two light inputs on separate photodiodes and can output the sum, the difference, and the individual outputs. We

place a polarizing beam splitter (PBS) before this device to change the transmission in each channel if the incoming light's polarization shifts.

This device's sum output acts the same as a simple photodiode: regardless of how the PBS splits the light, the sum of the two resulting beams will be the same as what hit the PBS initially. The difference signal sets the device apart because we can start by balancing the input to each photodiode with a half-wave plate (setting the difference to zero) and measure how much the difference changes when we shift a parameter. In this device, the notable use is to investigate how the ^{87}Rb vapor rotates laser polarization.

Another pre-cell photodetector is in place to independently measure the polarization of the LCVR. While we calibrate the LCVR in Section 4.1 to know what input voltage translates to what polarization, we will later add an error signal from the PID controller to the LCVR input, thus we cannot know the current polarization without measuring it independently. This detector takes some of the light intensity through a non-polarizing beam splitter and acts like the difference signal on the post-cell detector while avoiding all atomic effects in the cell. The equation we use to calculate polarization from the difference signal is

$$\phi = \frac{1}{2}\arcsin\left(\frac{\Delta I}{\Sigma I}\right), \tag{3.2}$$

where ϕ is the angle from the $\vec{B}$-$\vec{k}$ plane in radians, ΔI is the difference signal, and ΣI is the sum of the two photodetector intensities.

3.2.3 Auxiliary Electronics and Software

Several additional electronics and computer software are vital for the magnetometer's operation and experimentation that each require a brief mention.

A four-channel oscilloscope is essential for locking the laser, system monitoring, and for most data acquisition. It is connected to a desktop computer such that a

MATLAB program can extract any waveform data displayed on the scope screen. Its measurement abilities are common knowledge, but on the topic of locking the laser, the waveform in Fig. 3.4 is directly viewed on an oscilloscope triggered at the laser's modulation frequency. At this trigger frequency we also see the output of the final photodetector in line with the DAVLL signal such that the characteristic dips in transmission that occur in the neighborhood of an EIT peak line up with the linear zero-crossing of the DAVLL signal. This helps find the location to lock the laser frequency at before sweeping frequency around the resonance to find EIT peaks.

RF sweeps are essential in finding the exact small frequency range where CPT occurs. After laser lock, a Python program on the same desktop computer can perform RF sweeps around the central resonance in frequency ranges from 1 Hz to 10 MHz. This is done by precisely modifying the laser input current.

Two standard temperature controllers are used. The first is used to achieve rubidium resonance frequency in the laser; a combination of temperature and current supplied determines the output light frequency. The second sets the temperature of the main rubidium cell, where the temperature determines the cell's vapor pressure. Again, the latter controller holds the cell temperature at 56.0°C.

The LCVR that we use pairs with a digital interface controller that ensures that a clean and consistent signal is sent to the polarizer. Software pairs with the controller to directly control the input to the LCVR, however in a manner not easily customizable. Instead, we use external input, through this controller, from a function generator. Often we want to use it to set a constant DC offset, sweep through a range of voltages, or to perform high-frequency sinusoidal fluctuations. Two signals of these types can be added together, creating a slow sweeping, a fast modulating, or constant offset signal sent to the LCVR. This signal may be summed with the error signal from the PID controller to lock the LCVR at a polarization with peak transmission.

Chapter 4

Experimental Methods and Findings

The research accomplished falls under several short-term projects to understand or improve parts of the magnetometer. Each project's goals, methods, and findings are explained in separate sections of this chapter. To motivate the need for each section, here is the proposed operational procedure for measuring the magnitude and direction of a magnetic field using this magnetometer.

Magnitude

1. Sweep the laser frequency over a range encompassing the $a_{\pm 2}$ EIT peaks, about 2.5 MHz in Earth's field of about 0.5 Gauss.

2. Measure the difference in frequency between the peaks, then use equation $B = \Delta\nu/4\gamma$ to calculate the field strength ($n = 4$ since the $a_{\pm 2}$ peaks are 4 peaks apart).

Azimuthal Angle ϕ

1. Move the laser frequency to one of the $a_{\pm 2}$ EIT peaks.

2. Dither the LCVR polarization to find where the lock-in amplifier crosses zero (from positive to negative), indicating a point of maximum transmission.

3. Activate the PID controller which will provide an error signal based on the lock-in amplifier and lock the LCVR's polarization to the $\vec{B}$ - $\vec{k}$ plane.

4. Measure the polarization independently with a pre-cell detector to determine the angle ϕ.

Polar Angle θ

1. Move the laser frequency to one of the $a_{\pm2}$ EIT peaks.

2. Again lock the LCVR to maximum transmission, in the $\vec{B}$ - $\vec{k}$ plane. Measure the peak amplitude.

3. Next lock the LCVR to minimum transmission to measure the background level. Take a ratio of the peak amplitude and background to determine the angle θ.

4.1 LCVR Calibration

The LCVR is a new addition to this ongoing project, so the first task is to find a calibration equation that maps input voltage to angle and vice versa. There are many caveats to this task that make it far more in-depth than it first appears.

The first attempt at a calibration uses the LCVR controller software to control its input voltage. The laser's temperature controller is switched off so that no rubidium resonances interfere with raw transmission data. A polarizing beam splitter is placed after the LCVR but before the ^{87}Rb cell to observe the change in the LCVR's effective polarization angle. If the LCVR polarizes along the PBS's polarization axis the full laser intensity will be transmitted, but if it is polarized perpendicularly, no intensity should transmit. The transmitted intensity follows the equation

$$I = I_0 \cos^2 \phi, \tag{4.1}$$

where I is the transmitted intensity, I_0 is the incident intensity, and ϕ is the relative angle between the polarization axes of the incident light and PBS. Solving Eq. 4.1 for ϕ and 'unwrapping' the results, Figure 4.1 shows the first calibration result. Through a PBS, the intensity measured at supplementary angles (like 80° and 100°) will be the same, so the data-unwrapping process correctly maps each intensity to the correct angle of LCVR rotation. This produces a curve we can fit with an equation. Similar data that has not been unwrapped looks like Fig 4.3.
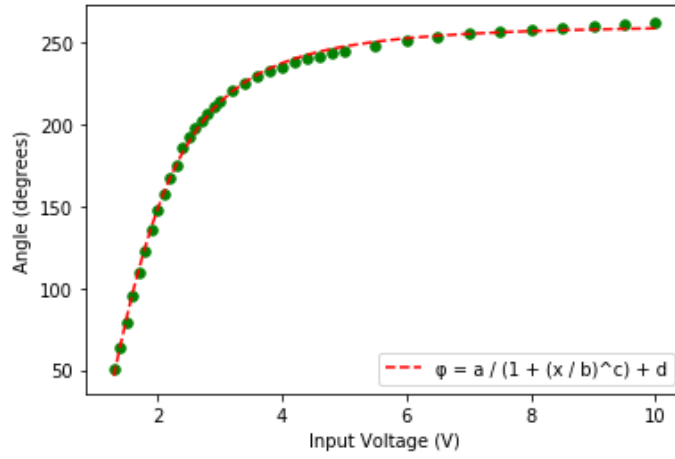
Figure 4.1: Input voltage to LCVR angle calibration by stepping to constant voltages and recording photodiode output. In the fit function, $a = -360 \pm 20$, $b = 1.50 \pm 0.06$, $c = 2.72 \pm 0.09$, and $d = 261 \pm 1$. These are statistical uncertainties.

This result is very important: it tells us that voltage does not linearly map to change in angle for the full 0-10 V range. It instead has a best fit function of the form

$$\phi_{LCVR} = \frac{a}{1 + (V_{in}/b)^c} + d, \tag{4.2}$$

which also leads to a function mapping angle to voltage

$$V_{in} = b \left(\frac{a}{\phi_{LCVR} - d} - 1 \right)^{1/c}, \tag{4.3}$$

where a, b, c, and d are fit parameters found every time a calibration is run and V_{in} is the input voltage. Unfortunately, this function becomes less accurate as we

reach higher voltages. Fortunately, a full range of 180° occurs before the function loses accuracy: this means we can instead survey a smaller range of voltages to find a better fit function. This choice is made more appealing by the fact that the function's rate of change at higher voltages is very small; the range of angles covered past 4 V is very small and thus not very useful for this project.

Another important note to make is that the numerical angles listed on the y-axis are calculated relative to the polarization axis of the polarizing beam splitter temporarily inserted between the LCVR and rubidium cell. This arbitrary choice of axis is the x direction in Fig. 2.4, perpendicular to the table.

We now want to observe the LCVR's calibration when the voltage is swept rather than individually set to different angles to observe the LCVR's performance. We used the LCVR controller software to set triangular wave voltage sweeps and recorded the output on an oscilloscope. Since the molecules in the LCVR have a non-trivial response time to changes in input voltage, the sweep was varied with three parameters: sweep direction (low to high versus high to low voltage), sweep range (about 1 V about a 90° range versus 10 V), and sweep time (1, 2, 4, 8, and 16 seconds). The results are in Fig. 4.2.

In Fig. 4.2(a), we have an increasing voltage sweep from 0-10 V. As expected, each of the five sweep times follows the same shape as our initial calibration in Fig. 4.1; they appear visually different at first because the data was not unwrapped (for example, the second 60° in Fig. 4.2 is equivalent to 120° in Fig. 4.1). A clear inconsistency appears in the low voltages: short-time sweeps have a hard time catching up to the angle they should be at when the voltage is swept so fast. The 16 second sweep (purple) encapsulates what the sweep should look like best, starting flat and unchanging before 1 V and then changing angle rapidly immediately after. The 1 second sweep (blue) attempts to adjust from 10 V at the end of the last sweep to 0 V, but clearly not fast
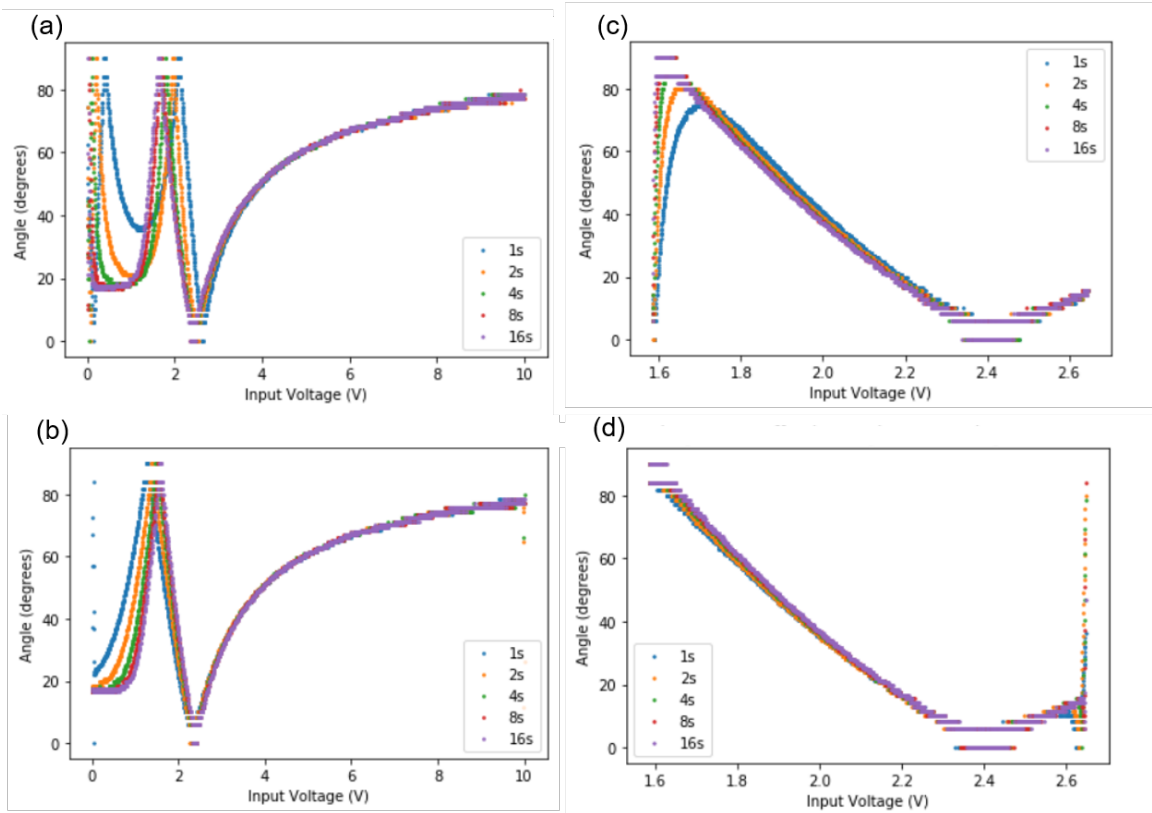
28

Figure 4.2: Input voltage to LCVR angle through twenty different sweep setting combinations. No 'unwrapping' is done, due to this analysis being qualitative. Sweep time colors are as follows: blue = 1 s, orange = 2 s, green = 4 s, red = 8 s, purple = 16 s. (a) has increasing voltage from 0-10 V, (b) has decreasing voltage from 0-10 V, (c) has increasing voltage from 1.584-2.646 V, and (d) has decreasing voltage from 1.584-2.646 V. The second range is approximately 90° of rotation.

enough to even reach the 20° starting point it should be at from 0→1 V. This means it and the other faster sweeps continue to lag behind even on the down-slope starting at 2V. Once the polarization angle starts changing at a slower pace, however, all sweep times have their angles synchronized to where they should be for the remainder of the higher voltages.

In Fig. 4.2(b), the same 0-10 V range is swept, but this time starting at the high 10 V and sweeping down to 0 V. Again, there is inconsistency between the fast and slow sweeps, but this inconsistency is visually smaller. This tells us relaxing the voltage

causes a faster reaction from the LC molecules than increasing it.

Fig. 4.2(c) sees an increasing voltage sweep again but with a smaller voltage range encapsulating approximately 90° of rotation. This smaller range ensures the sweeps are changing voltage at a slower rate, giving more time for LC molecule reaction. There is certainly more coherence between sweep times in this case versus the same downward slope around 2 V in Fig. 4.2(a). However, the faster sweeps of 1 and 2 s still lag behind when the sweep starts.

Finally in Fig. 4.2(d), returning to the better downward sweep but in this smaller range, all sweep times result an almost immediate shift to the correct angle when the sweep starts on the right. The 'trace' points on the right (where the LCVR is rapidly attempting to reach the correct angle) are much more brief than those in Fig. 4.2(c), and the visual lag behind the 16 s sweep is much smaller. This reinforces the need to sweep from high to low voltage going forward, in a small range if possible, and for longer times.

With this information, a more definitive calibration curve is possible. We now know a decreasing voltage sweep is ideal, slower sweep speeds are ideal (though long sweeps are unrealistic for the final device), and a narrow voltage range gives sweeps more accuracy. This next sweep is 4 seconds long, as it is a realistic time scale while being long enough to achieve a high degree of volt-to-angle accuracy. It takes the range of 3 V between 1 V and 4 V, as this encapsulates more than 180° of rotation and avoids less-useful higher voltages.

Fig. 4.3 is the result of the calibration sweep. The points follow a curve that the same function type as in Fig. 4.1 follow. The primary difference, however, is that the lack of deviating points at large V allows fit parameters which fit the data much better. This is numerically evident from the order-of-magnitude smaller errors on all the fit parameters between Fig. 4.3 and Fig. 4.1.
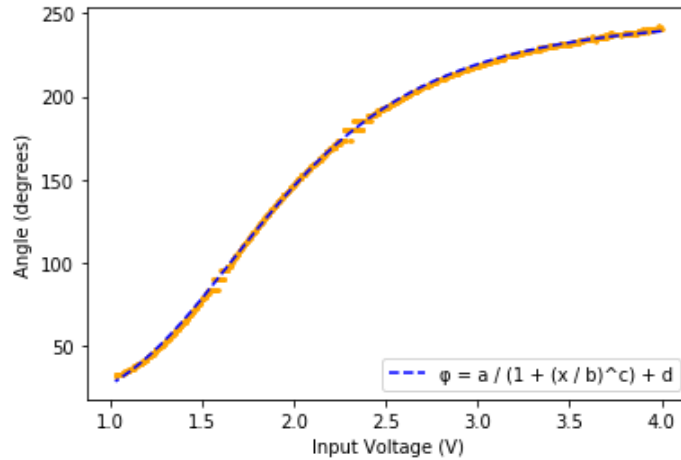
Figure 4.3: Input voltage to LCVR angle achieved through a decreasing voltage sweep from 4 V to 1 V for four seconds. In the fit function, $a = -239.1 \pm 0.4$, $b = 1.868 \pm 0.001$, $c = 4.12 \pm 0.01$, and $d = 249.3 \pm 0.1$. The decreased point resolution around 90° and 180° is due to the nature of the arc-cosine conversion from intensity voltage to angle.

This sweeping calibration method is much faster than taking individual points like in Fig. 4.1, so every time a new voltage source or form of signal modulation was introduced, we ran a calibration sweep to check the same function fit.

4.2   Applying Modulation

The lock-in amplifier requires a signal with amplitude modulation as input, so the next order of business is to test how the LCVR responds to an input signal with fast modulating amplitude. First, we found in Section 4.1 that the LCVR has a nontrivial response time to a new applied voltage, so it is reasonable to guess that modulating voltage a large amount with high frequency will result in the device lagging. It will be unable to achieve the full range of rotation for each period of the fast modulating signal.

To test this, we choose four DC offsets for the fast modulating signal and calculate the peak-to-peak modulation amplitude that would cause about ±5° of rotation about

each DC offset. We place a half-wave plate and PBS after the LCVR and turn off the laser's temperature control to ensure no EIT interaction. The half-wave plate is rotated to the point where about half of the maximum intensity output is achieved at each DC offset; this is to ensure the modulation of angle achieves the maximum intensity difference when rotated in a 10° range. We disconnect the LCVR's controller from the controller software and start using a more versatile function generator for the LCVR. We record the peak-to-peak amplitude of the photodetector's output signal when this modulation signal's frequency is set to several frequencies.
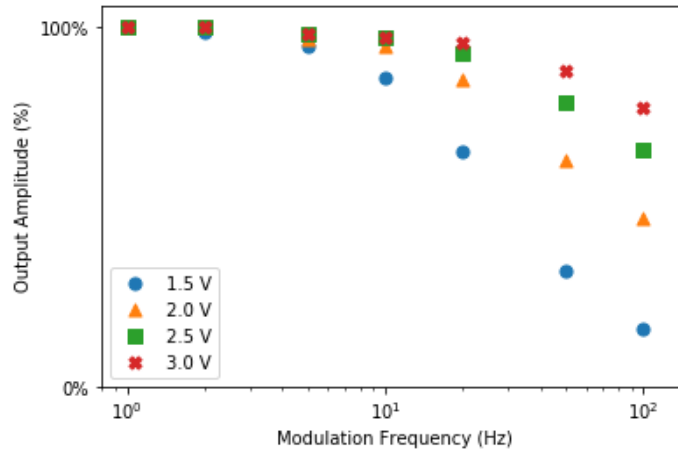
Figure 4.4: Effect of increasing modulation frequency on the peak-to-peak amplitude of the photodetector output signal for various DC offset voltages.

In Fig. 4.4, we observe what was expected: increasing the frequency of the modulation input decreases the effective rotation range achieved. The more interesting result here is that a higher DC offset for modulation means the overall loss of amplitude is decreased. When at a DC offset of 1.5 V, frequencies not far above 100 Hz cause too significant a loss for the modulation to be useful, falling lower in amplitude than noise. However, working with larger DC offsets allows higher usable frequencies because the loss isn't as destructive.

An additional reason to use smaller frequencies is that the output signal is dis-

torted and no longer sinusoidal above about 50 Hz. This is traced back to the response time: increasing the voltage to the LCVR shifts its molecules into position quicker than decreasing the voltage, and thus modulation shoots transmission fast in one direction and slow in the other. As a result, we decide to use a modulation frequency of 18 Hz. The lock-in amplifier can pick it up, it is sufficiently small to produce an easily detectable peak-to-peak amplitude, and no distortion occurs.

4.3 EIT Peak Frequencies

For many experiments in this work the magnetic field is held at constant magnitude and in transverse orientation ($\theta = 90°$) for maximum output amplitude on the $a_{\pm 2}$ peaks. This corresponds to inducing a 500 mG field through the set of Helmholtz coils parallel to the ground and transverse to the laser field. The frequency separation of EIT peaks is constant under this condition (the $a_{\pm 2}$ peaks are 700 kHz from the central peak), as determined by the theory in Section 2.3, thus we can predict the exact frequency of each peak.

Any half-wave plate or PBS between the LCVR and the cell is removed, the laser temperature control is switched on, and the laser is locked at the $F = 1, 2 \rightarrow F' = 2$ transition. The RF sweeping Python program is used to change the modulation frequency with ease and precision.

The lock-in amplifier finds the derivative of the photodiode output, so determining where the lock-in output crosses zero tells us where the peak position is. Unfortunately, experimentation with the lock-in output has the zero-crossing shifted a few kHz from the photodiode output's EIT peak, as seen in Fig. 4.5, due to the asymmetry of the EIT peak.

However, switching the lock-in's setting to measuring the second harmonic of the signal manages to fix this issue; this setting is less sensitive to the asymmetry. The
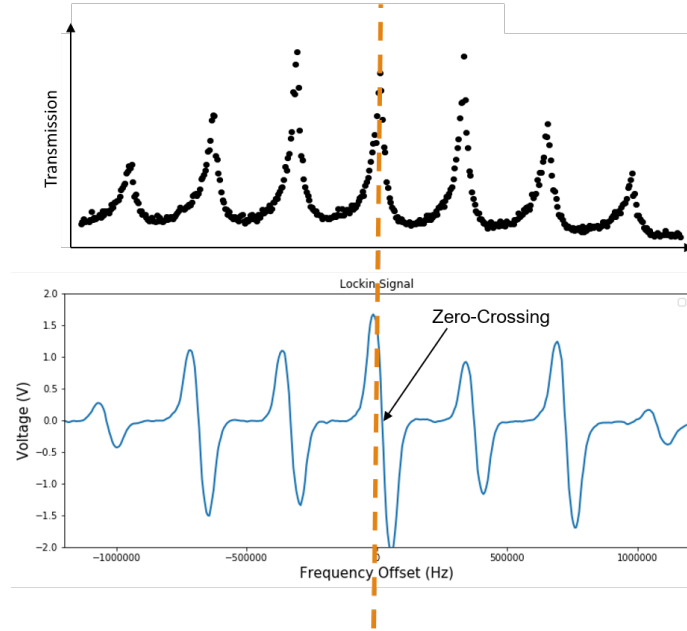
Figure 4.5: The lock-in amplifier's first derivative signal zero-crossing does not align with the EIT transmission peak during testing.

second harmonic instead finds the second derivative of the input, meaning the EIT peak will not align with a zero-crossing, but instead with a minimum or maximum due to the zero-crossing having a large slope. By first sweeping RF frequencies in a large enough range (±1.25 MHz) to see all seven EIT peaks and then "zooming in" by sweeping a small (±5 kHz) range around a peak position, lock-in output data gives a large positive peak with a maximum at the EIT resonance (Fig. 4.6).

Resonance	Frequency ($\pm2 \times 10^{-6}$ GHz)	Δ from Central Resonance (MHz)
-3	6.8336367	-1.0445
-2	6.8339844	-0.6968
-1	6.8343325	-0.3487
0	6.8346812	0
+1	6.8350323	0.3511
+2	6.8353805	0.6993
+3	6.8357316	1.0504

Table 4.1: EIT resonance frequencies for a 500 mG transverse field. The difference from the central resonance is listed in the third column.
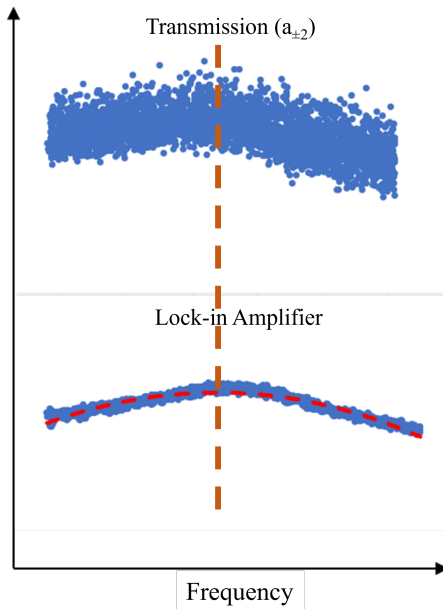
Figure 4.6: The lock-in amplifier's second harmonic signal peak lines up much closer to the EIT peak. This is in a much smaller frequency range than Fig. 4.5, thus the offset is much smaller than what is produced with the first harmonic.

Table 4.1 displays the frequencies of the EIT resonances at this field strength. Nontrivial noise existed in the lock-in second harmonic signal, so the exact location of the peak was found through curve fitting and warrants significant uncertainty. One tip that this data is imperfect (though a useful guide nonetheless) is that the magnitude of the difference from the central resonance for the peaks $\pm 1, \pm 2$, and ± 3 do not each match. It is arguable that parallel magnitudes here fall into each others' uncertainty ranges, but there are certainly methods to decrease the noise in the lock-in signal and achieve more accurate frequency measurements. Accuracy here leads to more accurate measurements of the scalar magnetic field.

4.4 Magneto-Optical Rotation caused by Rubidium

We discovered evidence of nonlinear magneto-optical polarization rotation (NMOR) in our system, gradually rotating the laser's polarization through the rubidium cell. We observe two forms of NMOR, single- and two-photon effects. Single-photon NMOR is related to the larger hyperfine energy level and its Zeeman sublevels: there is a difference in the index of refraction for left- and right-hand circularly polarized light between the sublevels. Two-photon NMOR is related to EIT, where a small difference in the g-factor for the two greater hyperfine levels breaks the symmetry between "symmetric" EIT peaks. Since we want to achieve a polarization locking mechanism so as to hold our LCVR at the polarization with maximal transmission, we need to characterize this effect so it can be accounted for in future measurements.

This is where the difference signal on the four-channel photodiode is useful: we can measure how the signal moves above or below its off-EIT-resonance zero point, indicating the polarization has rotated. The PBS splitting the post-cell signal into its horizontal and vertical polarization vector components enables the difference measurement.

First, we observe this polarization shift effect on the a_{+2} resonance over a 180° range of field rotation in θ, from parallel to the laser wave-vector to anti-parallel. We achieve this range by flipping the longitudinal Helmholtz coil input wires, creating a field in the opposite direction while still manipulating the transverse coils to reach any angle desired.

Figure 4.7 shows us that fields closer to parallel to $\vec{k}$ and fields closer to antiparallel to $\vec{k}$ rotate the laser's polarization in opposite directions. This gives a reliable method to differentiate between supplementary θ angles, like 50° and 130°, as the EIT peaks
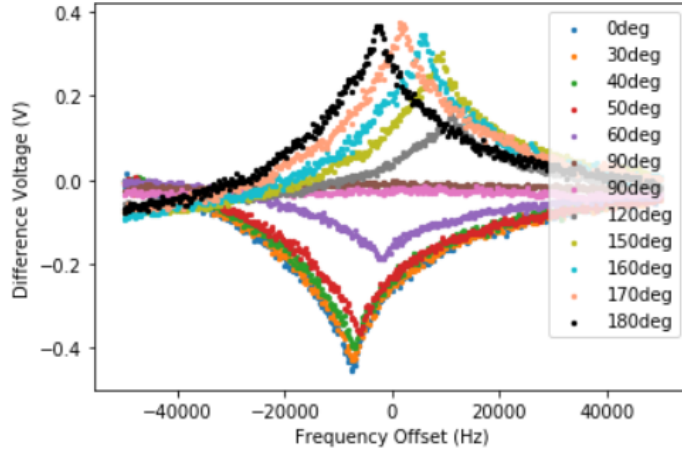
Figure 4.7: The differential photodiode's measurements of the rubidium cell's two-photon NMOR at different field θ angles for the a_{+2} resonance. There are two traces at 90° to show no difference between the angle when dialed with flipped coils. The shift in peak position is due to imprecise field coil alignment.

normally look the same at these angles. Two-photon NMOR therefore fulfils one of the greater magnetometer project's goals of breaking the z-axis symmetry. The x- and y-axis symmetries have yet to be broken, however.

We attempted the same measurements by flipping the transverse field coils in a 180° range. Unfortunately, there was no reliable difference between supplementary angles in this range and thus there is no useful method to find here. We expected this result, however, as we can imagine the polarization shift "peak" in Figure 4.7 oscillating in amplitude like a cosine wave: the wave crosses zero amplitude at 90° and hits maxima and minima at 0° and 180°, so we cannot tell if we are at $-10°$ or 10° by looking at amplitude alone.

Now, we would like to observe if this effect occurs for all the EIT peaks, taking full seven-resonance polarization difference traces at many field angles.

Figure 4.8(b) shows each of these traces stacked on top each other, from $\theta = 0°$ to $\theta = 90°$, clearly showing that at each of the EIT resonant frequencies, a polarization shift occurs. The magnitude of this shift is dependent on the peak and field angle.
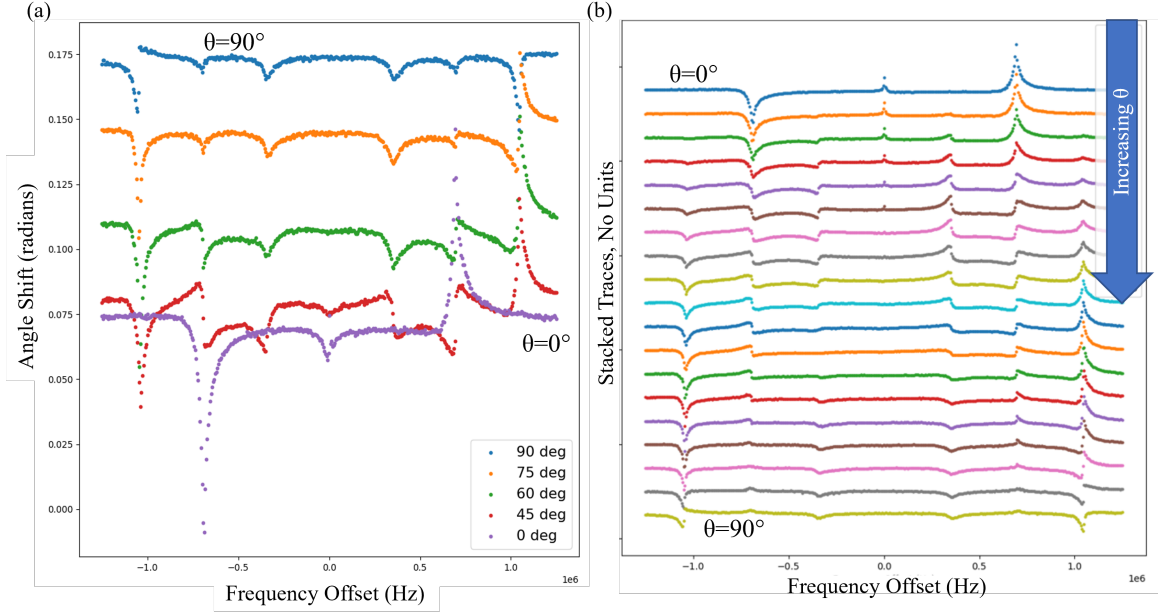
37

Figure 4.8: The differential photodiode's measurements of the rubidium cell's rotation of polarization at different field θ angles over the EIT spectrum range. Each trace has seven deviations from its unique "zero" level, each where an EIT peak normally occurs. Figure (a) is the data at its natural shifting offset (investigated subsequently) and figure (b) is a plot of many traces in increments of 5° separated for visual clarity of rotation trend for each peak.

Mirrored peaks ($a_{\pm 2}$ for example), have shifts in opposite directions, but of the same magnitude.

Figure 4.8(a) is the actual measured signal with original background included for chosen traces. As we move from $\theta = 0°$ to $\theta = 45°$, not much extra rotation occurs in the background, off of EIT peaks. However, moving from $\theta = 45°$ to $\theta = 90°$, the polarization has been rotated to start with before the two-photon NMOR occurs. This indicates that there is another source of NMOR with changing field angle, single-photon NMOR. We therefore want to observe the polarization rotation of the entire single-photon absorption resonance at a range of field angles.

Figure 4.9 displays how single-photon NMOR rotates the polarization dependent on θ. Figure 4.9(a) sweeps the laser in a much larger range than EIT peaks occur, so
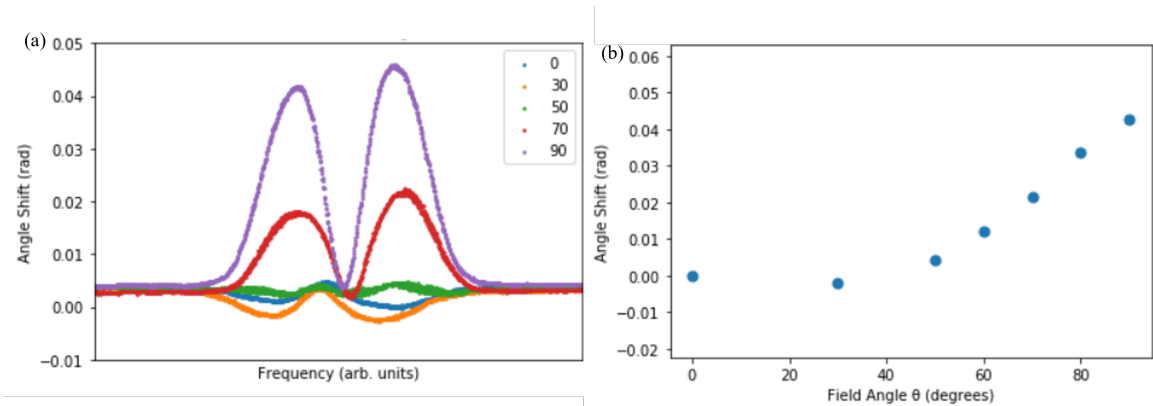
Figure 4.9: The differential photodiode's measurements of the rubidium cell's rotation of polarization at different field θ angles for the full absorption resonance we lock to. The rightmost difference peak in (a) is our resonance, however the left peak acts in a similar way. (b) tracks the resonance's difference amplitude and plots how the polarization angle shifts as the field angle changes.

they are not visible. The rightmost rotation peak in Fig. 4.9(a) is the frequency location of the resonance we lock to, and EIT occurs in a small frequency range around the peak. The leftmost peak is a nearby resonance we don't interact with that acts in a very similar way dependent on field angle. Figure 4.9(b) tracks the top of the resonance peak for visual clarity as we move through θ angles. There is a clear small rotation in one direction before rotating significantly in the other as we approach $\theta = 90°$. It is important that these data are recorded as this shift needs to be taken into account when calibrating laser polarization-dependent methods for determining magnetic field angle. Without subtracting shifts like this, field direction measurements can be up to 0.05 radians off (about 2.9°), defeating the goals of accuracy this device requires.

4.5 Insights from the Lock-in Amplifier

As explained in Section 3.2.1, we must understand the lock-in amplifier's interaction with our system to enable polarization peak locking. We primarily work around

39

polarization angles that cause maximum transmission (peaks in Fig. 2.5 or Fig. 4.15) where the lock-in signal should be zero, so the object of this section is to determine the sensitivity of the lock-in signal to small changes in polarization off a peak. This is examined through both EIT transmission peaks dependent on polarization and transmission peaks based on the alignment of the LCVR with a PBS. For reference, we achieve maximum transmission at the $a_{\pm2}$ EIT peaks when $\phi = 0°$ and $\phi = 90°$, as seen in Fig. 4.10.
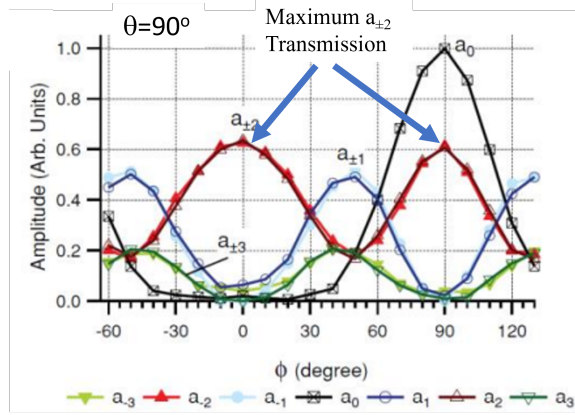

Figure 4.10:   The locations of maximum transmission for the $a_{\pm2}$ peaks are $\phi = 0°$ and $\phi = 90°$.

The lock-in amplifier requires a modulating signal to lock onto, so a sine wave is required with a DC offset $V_{DC,max}$ that rotates the LCVR to a transmission peak angle ϕ_{max} and modulates around it. We connect the photodetector output to the lock-in input and the LCVR's input to the lock-in's carrier wave reference input, as the photodetector output will have the same frequency as the frequency the LCVR rotates. The frequency of this carrier signal is important due to a key finding from this study: the LCVR is relatively slow. It has a reaction time to instantaneous input voltage changes close to 20 ms, meaning only frequencies of 50 Hz or lower are theoretically viable to achieve predictable signals. If the frequency is set much higher, a destructive effect occurs where LCVR DC offsets cause the lockin signal to move

up and down slowly with no recognizable pattern. The amplitude of the drifting increases about 10 times from one transmission peak to the next. The studies in Sections 4.1 and 4.2 reinforce that the LCVR's reaction time is a limitation.

Choosing a LCVR modulation frequency of 23 Hz, we examine the behavior of the lock-in signal at the a_0 and a_{+2} EIT peaks. Referring to Fig. 4.10, the a_{+2} peak has two angles at which transmission is a maximum in this transverse field, while the a_0 peak only has one. We correct for this by setting the polarization with a $V_{DC,max}$ that solely maximizes the a_{+2} peak and adding a half-wave plate after the LCVR to rotate the polarization 90° back to maximum transmission. We take 50 second traces of the lock-in signal at each peak in Fig. 4.11(a) and (b) and Fourier transform the signals in Fig. 4.11(c) and (d). A third signal taken at the first peak using purely cross-polarization (no atomic interactions) is used as a control.

In Fig. 4.11(a) and (b), the higher DC offset causing significantly more drift is very apparent when comparing the Peak 2 ($\phi = 90°$) signals to their Peak 1 ($\phi = 0°$) counterparts. Another important observation, however, is that this drift occurs for both the a_0 and a_{+2} regimes similarly, removing the m-value as a possible cause. Finally, there is no clear difference between the $\phi = 0°$ traces using the vapor cell interaction versus using PBS cross-polarization; the only difference in these example signals is the 'random' nature of the fluctuations causing a new trace shape each time.

Figures 4.11(c) and (d) display the Fourier noise frequency breakdown of the signals. It shows the amount of noise present at each of the frequencies that makes up the total signal. Again, the results in (c) and (d) are similar, ruling out m dependence as a cause for the drift. We use a lock-in integration time constant of 1 second in these measurements, which explains why the noise for frequencies above 1 Hz is low. That is primarily electronic noise from the lock-in output, as nothing from the optical signal surpassing 1 Hz passes through the lock-in.
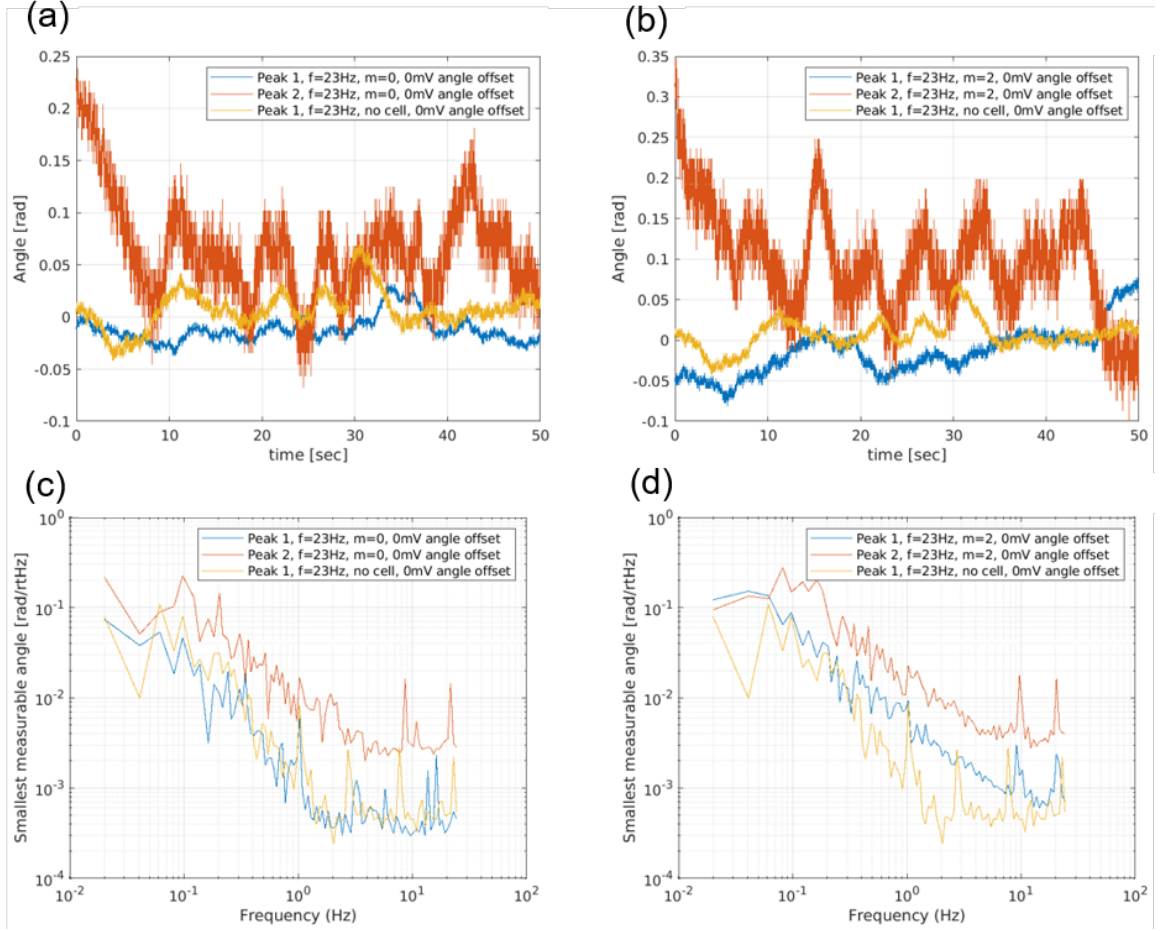
Figure 4.11: (a) a_0, 50-second traces of lock-in amplifier outputs at two transmission maxima, $\phi = 0°$ and $\phi = 90°$, Peak 1 ($\phi = 0°$) having traces using both EIT and cross-polarization. (b) is the same as the first but with a_{+2}. (c) Fourier transform frequency analysis of plot (a). (d) Fourier transform frequency analysis of plot (b).

At frequencies below 1 Hz, we start to capture fluctuations in polarization. Noise here is significantly higher (the plots are logarithmically scaled), suggesting that there are fluctuations in the LCVR or rubidium cell affecting the stability of polarization.

To measure the lock-in amplifier's performance, we employ two other detectors that could potentially perform the same task: the two-channel difference output of the post-cell photodetector and the pre-cell polarization detector, both represented in Fig. 3.1.

We record the output of each of the two detectors and the total transmission for 100 seconds while on and slightly off the first polarization peak, in plane with $\vec{B}$ and $\vec{k}$. Half-wave plates in front of the pre- and post-cell detectors are rotated such that the difference measured is zero; this indicates the light is re-polarized at a 45° angle. In this regime, the zero-crossing of these two signals is at maximum transmission, just like the lock-in amplifier's signal.
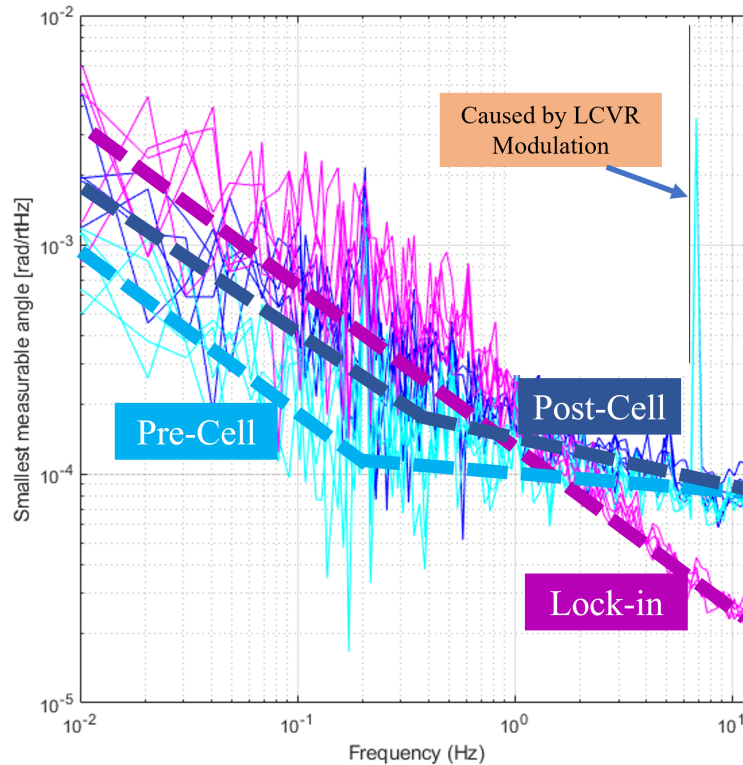

Figure 4.12: Fourier transform frequency-noise analysis of the lock-in amplifier, pre-cell detector, and post-cell detector signals over 100 seconds. Approximate trendlines overlay each trace for clarity.

Figure 4.12 shows the Fourier frequency-noise breakdown of this data. On the bottom right of the plot, the lock-in noise at higher frequencies is significantly smaller due to the 1 second time constant filtering out faster noise above 1 Hz. A massive peak in noise occurs for both the post-cell and pre-cell detectors on the right side

of the plot due to necessary modulation used on the LCVR's input as described in Section 4.2. Given the lock-in's drift discovered earlier in the section, it is unsurprising that it performs worse than the two detectors at sub-1 Hz frequencies. However, we don't need to be able to resolve a signal for the polarization lock over 10 let alone 100 seconds, there needs to be a smooth response in short periods of time to give useful feedback. Therefore, we use the lock-in amplifier and not a differential detector for the source of feedback.

Another interesting point to observe is the fact that the pre-cell detector resolves signals noticeably better than the post-cell detector over all frequencies. We discovered in Section 4.4 that this is due to various atomic effects inside the rubidium cell including nonlinear magneto-optical rotation that cannot be remedied without loss of performance.

4.6 Polarization Lock for the Azimuthal Angle

With the lock-in amplifier optimized and chosen as the preferred feedback mechanism for the LCVR locking loop, we build a feedback loop by incorporating a proportional-integral-derivative (PID) controller and a summing circuit. The PID controller uses the absolute error from zero, information about the error over time, and the trajectory of the error for the future to rapidly output a voltage that is the difference between the LCVR's current voltage and that which would bring it to the transmission peak. We sum the hard-set function generator voltage and the PID voltage and apply the sum to the LCVR. The result is, when pre-set within a generous range of the polarization peak, the LCVR is locked to the rotation of maximum output.

There are a few key effects of the lock. First, it does some work to correct noise and imperfections in the LCVR. Though it is not perfect in this regard, this noise has not been a major issue. Second, if we manually change the raw LCVR input on

the function generator, the lock will ensure the sum of the new input and the error signal is the same as it was before changing it. Third, to be subsequently tested, the lock should correct the LCVR polarization when the azimuthal ϕ angle is changed with the magnetic field coils rather than with the LCVR.

To test the last point, we use Eq. 3.2, or its approximation

$$\phi \approx \frac{1}{2}\left(\frac{\Delta I}{\Sigma I}\right),\tag{4.4}$$

to calculate the angle the polarization has rotated. Again, ϕ is the angle in radians, ΔI is the difference between the intensities measured by each of the balanced photodetectors in either the pre- or post-cell setup (effectively the output of the difference signal), and ΣI is the total transmission between the two detectors. The sum is easily measured on the post-cell detector, but for the simplistic pre-cell detector we must rotate the half-wave plate to make the difference zero and cover one detector. We can double the output to find the sum. The reason for the approximation from Eq. 3.2 is that arc-sine is near linear for a significant domain around zero.

The polarization lock is set at $\phi = 0°$ and left active with only this initial setting. 100 second traces of both detector difference and sum signals are taken at field angles $\theta = 90°$, $\phi = 0°, 2°, 4°, 6°, 8°, 10°$. These are averaged and plugged into Eq. 4.4.

From Fig. 4.13, the average difference between angles measured by the pre-cell detector was $2.05° \pm 0.06°$. The uncertainty is derived from another fourier transform noise breakdown, in which the average smallest measurable angle is about $\pm 0.06°$. This result assures that the feedback loop sets the LCVR to the correct angle to stay in-plane with the magnetic field and laser.

On the other hand, the post-cell detector consistently reads the angle shift as much larger than it should have. In this run, we calculated about $4.32°$ of rotation on average. There are certainly effects like NMOR discussed in section 4.4 which
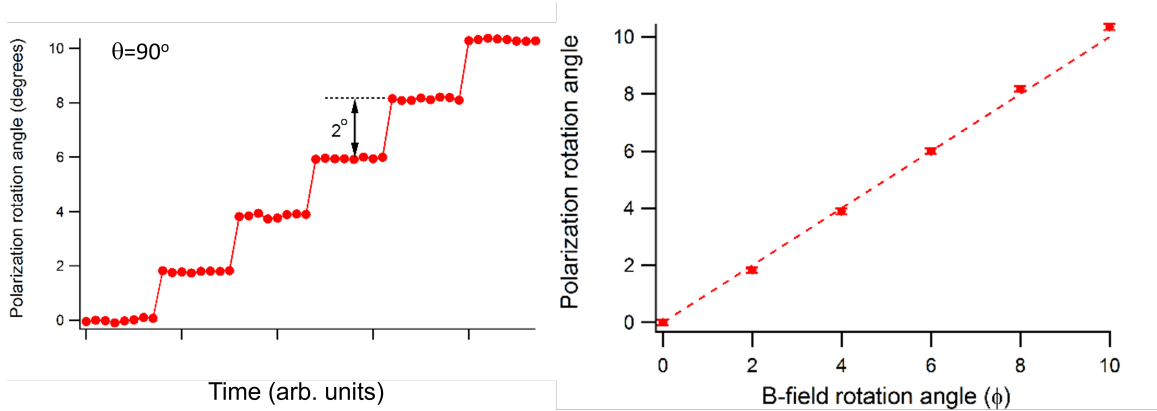
Figure 4.13: Pre-cell detector measurement of ϕ when field is rotated in steps of 2° in the xy-plane while $\theta = 90°$. The left plot shows the measured ϕ while the field sits at various angles, each step visible over time. The right plot is a traditional plot of intended rotation versus average measured result.

affect polarization readings after the rubidium cell. Thus we must rule the post-cell detector out as a viable candidate to report on polarization.

Next we must observe the lock's effectiveness when the θ angle shifts closer to 0, where the field $\vec{B}$ is parallel to the laser vector $\vec{k}$. In theory, it should start performing worse when closer, as the change in EIT peak amplitude over a sweep in ϕ becomes smaller as shown in Fig. 2.5(b). This gives the lock much less information as to where maximum transmission occurs, especially given any noise present.

Figure 4.14 shows the results from taking a trace of the pre-cell detector at field angles between perpendicular (90°) to parallel (0°) with the laser vector. There is a very clear increase in the smallest measurable angle as the field moves towards 0°, indicating a large loss in sensitivity. In angles closer to 0°, the detector's difference signal moved dramatically up and down as the locking mechanism struggled to find the exact angle to lock to. This result shows that, using this locking mechanism, the magnetometer will ideally need to measure ϕ when at or closer to $\theta = 90°$.
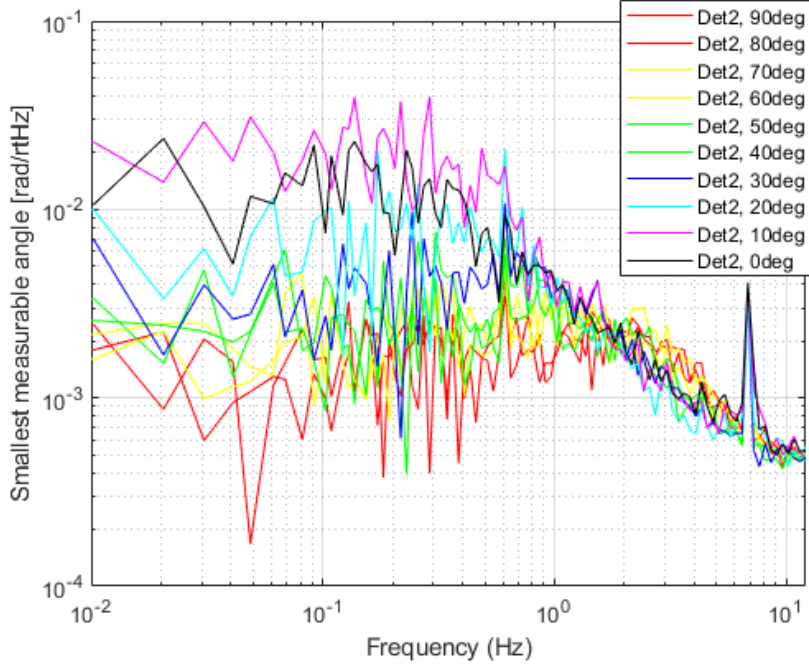
Figure 4.14: Fourier transform frequency-noise analysis of the pre-cell detector as θ decreases from 90° to 0°, longitudinal with the laser.

4.7 Measuring the Polar Angle

While the ability to measure θ has not been demonstrated yet, work done this year has suggested some viable methods. First, the stated method for determining θ in Section 2.4 is to compare the difference in maximum and minimum intensities achieved on the $a_{\pm 2}$ peaks, since this difference decreases and theoretically goes to 0 as θ goes to 0°.

Figure 4.15 shows how the transmission intensity changes as ϕ is rotated when at different θ. This graph does not look like the theoretical Fig. 2.5 because there is a large background caused by optical pumping. Fig. 2.5 measures the EIT peak amplitude while Fig. 4.15 measures total transmission. Optical pumping (somewhat like coherent population trapping) moves atoms to a higher energy state, rendering them unable to absorb more photons; the transmission increases accordingly depend-
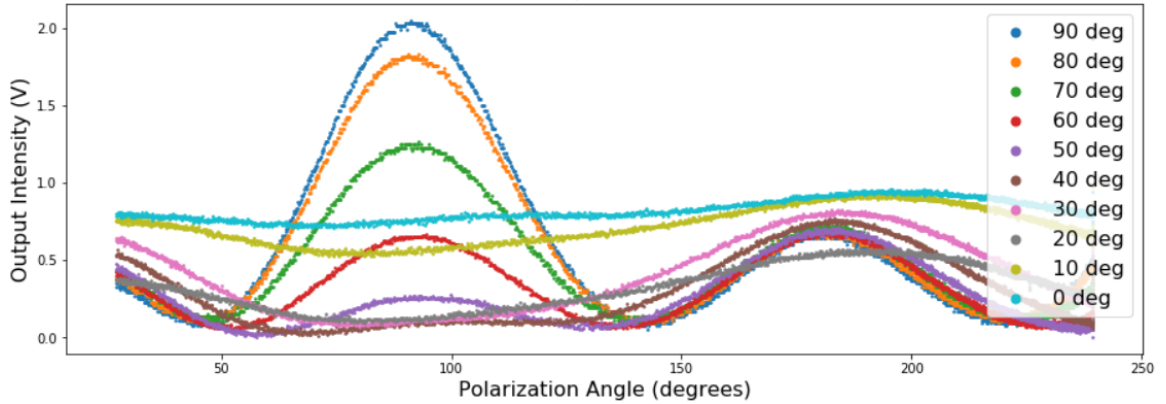
Figure 4.15: Transmission intensity of the a_{+2} peak as ϕ is rotated with the LCVR at different θ. Note the x-axis is the LCVR's polarization angle relative to the x-axis in Fig. 2.4, so 90° is in-plane with the $\vec{B}$-$\vec{k}$ plane.

ing on the laser polarization and magnetic field. We need to determine a method that bypasses the effect of optical pumping to determine θ.
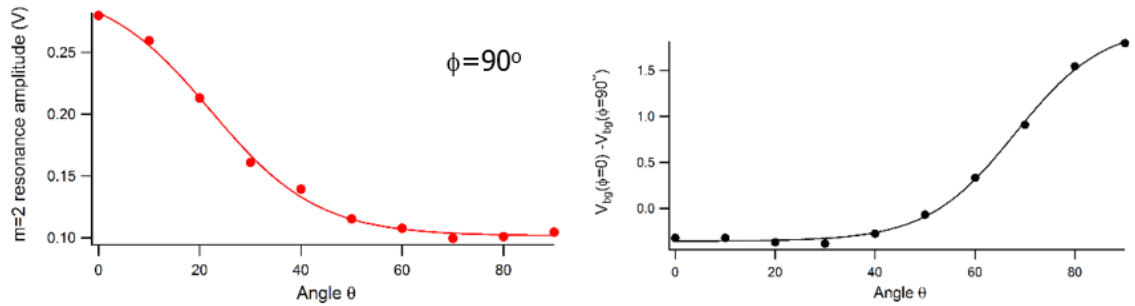

Figure 4.16: EIT peak a_{+2} amplitude at $\phi = 90°$ (perpendicular to the $\vec{B}$-$\vec{k}$ plane) (left) and background level (right), at various θ.

As such, Figure 4.16 has the amplitude of the a_{+2} EIT peak at $\phi = 90°$ and the background level in the same regime. The plots are nearly opposites of one another, with one having a high derivative when the other doesn't, and vice versa. This opens up a promising opportunity to measure θ with precision using a combination of the $\phi = 90°$ peak intensity and the measured background level.

In Section 4.4, we discovered a method to differentiate between θ with a component in Fig. 2.4's positive z direction versus its negative z direction. This is achieved

with nonlinear magneto-optical rotation. Unfortunately, no method exists at the moment to differentiate between the positive and negative x and y directions. Future work on this magnetometer will search for more effects that differentiate geometrical octants and also use promising machine learning methods to break the symmetry and determine the exact field direction.

Chapter 5

Conclusions

Our role in the first year of this four-year project was a success. We successfully built, tuned, and modified the vector atomic magnetometer to be a useful baseline for collaborators to refine precision and automate the process of magnetic field vector measurement.

Installing and calibrating the Liquid Crystal Variable Retarder enabled fast, precise control of the field's apparent azimuthal angle ϕ both manually and under a feedback-driven polarization locking scheme. Observing the system's response to added quick polarization modulation allowed the employment of the lock-in amplifier, critical for the polarization locking feedback loop. Establishing the frequencies where we find EIT peaks under a known field strength informs the calibration of field coil inputs to create angular fields and ensures easy hopping between peaks. Recording the effect of nonlinear magneto-optical polarization rotation at EIT peaks and on the overall resonance under different field angles informs future calculations and measurement algorithms and provides a simple method for distinguishing fields at supplementary angles. Establishing the functional polarization locking mechanism fulfils the vital function this device requires to measure both characteristic angles of magnetic fields, especially the azimuthal ϕ. Finally, we have made substantial progress in demonstrating the best method for precise measurement of the polar θ.

Future work to make this magnetometer precise, autonomous, and smaller is already underway. More sensitive methods of measuring θ are under investigation, using every EIT peak rather than the $a_{\pm 2}$ only and applying machine learning methods. These methods may be useful in breaking the remaining two x and y symmetries in field direction, while analog methods like NMOR analysis may still be discovered. The research described herein is vital to prove the magnetometer's concept and provide a working prototype to those who will make this chip-sized vector atomic magnetometer possible.

References

[1] Alan Edelstein. "Advances in magnetometry". en. In: *J. Phys.: Condensed Matter* 19.16 (Apr. 2007). Publisher: IOP Publishing. ISSN: 0953-8984. DOI: 10.1088/0953-8984/19/16/165217.

[2] Corey J. Cochrane et al. "Vectorized magnetometer for space applications using electrical readout of atomic scale defects in silicon carbide". en. In: *Sci Rep* 6.1 (Nov. 2016). ISSN: 2045-2322. DOI: 10.1038/srep37077.

[3] Kevin Cox et al. "Measurements of the magnetic field vector using multiple electromagnetically induced transparency resonances in Rb vapor". In: *Phys. Rev. A* 83.1 (Jan. 2011), p. 015801. DOI: 10.1103/PhysRevA.83.015801.

[4] B. D. Agap'ev et al. "Coherent population trapping in quantum systems". en. In: *Phys.-Usp.* 36.9 (Sept. 1993). DOI: 10.1070/PU1993v036n09ABEH002306.

[5] V. I. Yudin et al. "Vector magnetometry based on electromagnetically induced transparency in linearly polarized light". In: *Phys. Rev. A* 82.3 (Sept. 2010). DOI: 10.1103/PhysRevA.82.033807.

[6] O. Gravrand et al. "On the calibration of a vectorial 4He pumped magnetometer". In: *Earth, Planets and Space* 53.10 (Oct. 2001), pp. 949–958. ISSN: 1880-5981. DOI: 10.1186/BF03351692.

[7] W. Happer and H. Tang. "Spin-Exchange Shift and Narrowing of Magnetic Resonance Lines in Optically Pumped Alkali Vapors". In: *Phys. Rev. Lett.* 31.5 (July 1973). Publisher: American Physical Society, pp. 273–276. DOI: 10.1103/PhysRevLett.31.273.

[8] Nathan T Belcher. "Development of a Prototype Atomic Clock Based on Coherent Population Trapping". Senior Thesis. William & Mary, May 2008.

[9] L. Essen, E. G. Hope, and D. Sutcliffe. "Hyperfine Splitting of Rubidium-87". en. In: *Nature* 189.4761 (Jan. 1961), p. 298. ISSN: 1476-4687. DOI: 10.1038/189298a0.

[10] S. Brandt et al. "Buffer-gas-induced linewidth reduction of coherent dark resonances to below 50 Hz". In: *Phys. Rev. A* 56.2 (Aug. 1997). DOI: 10.1103/PhysRevA.56.R1063.

[11] C. Svelto et al. "High-resolution spectroscopy of the 39K transitions at 770 nm and 13C2H2 saturated lines by a solid-state laser at 1.54 μm: toward an accurate frequency standard in the optical communication band". In: *Instrumentation and Measurement, IEEE Transactions on* 51 (Sept. 2002), pp. 764–769. DOI: 10.1109/TIM.2002.803302.