Infrared ECDL Locking Techniques for a Two-Step Transition in Lead

by
Julia M. Matin

Professor Protik K. Majumder, Advisor

Independent Study Final Report

WILLIAMS COLLEGE
Williamstown, Massachusetts
June 2025

Contents

Chapter 1
Introduction

1.1 Motivation

The Standard Model (SM) is currently the best explanation of the universe on subatomic scales. Generally, it encapsulates our understanding of elementary particles and how they interact via three of the four fundamental forces. The Standard Model, however, is incompatible with other well-established physics. Among other things, it falls short at describing the matter-antimatter asymmetry, neutrino oscillations, dark energy, and dark matter. Similarly, it is inconsistent with general relativity, failing to incorporate gravity as it does with forces which operate on comparably smaller scales. These inconsistencies necessitate physics "Beyond the Standard Model" (BSM).

Many BSM experiments require expensive equipment and large-scale infrastructure, typically built deep underground to create a stable environment. CERN's Large Hadron Collider (LHC) accelerates particles close to the speed of light, creating high-energy collisions to search for new particles and improve our current understanding of fundamental forces. The IceCube Neutrino Observatory drills thousands of optical sensors deep into Antarctic ice to observe neutrino oscillations and search for decay products of potential dark matter candidates. Atomic, Molecular, and Optical (AMO) physicists, by comparison, use precise, low-energy tests to probe new physics. The Majumder Lab in particular employs table-top absorption spectroscopy experiments—using lasers, furnaces, and small atomic samples—to study the structure of heavy, multi-valence-electron atoms. We measure characteristics such as electronic transitions amplitudes, polarizabilities, and hyperfine splittings to extreme precision, measurements which serve as experimental benchmarks to test the accuracy of ab-initio theory calculations.

1.2 Prior and Ongoing Work

BSM interactions are thought to significantly intensify with atomic number, as we deviate farther from the well-established picture of hydrogen. Discrepancies between theory work founded on first principles and precise, experimental measurements of heavy atom characteristics are therefore a good place to search for new physics. For this reason, the lab partners

with the theoretical physics lab led by Marianna Safronova at the University of Delaware. Beginning from first principles, the Safronova group applies mathematical approximations to make otherwise unsolvable problems tractable. Previous measurements from the Majumder lab include transition amplitudes and atomic polarizabilities resulting from the Stark effect, both of which are good tests of wavefunctions far from the nucleus. The lab also measures isotope shifts and hyperfine structure, which provide insight into the wavefunction close to the nucleus. Heavy atoms are similarly good testbeds for parity non-conservation (PNC) effects, which refer to spatial asymmetry under coordinate inversion (i.e. $\vec{r} \to -\vec{r}$). This symmetry-breaking is a unique feature of the weak interaction, and its effects grow rapidly with atomic number.

Previously, the subject of the lab has been indium (Z = 49) and thallium (Z = 81), both relatively soft metals with three valence electrons. In recent years, the lab has shifted its attention to lead, which was a natural next choice given its chemical similarities to indium and thallium, as well as its convenient overlap in laser wavelengths required for spectroscopy.

1.2.1 Properties of Lead

Lead has 82 protons and 4 valence electrons, which puts it in the Group IV elements [8]. Since radioactivity in the periodic table is typically considered to begin with Polonium (Z = 84), lead is one of the heaviest, stable elements[1]. There are four naturally occurring lead isotopes: ^{204}Pb (which makes up 1.4% of lead in natural abundance), ^{206}Pb (24.1%), ^{207}Pb (27.1%), and ^{208}Pb (52.4%)[2]. We have all of these isotopes in various combinations in the lab, samples sent to us in vacuum-sealed quartz cells which we can vaporize in tabletop furnaces. Lead has a melting point of 327.5°, and temperatures greater than 750° are sufficient to vaporize it at low pressure. Each isotope has slightly different energy levels, referred to as "isotope shifts" which arise from two main effects: the mass shift, due to differences in nuclear mass, and the field shift, resulting from differences in nuclear charge distribution. Only ^{207}Pb, however, exhibits hyperfine structure, an interaction between the magnetic moment of the nucleus and the magnetic field produced by the electrons, which results in a subtle splitting of energy levels.

The lab is currently studying a pair of two-step transitions, sequential excitation processes in which an atom absorbs two photons to reach a higher-energy state via an intermediate level. The transition examined in this report begins with the $6s^2 6p^2\ {}^3P_0 \to\ {}^3P_1$ magnetic dipole (M1) transition using a 1279 nm laser followed by the $6s^2 6p^2\ {}^3P_1 \to\ 6p7s\ {}^3P_0$ electric dipole (E1) transition using a 368 nm laser (Fig. 1). The goal of this project was to explore two schemes for stably locking the laser driving the M1 transition, thereby enabling precise future measurements of the E1 transition.

[1] Before it was discovered to be toxic, lead's durability and corrosion-resistance caused it to be common for material for pipes dating back to Roman times. The symbol for lead Pb actually comes from its Latin translation plumbum, also the root of the word plumber.

[2] The missing isotope ^{205}Pb is radioactive and therefore not naturally occurring.

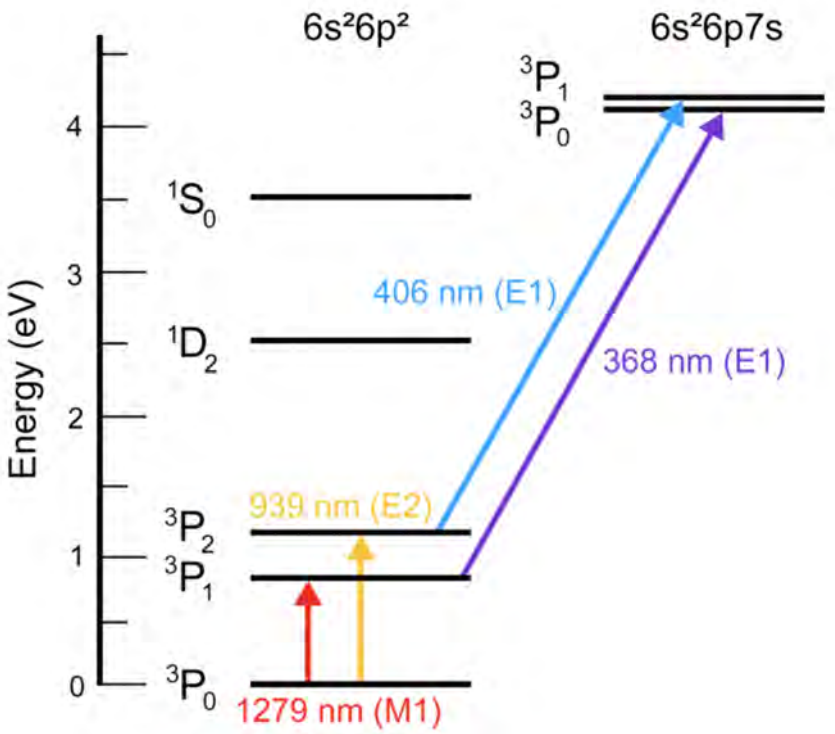

Figure 1: Low-Lying Energy Levels in Lead (figure created by Robin Wang '24). This report will focus on stabilizing the 1279 nm M1 transition in order to make measurements of the 368 nm transition.

Chapter 2
Atomic Theory and Spectroscopy Setup

In this chapter, we provide an overview of the atomic theory relevant to the experiment, aiming to offer a broad conceptual understanding rather than detailed derivations[3]. We will also lay out the general mechanics of the two-step excitation scheme, including how we solve the issue of Doppler broadening and velocity-changing atomic collisions. Finally, we emphasize the importance of stably locking the first step of the transition in order to precisely measure characteristics of the second step.

2.1 Atomic Structure

The first foundational idea in quantum mechanics was that energy exists, not in a continuum as classical physics assumes, but in a finite number of quantized states localized in space. At the turn of the twentieth century, Max Planck introduced energy quanta to understand the unexpected drop-off in intensity emitted from blackbodies at high frequencies, resolving the so-called "ultraviolet catastrophe." Not long after, Einstein applied the idea of energy quanta to explain the photoelectric effect, in which the energy of electrons ejected from a metal depends, not on the intensity of incoming light, but on the number of photons it contains. These two discoveries initiated a radical rethinking of atomic structure, grounded in the quantization of energy levels.

This new image of the atom became known as the Bohr model. In this model, electrons occupy discrete, quantized energy levels characterized by the principal quantum number n, an integer that increases with distance from the nucleus. The farther these energy levels are from the nucleus, the smaller their spacing becomes due to the weakening of the Coulomb potential. Eventually, they converge to the atom's binding energy, beyond which electrons are free and *can* exist in an energy continuum.

2.1.1 Hydrogen

To further understand the structure of heavy, complex atoms like lead, it is useful to begin with a simpler model: the hydrogen atom. Along with a few idealized systems—the free

[3]For more in-depth derivations, see [3]

particle, the particle in a box, and the particle in a harmonic oscillator potential—the hydrogen atom is presented extensively in any introductory quantum mechanics textbook given that it is one of the few systems with an exact analytical solution to the time-independent Schrödinger equation:

$$\hat{H}\,\Psi = E\,\Psi \tag{1}$$

Where the Hamiltonian operator $\hat{H}$ is a sum of the kinetic and potential energies of the electron:

$$\hat{H} = \hat{T} + \hat{V} = -\frac{\hbar^2}{2m_e}\nabla^2 - \frac{e^2}{4\pi\varepsilon_0 r} \tag{2}$$

Solving the Schrödinger equation is essentially a two-fold task. First, we would like to know about the probability density of the electron (i.e. where in space it is likely to be), which is equal to the absolute square of the wavefunction ψ:

$$P(\vec{r}) = |\psi(\vec{r})|^2 \tag{3}$$

Additionally, we want to find the discrete energy levels E_n that the electron may occupy, which come from finding the eigenvalues of the Hamiltonian. This is where our primary interest in spectroscopy lies. Neutral hydrogen has an especially simple potential landscape that allows us to analytically solve for these values exactly. It is simply the Coulomb interaction of the proton and electron:

$$V(r) = -\frac{e^2}{4\pi\varepsilon_0 r} \tag{4}$$

Plugging this potential into the Hamiltonian, we can solve for exact the energy levels of hydrogen:

$$E_n = -\frac{m_e}{2\hbar^2}\left(\frac{e^2}{4\pi\varepsilon_0}\right)^2 \frac{1}{n^2} = -\frac{13.6\,\text{eV}}{n^2} \tag{5}$$

2.1.2 The Many-Body Problem

Now, what if we wanted to solve the Schrödinger equation for slightly more complex atoms? With each additional electron, we must add new terms to our Hamiltonian. For example, the Hamiltonian for H^- (which still has a single proton but two electrons, making it a "three-body" problem) is:

$$\hat{H} = -\frac{\hbar^2}{2m_e}\nabla_1^2 - \frac{\hbar^2}{2m_e}\nabla_2^2 - \frac{e^2}{4\pi\varepsilon_0 r_1} - \frac{e^2}{4\pi\varepsilon_0 r_2} + \frac{e^2}{4\pi\varepsilon_0 r_{12}} \tag{6}$$

where the first two terms come from the kinetic energy of the two electrons and the second two terms come from the electrostatic potential between each electron and the proton. Although this expression may not look much more complicated than our original hydrogen Hamiltonian, the final term representing the Coulomb repulsion force between the two electrons makes it already too complicated to solve the Schrödinger equation exactly. We may use approximation methods (the Hartree-Fock method, for example), yet each of these approximations introduces uncertainty. If no exact solution exists for these simple atoms, then what happens when we have three, four, or even eighty-two electrons? Each additional electron adds an another layer of complexity that makes the task increasingly difficult. Luckily

for us, the mathematics of this problem falls to the Sanfranova lab, while we rely on experimental techniques of absorption spectroscopy to measure these transitions directly.

2.1.3 Hyperfine Structure

Beyond the gross atomic structure described above, smaller atomic interactions can break degeneracies—that is, multiple distinct quantum states with equal energy—creating finer energy level splittings. One such interaction gives rise to hyperfine structure, a smaller yet measurable energy splitting due to the interaction due to the motion of an electron around the spinning nucleus. Just as electrons moving in a solenoid produce magnetic fields, so too do electrons orbiting a nucleus, though on a much smaller scale. The resulting magnetic dipole moment can then interact with the magnetic moment of the nucleus. This results as an extra term in the Hamiltonian:

$$\hat{H}_{\text{hf}} = A\,\vec{I}\cdot\vec{J} \tag{7}$$

where A is the hyperfine constant, $\vec{I}$ is the nuclear spin, and $\vec{J}$ is the sum of the electron's angular momentum and spin ($\vec{J} = \vec{L} + \vec{S}$). This results in correction to the energy:

$$\Delta E_{\text{hf}} = \frac{A}{2}\left[F(F+1) - I(I+1) - J(J+1)\right] \tag{8}$$

Where the total angular momentum of the atom F is:

$$F = |I - J|,\ |I - J| + 1,\ \ldots,\ I + J \tag{9}$$

This interaction only affects ^{207}Pb because it is an odd isotope and therefore has a nonzero nuclear spin. The $^3P_1$ level ($I = 1/2$ and $J = 1$) splits into $F = |1/2 - 1| = 1/2$ and $F = 1/2 + 1 = 3/2$. The $F = 1/2$ and $F = 3/2$ levels have energies of A greater than and A/2 less than the $^3P_1$ level, respectively. We can use the infrared laser to excite electrons to either level, using slightly different wavelengths, and then scan the UV laser to complete the two-step transition.

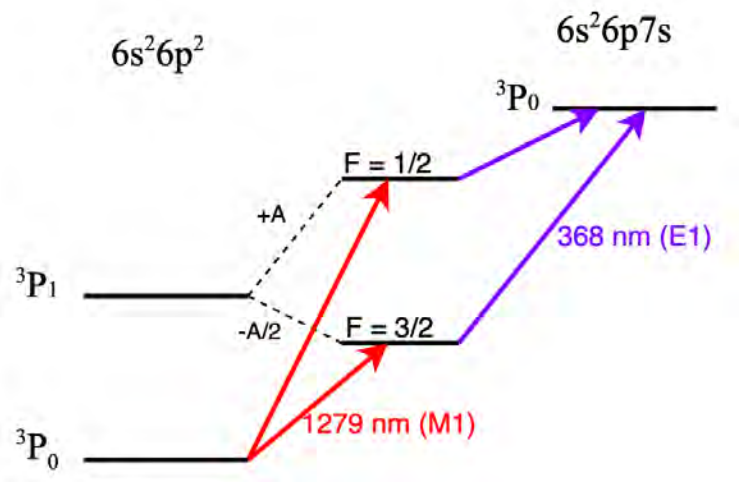

Figure 2: ^{207}Pb Hyperfine Splitting Diagram

2.2 Absorption Spectroscopy

Atomic absorption is a good illustration of both energy conservation and the quantization of atomic energy levels. As discussed earlier, each element in the periodic table has a unique and discrete set of electronic energy levels, an atomic fingerprint which differentiates hydrogen from helium, helium from neon, and so on. If a photon hits an atom with an energy level spacing equal to the photon's energy, where we can relate photon energy and frequency or wavelength using

$$E = hf = \frac{hc}{\lambda} \tag{10}$$

then the electron can absorb the photon's energy and jump to a higher energy level within the atom. If we shine light containing a continuum of photon wavelengths at a population of vaporized atoms, therefore, only the wavelengths corresponding to difference in the atoms' energy levels will be absorbed. The resulting spectrum of exiting light will contain dark lines or dips in intensity which provide insight into the energy levels of an atom. This is the basis of absorption spectroscopy. Similarly, emission spectroscopy relies on electrons jumping from higher to lower energy levels, emitting photons in the process, although this technique requires a previously excited sample of atoms to fluoresce and is therefore less common[4]. Spectroscopy, therefore, is an incredibly powerful tool that allows astronomers to decode stars, physicians to analyze the chemical makeup of tissues in diagnostic procedures, and atomic physicists to indirectly measure quantized energy states.

2.2.1 Two-Step Transition

As discussed earlier, we wish to look at a two-step transition in lead (Fig. 1). This transition begins with the $6s^2 6p^2\ {}^3P_0 \rightarrow\ {}^3P_1$ magnetic dipole (M1) transition which requires light around 1279 nm, followed by the $6s^2 6p^2\ {}^3P_1 \rightarrow\ 6p7s\ {}^3P_0$, electric dipole (E1) transition using light around 368 nm. Given that the primarily goal of this report is to present locking schemes for the 1279 nm laser, we will not delve too far into the full experimental setup. Instead, we will provide a broad overview of the experiment and key optics, emphasizing the critical role of frequency stabilization of the 1279 nm laser in enabling precise and reliable excitation of the M1 transition.

2.2.2 Optical Pumping to the $^3P_1$ State

At the temperatures required to vaporize lead at low vapor pressure (roughly 750°C = 1023K), only a small fraction of atoms naturally occupy the first excited state due to thermal excitation. We can estimate this fraction using the Boltzmann factor, which tells us that the ratio of naturally excited atoms to ground-state atoms is roughly $e^{-\frac{\Delta E}{k_B T}} \approx 2 * 10^{-6}$. To increase the population in the $^3P_1$ state, we use a 1279 nm laser to optically pump atoms up from the ground state. In doing so, we artificially enhance the number of atoms available for the second-step transition. We can then overlay the UV beam with a dichroic mirror, which

[4]We can also induce these electrons to drop down by hitting them with a photon in a process called stimulated emission. This is the fundamental principle behind laser operation. (LASER = Light Amplification by Stimulated Emission of Radiation.)

selectively reflects some wavelengths while transmitting others[5]. The overlapping beams can then excite the same population of atoms, increasing the size of our absorption dip and enabling a more precise investigation of the E1 transition.

2.2.3 Doppler Broadening and Narrowing

This optical pumping scheme frees us (mostly) of the effects of Doppler broadening. Imagine a population of hot, vaporized atoms moving at high velocities in various directions. Due to the Doppler effect, atoms moving toward the incoming laser light perceive it at a shorter wavelength (blue-shifted), while those moving away perceive it at a longer wavelength (red-shifted). We can quantify the shift in terms of the lab-frame photon frequency f and the atom speed v in the nonrelativistic limit (i.e. $v << c$):

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

This velocity-dependent frequency shift broadens the absorption profile, spreading otherwise sharp spectral features into a wider curve. Without going into too much detail, we can generally describe the broadened dip as a Gaussian[6]. The higher the temperature, the faster the atoms move and the more pronounced the broadening becomes, an effect which obscures the true lineshape and prevents precise transition measurements.

By optically pumping at a fixed wavelength around 1279 nm, we selectively excite atoms that are nearly stationary relative to the laser beam, as only this velocity class of atoms resonate precisely with that wavelength. We can then scan over the second transition using a UV laser in the neighborhood of 368 nm and observe Doppler-narrowed spectral features, allowing for much more precise spectroscopy. This technique is also called "hole-burning," as we are effectively removing a hole from the ground-state velocity distribution. We can similarly mitigate Doppler broadening by sending an overlapping 1279 nm beam in the opposite direction. This co-and-counter-propagating beam configuration ensures that only atoms with near-zero velocity along the beam axis are selected, further narrowing the velocity distribution.

2.2.4 Laser Chopping

Our optical pumping scheme only partially solves the issue of Doppler broadening. According to the kinetic theory of gases—a statistical mechanical description of elastic collisions between moving particles in an ideal gas—collisions between moving atoms and our "stationary" excited atoms can alter their velocities. These velocity-changing collisions can then reintroduce Doppler broadening to our previously narrowed lineshape. To fix this problem, we rapidly switch the 1279 nm beam on and off. This modulation temporarily de-excites our atoms, preventing collisions. We can then use a lock-in amplifier to look only at signals when the beam is hitting the atoms, filtering out much of the noise.

[5]Here, we dichroic mirror specifically designed to transmits IR light while reflecting UV light.

[6]The true picture is actually much more complicated than I describe here. Other mechanisms such as pressure, power, and finite-lifetime broadening also affect the shape of the absorption dip. Similarly, the lineshape is not only a Gaussian but a mixture of a Gaussian and a Lorentzian called a Voigt convolution. To learn more about these effects and the true lineshape, see [1, 2, 6]

Acousto-Optic Modulator

To chop the pumping beam at high frequencies, we use an Acousto-Optic Modulator (AOM) made by Brimrose. This device relies on the Acousto-Optic Effect, in which light interacts with acoustic vibrations in a crystal, creating a periodic modulation of the refractive index that acts like a diffraction grating. Some AOMs operate in the Raman-Nath regime, in which several diffracted beams are produced. Ours, however, operates in the Bragg regime meaning the acoustic wave forms a well-defined grating and light is predominantly diffracted into a single first-order beam [5].

To align the AOM, we set its frequency to about 1 Hz. This allowed us to see the diffracted beam blinking on and off, while the zeroth-order beam varies in intensity as some of its light is allocated to the diffracted beam (these variations are difficult to see by eye). We could have used either of these beams, since we only care about the intensity of light that is turned on and off, so we chose the diffracted beam. We then fine-tuned the mount orientation to maximize diffraction efficiency, which is defined to be the ratio between the diffracted and incident power and increases for shorter optical wavelengths. The 939 nm laser, for example, achieves up to 60% efficiency whereas the 1279 nm laser maxes out at about 40%. After aligning the AOM, we set up the high-frequency lock-in with the AOM frequency as the reference to look only at the portion of the UV light exiting the furnace that has the IR beam chopped on. For this reason, we need a high-speed UV photodetector that can detect quick intensity changes.

2.3 Overall Setup and Remaining Components

We can now bring together the beam configuration, lead vapor cell, furnace, and AOM to describe the full experimental setup, while also introducing a few remaining elements including the laser and devices used to read frequencies. This information will be especially important moving forward, as it provides the foundation for precisely tuning the 1279 nm laser and evaluating the quality of its lock.

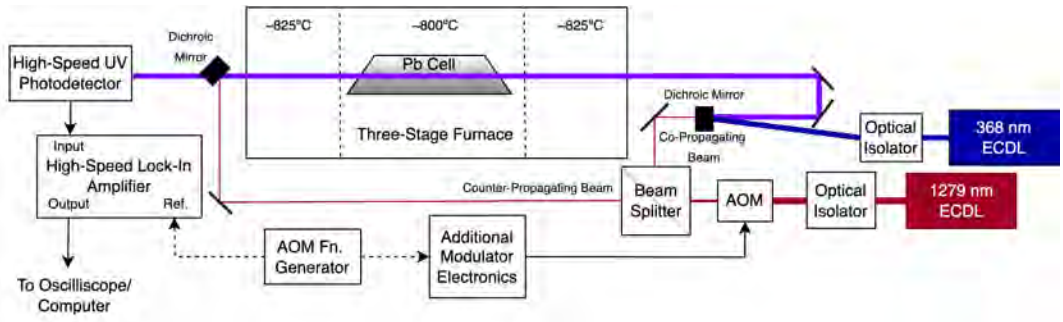

Figure 3: Diagram of Two-Step Transition Experimental Set-Up. By setting the outer furnace stages to a hotter temperature (roughly 25°C more than the center), we can adjust for cooling due to surrounding air currents, keeping the vapor cell close to a constant temperature throughout. Here, we have omitted details about modulator electronics for simplicity (for more information, see [2]).

2.3.1 External-Cavity Diode Laser

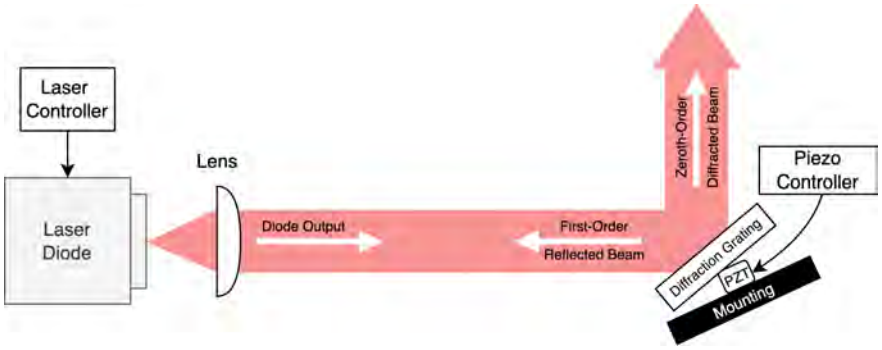

Figure 4: ECDL Diagram

Like all lasers, external-cavity diode lasers (ECDLs) emit a highly coherent, monochromatic light beam through the process of stimulated emission. ECDLs combine a semiconductor diode with an adjustable internal diffraction grating to achieve a narrow linewidth. The diode emits a range of frequencies which hit the diffraction grating. The zeroth-order diffracted beam escapes the laser, serving as the output, while the first-order beam is reflected back into the diode, providing feedback which reinforces the resonant frequency[7]. These lasers are useful for our purposes since they are tunable, allowing us to select out a specific frequency or scan over a range of frequencies to study spectral features. We use a commercial infrared laser (Sacher Lasertechnik TEC 150) for the first transition and a home-built ultraviolet ECDL for the second transition. Directly after each laser, we place an optical isolator which prevents light from reflecting back and potentially interfering with or damaging the ECDL.

There are multiple ways to adjust the laser's emitted frequency of light. The most direct and controllable method is to adjust the diffraction grating angle, shifting the wavelength of the zeroth-order beam. We do this by sending a voltage from the piezoelectric transducer (piezo or PZT), a device which converts voltage into mechanical movement by applying an electric field to a small ceramic piece sitting under the diffraction grating. The ceramic piece then expands or contracts, depending on the sign of the applied voltage, subtly shifting the grating angle. We can therefore scan over a range of frequencies by sending a triangle wave to the piezo. Since changes in the piezoelectric ceramic material depend on mechanical movements, however, the shape of the frequency response does not exactly match the triangle wave. Much like how the position of spring is not directly proportional to applied force due to initial friction, we see distortions in the tuning response. For this reason, frequency "upscans" and "downscans" are not perfectly symmetric.

The other way to adjust the laser frequency is by adjusting the current or temperature of the diode itself, both of which are controlled by the laser control module [4]. Increasing either one raises the temperature of the semiconductor, shifting its bandgap and causing the peak wavelength of the diode's gain curve—the distribution that describes how much

[7]This mechanism is a useful indicator of mode-hopping, in which the laser suddenly switches to a different frequency. Mode-hops are more likely to occur if we are scanning over a longer frequency range, so they should not be a problem when the laser is locked.

amplification the laser provides at each frequency—to shift to lower frequencies (redshift). Conversely, decreasing the diode current or temperature shifts the gain curve towards higher frequencies (blueshift). For reasons I will describe later, we adjust both the piezo voltage and the diode current in our locking scheme.

2.3.2 Measuring Frequencies

Measuring the absolute frequency of light is generally not as simple as reading a number off a device. We need a well-established frequency as a comparison, which, using principles of interference patterns, helps provide insight into our unknown laser frequency. One simple way of doing this is using our 1279 nm wavemeter, which uses the well-known wavelength of a helium-neon (HeNe) laser (632.8 nm) as a reference. If we line up the HeNe beam with the infrared beam, the wavemeter analyzes the resulting interference fringes to provide an exact value of the IR beam frequency (1279 nm = 2343959796716.18 Hz). Additionally, the moving wave display on the wavemeter is useful to know whether we are scanning single mode. If the waves appear to shake a little, the laser is multimode. Although the wavemeter appears to be as simple as reading a number off a screen, therefore, it is still using a well-known reference frequency.

Fabry-Pérot Cavity

Another way to know the frequency of emitted light is by using a Fabry-Pérot cavity (also known as a Fabry-Pérot interferometer), which consists of two parallel, partially reflective mirrors facing each other, forming a resonant cavity. When light enters the cavity, it undergoes multiple reflections between the mirrors, resulting in constructive or destructive interference depending on the wavelength and the spacing between the mirrors. We may describe the resonant frequencies in terms of the cavity length L and integer n:

$$\nu_n = \frac{cn}{2L} \tag{11}$$

This interference condition leads to the transmission of only specific wavelengths, effectively creating a high-resolution spectral filter. (The sharpness of our Fabry-Pérot peaks can also indicate whether the laser is scanning single-mode.) As we scan over a range of frequencies using the 1279 nm laser, we see fringes corresponding to the resonant frequencies which, for our particular cavity, are 500 MHz apart. This is a useful, known frequency which allows us to convert between voltage of absorbed light and frequency. For example, if we scan over a range of frequencies, getting an absorption dip with time on the x-axis and voltage on the y-axis, as in Fig. 5, we can then quantify the slope of the side of the absorption dip using the even spacing of our superimposed Fabry-Pérot peaks.

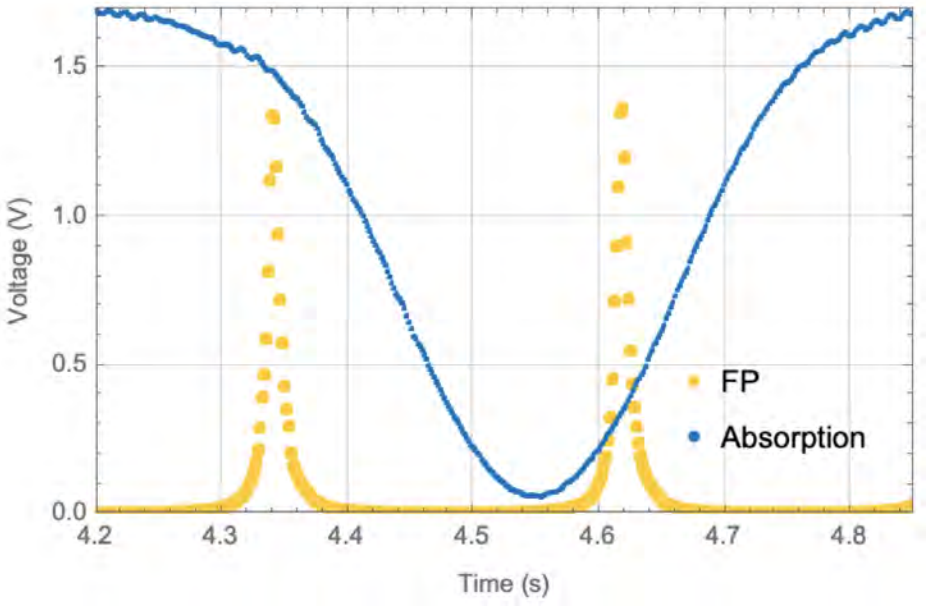

Figure 5: Sample Absorption Dip and Fabry-Pérot Peaks During Scan of 1279 nm Laser

If we multiply the inverse slope of the original absorption dip by a factor of 500 MHz per a complete Fabry-Pérot cycle time, we get a conversion factor in units of MHz/V:

$$\frac{df}{dV} = \left(\frac{dt}{dV}\right)\left(\frac{500\text{MHz}}{\text{FP cycle time}}\right) \tag{12}$$

Using Fig. 5 as an example, we can use the slope to the left of resonance to find a conversion factor of roughly 200 MHz/V. We will later use this conversion method to find the frequency noise and assess locking quality.

Well-Established Spectral Features

A third method of measuring frequencies is to compare them to well-established frequencies of spectral features in lead. For example, we know that the hyperfine splitting between the F = 1/2 and F = 3/2 levels of ^{207}Pb is 3583 MHz[8]. This value later comes in handy to assess locking quality when we no longer have the Fabry-Pérot on the table.

2.3.3 Importance of Locking the 1279 nm ECDL

As we can now see, stabilizing the laser that drives the M1 transition to a well-known frequency is critical to ensure the accuracy of subsequent measurements. Any instability or drift in laser frequency will optically pump a different velocity class of atoms into the $^3P_1$ state, potentially introducing unaccounted for errors in later E1 transition measurements. By carefully locking the M1 transition, we lay the groundwork for precise spectroscopy in the next phase of the experiment.

[8]Referring back to Fig. 2, we can write this splitting in terms of the hyperfine constant: 3853 MHz $= A + \frac{A}{2} = \frac{3A}{2}$.

Chapter 3
Locking the Infrared Laser

There are a few sources of noise that contribute to the instability of the 1279 laser. Hot air currents circulating near the furnace can create short-term, sporadic fluctuations. These can largely be suppressed by fitting end caps on the end of the cylinder extending from the furnace, letting the laser beam through while blocking unwanted air movement. The primary noise that locking will minimize are longer-term drifts arising from unpredictable mechanical and thermal fluctuations within the laser itself (Fig. 6). The 1279 nm diode operates close to room temperature (around 18°C) and has no insulating box, unlike the 368 nm ECDL which must be kept much cooler. Changes in the diode's temperature affects the bandgap of semiconductor junction which in turn affects the gain profile. As a result, even minor external thermal fluctuations can change the laser's emitted frequency significantly. For these reasons, it is important to find a way to reliably stabilize the IR laser frequency.

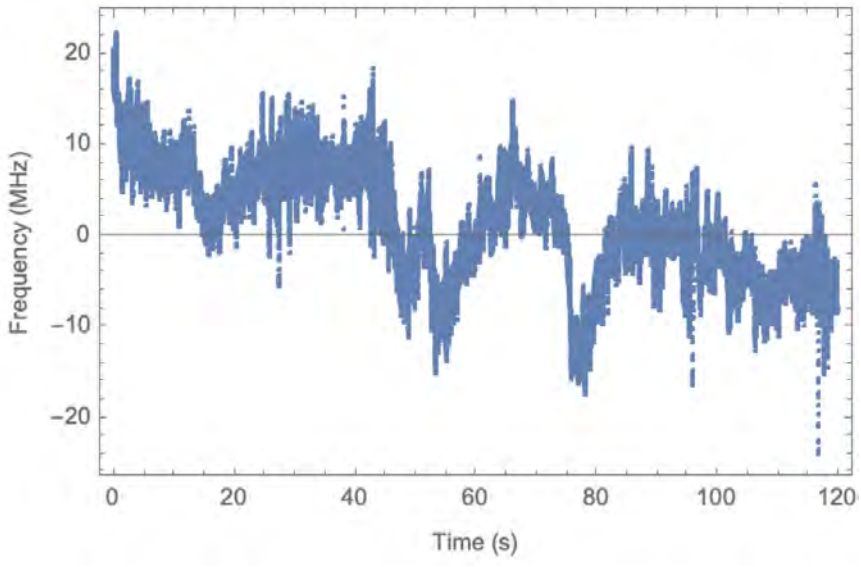

Figure 6: Frequency drifts of the 1279 nm over two minutes. On this timescale, the frequency drifts by tens of MHz, noise typical of commercial ECDLs. Quick variations come from air currents while the overall downward trend comes from thermal fluctuations in the ECDL.

3.1 General Locking Scheme

We explored two different locking schemes: one devotes a separate, smaller furnace entirely to locking while the other sends the locking beam through the larger furnace in parallel to the excitation beams. The basic set-up for both locking schemes is essentially the same. We isolate a small amount of the 1279 nm beam with a slide or a polarizing beam splitter and then pass it through a vaporized cell with a IR detector on the other side. This signal feeds into a set of electronic instruments which stabilize the laser's frequency.

3.1.1 Locking Signal

While it is possible to lock directly to the raw absorption dip, there are a few reasons why this signal is not optimal for our purposes. Since the only information we feed into our locking electronics is the strength of the input signal, if we try to lock to a maximum or minimum, they would have no indication of whether to send positive or negative current to the piezo because both would have the same effect on the input signal. This means we can only lock to regions with a nonzero slope meaning resonance is off limits. Additionally, the absorption spectrum depends on both the frequency and intensity of laser light. If, for some reason, the intensity of light that hit the detector increased, the locking electronics would output a correction signal which would change the frequency as well. To avoid these issues, we use an alternative signal that isolates frequency-dependent changes for more reliable stabilization.

Frequency-Modulation[9] Spectroscopy
Much of spectroscopy relies on taking advantage of devices to extract useful information from the signal while leaving the rest behind. Frequency-modulation spectroscopy is a powerful and well-established tool that uses the combined efforts of a function generator (or other frequency-modulating device) and a lock-in amplifier to achieve a signal that is far more useful for locking than the raw absorption dip. To do this, we send an audio-frequency ($\sim$5 kHz)[10] and low-amplitude ($\sim$0.04 Vpp[11]) sinusoidal wave to the 1279 nm piezo which dithers the laser frequency. We then demodulate the signal using a lock-in amplifier with this modulation as the reference signal, so that we are only looking at the portion of the signal that oscillates with the same phase as our high-frequency sine wave.

We can define our modulation frequency $\delta f = m\sin(\omega t)$, where m is the modulation amplitude and ω is the modulation frequency. If we Taylor expand for small modulation frequencies, keeping only the first two terms, we end up with the following expression for the transmitted intensity I as a function of laser frequency f and our modulation:

$$I(f+\delta f) = I(f) + \delta f\frac{dI}{df} + ... = I(f) + m\sin(\omega t)\frac{dI}{df} + ... \tag{13}$$

[9] I will avoid calling this "FM" spectroscopy, which has come to be associated with the bandwidth in the range of 100 MHz as a result of FM radio. Here, we are working in the "audio" frequency range (ie. the range that the human ear can sense)

[10] This frequency actually is on the lower side for frequency-modulation spectroscopy. Generally, the higher the modulation, the cleaner the signal, but we are (generally) limited to frequencies that can be produced by a function generator. More on this soon.

[11] We will later quantify this input voltage in terms of change in laser frequency.

This equation consists of a term dependent of the un-modulated intensity and a term proportional to the derivative of the intensity oscillating at a frequency of ω. Similarly, expanding out to three terms would give us a term proportional to the second derivative of the intensity oscillating at 2ω, a fourth term would give us the third derivative oscillating at 3ω, and so on. As a result, FM spectroscopy is sometimes referred to as "derivative spectroscopy." For a single symmetric absorption dip, each subsequent derivative flips in symmetry: the first derivative is asymmetric, the second is symmetric, and so on. For our purposes, we only care about the term oscillating at ω which essentially captures the slope of the absorption dip (Fig. 7). The resulting dispersive line-shape turns out to be a useful signal for locking because it has three areas with non-zero slope and a maximum slope on resonance. It is also unaffected by constant offsets in laser intensity, making it ideal for locking purposes.

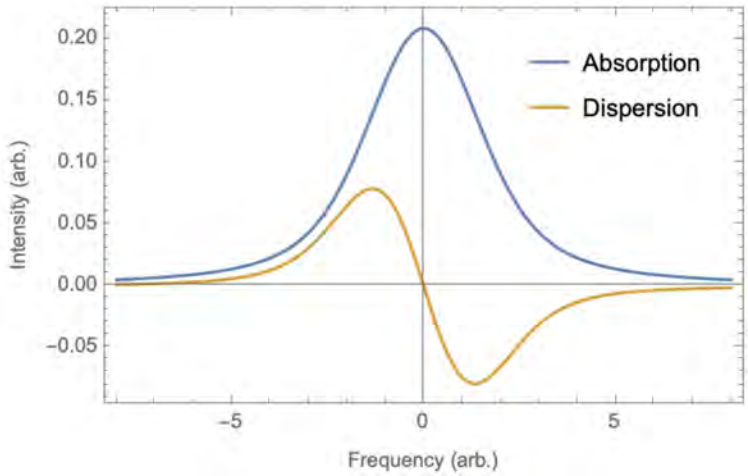

Figure 7: Model of Absorption and Dispersion Curves

3.1.2 PID Controller

To lock, we used a Proportional-Integral-Derivative device (PID), a feedback-based loop controller which takes in an error signal from the locking beam and outputs a correction signal to minimize drifts over time. The fundamental theory of the PID dates back to early automatic ship steering, which we can use as a good visual representation of how the controller operates.

Imagine we have a ship that is meant to head North. Environmental factors such as wind and sea currents can disturb our heading, tilting us slightly to the East or West. A PID can steer us back on course, using three adjustable parameters: proportional (P), integral (I), and derivative (D). The proportional term simply produces an output that is directly proportional to the magnitude of the difference between our current and desired heading. If unpredictable water currents push our heading slightly East, this component will steer us West. The integral term reflects the accumulation of past errors, correcting steady-state errors and smoothing out any small fluctuations. If the wind persistently steers us to the East, this term adds a steady correction to the West. Conversly, the derivative component estimates future fluctuations by assessing the rate of change of the error. As we get closer

to our desired heading, this term slows down the steering correction, dampening overshoots. We can combine these three terms into the following equation[12]:

$$Correction\ Signal = P[\epsilon(t) + I \int \epsilon(t)dt + D\frac{d\epsilon(t)}{dt}] + Offset \tag{14}$$

Where the error is the difference between the "setpoint" (the desired voltage which must be programmed in) and "measure" (the current voltage reading).

The PID therefore gives us robust and precise way of minimizing errors. It is currently used to solve many control problems in modern industry, regulating quantities such as temperature, current, and pressure. Many PID systems are hidden in other components of the experiment. Both laser controllers have a built-in PID to stabilize temperature to high precision, and the commercial furnaces use have automated temperature-control PIDs as well. For the purpose of locking the 1279 nm laser, we use the PID to monitor the light exiting the furnace, looking for indications of frequency drifts. The PID then sends a correction voltage to the piezo which subsequently adjusts the internal diffraction grating of the 1279 nm ECDL to keep it outputting at a constant frequency.

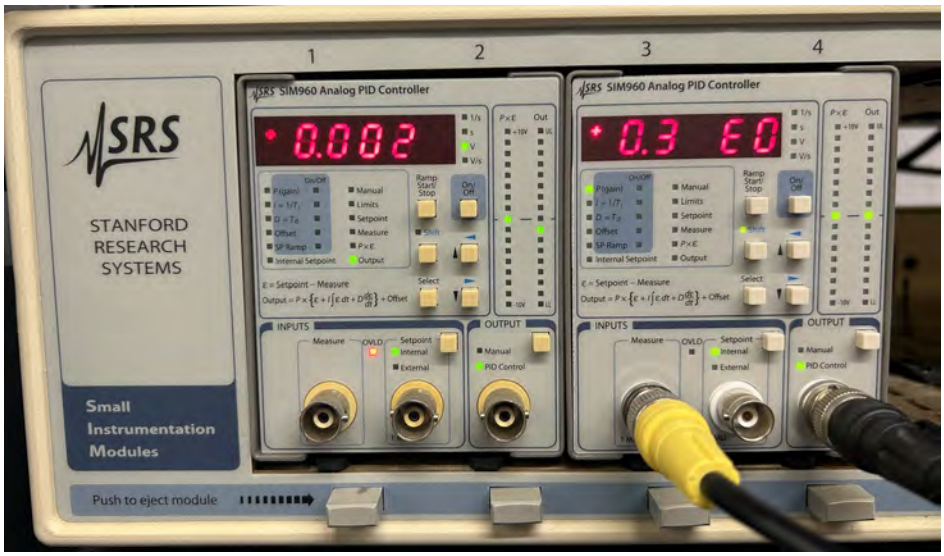

Figure 8: Photograph of Two SRS SIM960 PID Controller Modules.[13]I found it useful to monitor the error signal while adjusting parameters using the BNC outlet located on the back.

Tuning

Before choosing the parameters of the PID, manually tune the laser so that it is single-mode and sitting where you would like to lock on the absorption dip. Choose the setpoint

[12]Interestingly, if we set $x = \int \epsilon\, dt$ and shift a few terms around, we get an expression that closely resembles the damped harmonic oscillator: $Output - Offset = P\,[D\,\ddot{x}(t) + \dot{x}(t) + I\,x(t)]$. As the PID corrects for errors, it follows a similar pattern of damped oscillations until it reaches the desired value.

[13]The manual of this model can be found here: https://www.thinksrs.com/downloads/pdfs/manuals/SIM960m.pdf (or you can look for the dusty physical copy on the lab shelves).

of the PID to be about the same value as the measure quantity. If the difference between the setpoint and measure exceeds 1 V, the LED overload indicator will light up and the PID will have difficulty returning the signal to the setpoint.

Now it is time to choose the parameters. First, set the P-value as low as possible and select its polarity ($\pm$) so that the feedback is negative. Generally, positive polarity for the PID should be used for non-inverting processes (where a positive change in input corresponds to a positive change in signal) while negative polarity should be used for inverting processes (where a positive change in input corresponds to a negative change in signal). If the slope of our input signal is positive we have a non-inverting process since positive shifts in voltage induce positive shifts in frequency, and we should choose positive P-polarity. For negative slopes, we should select the negative P-polarity. Once the sign is chosen correctly, begin to increase the P-value by small amounts until you observe oscillations, and then back off until they go away. An optimal P-value will effectively reduce the error with minimal oscillation. Next, set the I-value to something small. Again, slowly increase it until you see oscillations, and then back off. An optimal I-value is large enough so that it does not overshoot the setpoint too much and small enough so that it is not too slow in responding to errors. In this way, choosing the I-value is much like choosing the optimal time constant of a lock-in amplifier or achieving critical damping in a simple harmonic oscillator. For our purposes, the D-value was not useful, since we are trying to account for random, rather than predictable, fluctuations, so we chose to leave it off. To test if the PID is locking properly, change the setpoint by something small. The signal should abruptly shift to the new value and quickly smooth out any oscillations.

3.2 Single-Furnace Locking

This scheme integrates the locking into the rest of the experiment, passing the locking beam in parallel with the excitation beams in the same vapor cell (Fig. 10). First, a polarizing beam splitter situated right before the AOM sends a small amount of the light off in the perpendicular direction. We will use this light for locking. With a few cleverly placed mirrors, we join the co-propagating and locking beams to run in parallel about a centimeter apart through the cell. This process takes some time and patience. It is easiest to begin by aligning the beams with the furnace out of the way, first getting one beam in place and then using it as a guide for the other. John and I also found it useful to lightly draw two dots a centimeter apart on a photosensitive card and loosely tape the card to the target mirror, which was far away from the mirrors we needed to adjust. Once the beams were aligned, we used an IR photodetector like normal, sending locking signal to the lock-in. Besides aligning the beam, it can be difficult to set this up without some of the AOM-chopped light hitting the photodetector. This setup, however, is more compact than using a separate locking furnace and seems to lock well if set up correctly.

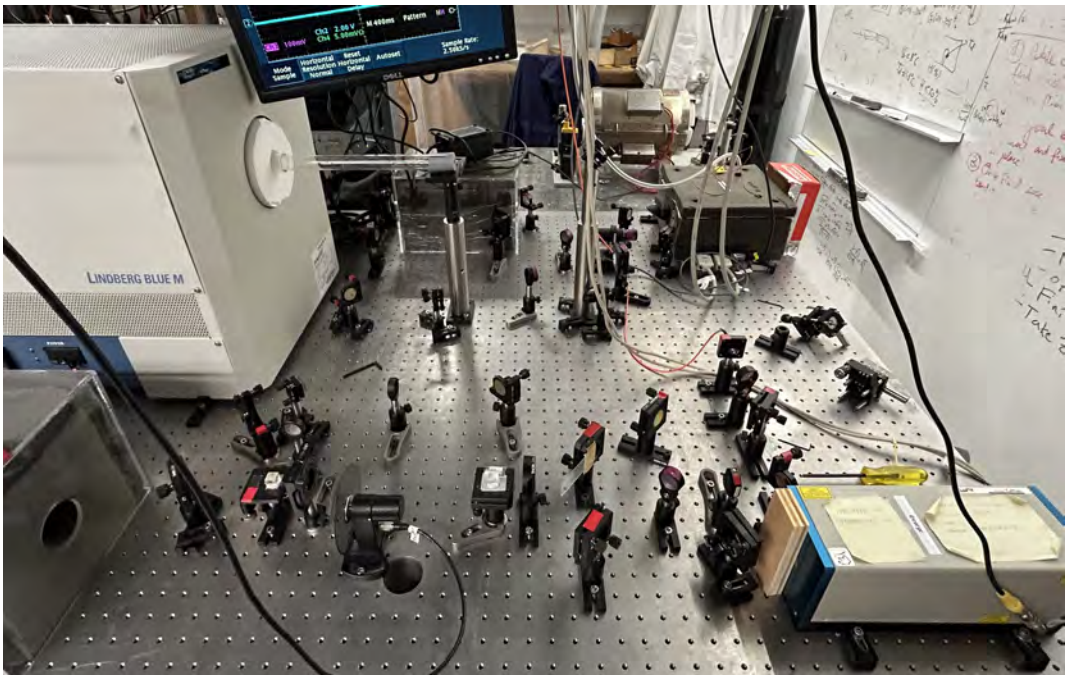

Figure 9: Photograph of the Single-Furnace Locking Setup. In this photo, we are using a chopper as a placeholder for the AOM.

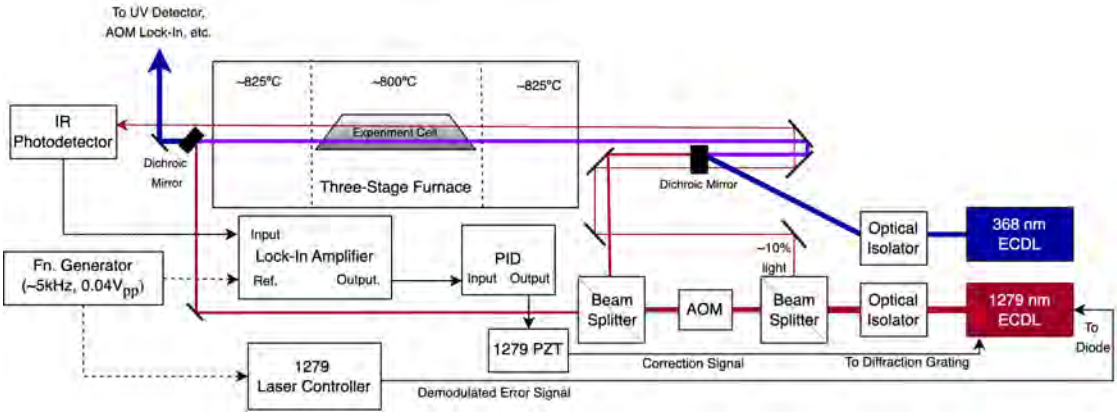

Figure 10: Single-Furnace Locking Diagram

Using this scheme, we took a number of scans to test the locking efficiency at different modulation amplitudes and laser light intensities to quantify an optimal lock. It is especially useful to know the minimum amount of light needed for a stable lock, since any remaining light can be used for the rest of the experiment[14]. We were able to lock the laser to within 1 MHz RMS noise using as little as 1.3 mW devoted to the locking beam, which is about 7% of the maximum 20 mW beam output. There also seemed to be no significant trend in locking quality as we varied modulation amplitude, so we set modulation amplitude to 0.04

[14]For all of these scans, we used the John's LabView PID program rather than the dedicated PID, though they work essentially the same.

V_{pp}, the lowest possible value on the function generator, for the remaining scans. All of these scans were done using a modulation frequency of 5 kHz.

3.3 Separate-Furnace Locking

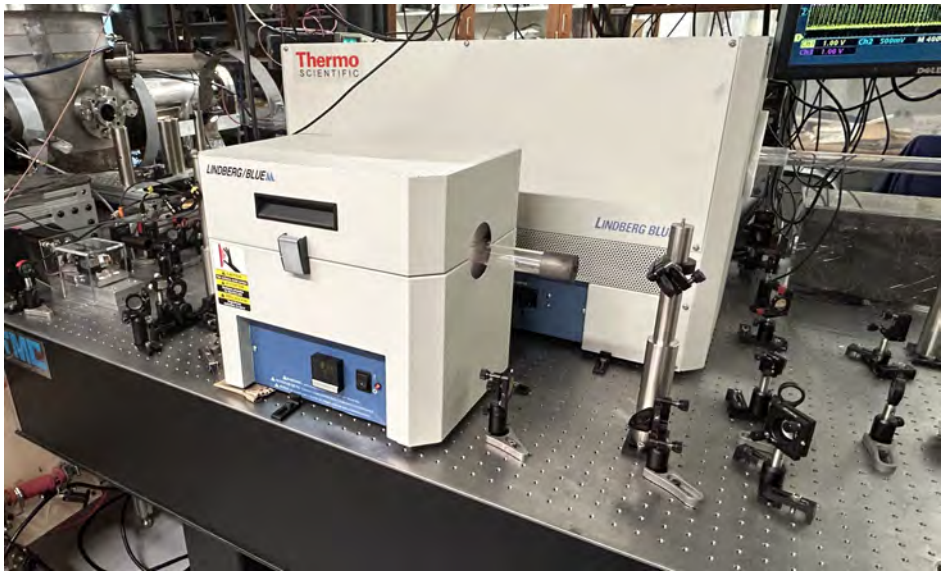

Figure 11: Photograph of Separate-Furnace Locking Setup

This setup is very similar to the one described in Robin's thesis [6], which also devoted a separate furnace to locking. Rather than using an electro-optic modulator (EOM) right before the locking furnace to modulate the locking beam, however, here we modulate the entire 1279 nm beam with a function generator. The previous EOM-locking technique was very good in that it was able to modulate only the segment of the beam dedicated to locking at a very high frequency (100 MHz), but it required use of the one (expensive) high-frequency lock-in in the lab, which is also needed to demodulate the AOM-chopped signal. Here, we implement a technique that allows us to use a function generator and any ordinary lock-in.

For this locking scheme, we use a slide to take a segment of the frequency-modulated 1279 nm beam before it reaches the AOM. The slide takes off about 8% of the beam which we then pass through a smaller tabletop furnace (Thermo Scientific Lindberg/Blue) which we dedicate entirely to locking. We place a natural abundance cell in this furnace, which gives us many locking regions with a non-zero slope13.

While setting up the smaller furnace, we noticed an odd dispersive signal and realized that the IR detector–which is sensitive to a range of wavelengths—was detecting the black-body radiation emitted by the furnace, which raised the entire signal by a constant amount. To solve this issue, we put a diffraction grating right after the furnace to pick out wavelengths in the range of the laser light and scatter the rest away from the detector. We also used end caps to minimize air current noise as we did with the larger furnace. Additionally, we were originally using different modules of the same piezo for both lasers, but cross-talk from the

scanning UV module caused an unintentional modulation in the locked IR voltage. To fix this, we simply switched the 1279 nm to the other piezo.

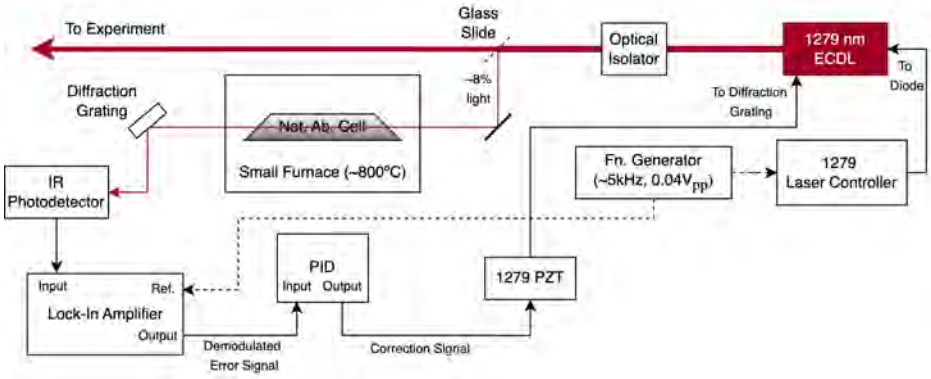

Figure 12: Separate-Furnace Locking Diagram

Using this scheme, we locked the 1279 nm laser with an RMS noise around 0.7 MHz[15], a value in the range of what other similar locking techniques have achieved. Without the Fabry-Pérot nearby, we used the well-known hyperfine splitting value of 3.583 GHz (see Fig. 13) to convert voltage to frequency, using the same technique described earlier when using the Fabry-Pérot.

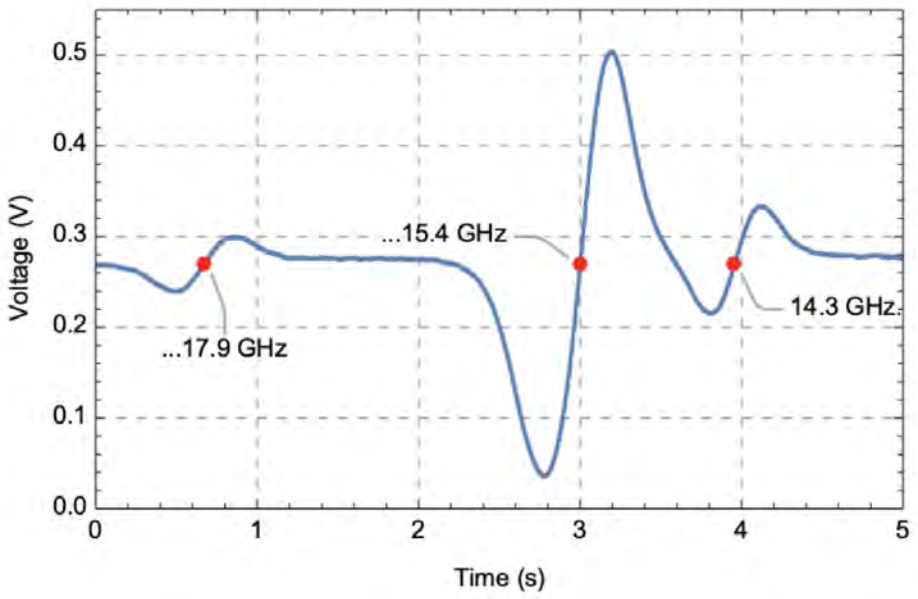

Figure 13: Natural Abundance Full Dispersion Curve. In red, we point out the last three digits of the frequency as read from the wavemeter, including the first decimal point (the second decimal place is not very reliable). The 17.9 GHz and 14.3 GHz points correspond to the 207 hyperfine splitting, with 3.583 GHz between them.

[15]This value may fluctuate depending on whether the PID settings have been properly optimized, but it generally stays in the neighborhood of 1 MHz or less.

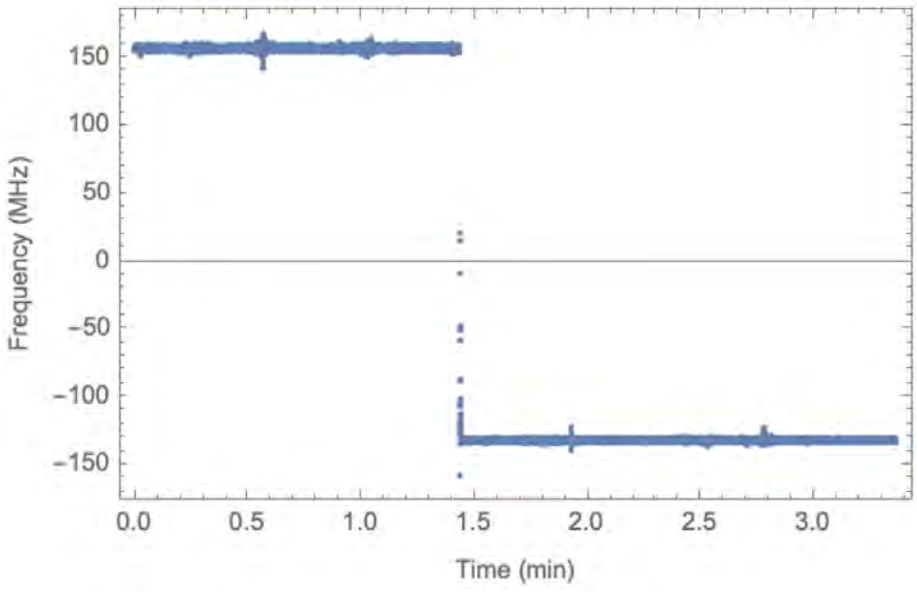

Figure 14: Demonstration of Lock with Two-Furnace Scheme. (locked to the center of larger central feature in Fig. 13). After about 1.5 minutes, the setpoint was changed by 300 MHz. The RMS noise before the setpoint was changed was about 1.1 MHz and the RMS noise after the setpoint was changed was 0.8 MHz. There is also evidence of acoustic oscillations from movement in the lab (see 0.5 min, for example) which shows us how sensitive the lock is to small disturbances.

3.3.1 A Test of Systematic Error

While the dispersion curve is a much better locking signal than the raw absorption dip, it is not perfect. It is possible that the zero-crossing of the dispersion curve is drifting over time, a systematic error that would go unaccounted for in our earlier noise assessment. To test for this drift, we took a number of upscans and downscans using the natural abundance cell, each five minutes apart over the course of 25 minutes (Fig. 13), modeling of the systematic error test done in [7]. We then averaged the top and bottom of the central feature to check for fluctuations in the zero-point. If the 1279 nm laser were locked to the center of this curve, any drifts would cause a systematic change. Using the slope of the dispersion curve, we calculated the RMS noise over this timescale to be about 1 MHz which, given that scans in the final spectroscopy experiment would only last a few seconds, reassures us of the quality of our lock.

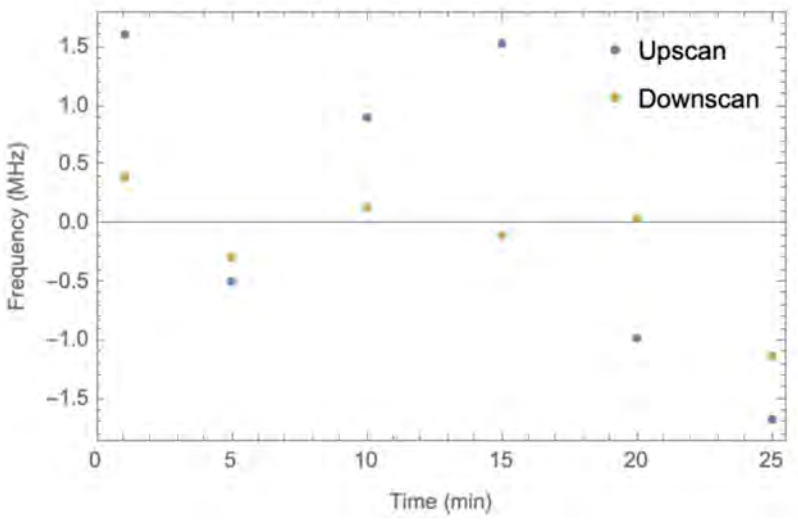

Figure 15: Longer-Term Dispersion Drift. Scans were taken 5-minutes apart over 25 minutes. At each 5-minute increment, we took both an upscan and a downscan. The zero-crossing of the dispersion curve varies slightly less for downscans than for upscans, though neither varies by significant amounts on the timescales that we care about.

3.3.2 Demonstration of Hyperfine Splitting

As a demonstration of the separate-furnace locking scheme's functionality, we locked to different features of the natural-abundance dispersion curve using an experimental cell containing ^{206}Pb and ^{207}Pb. This allows us to lock the 1279 nm laser to slightly different wavelengths and excite different velocity classes of atoms to the $^3P_1$ level. Using the high-speed lock-in to look only at the times when the excitation beam was chopped on, we observed the Doppler-narrowed lineshape and found evidence of hyperfine splitting of the 207 isotope.

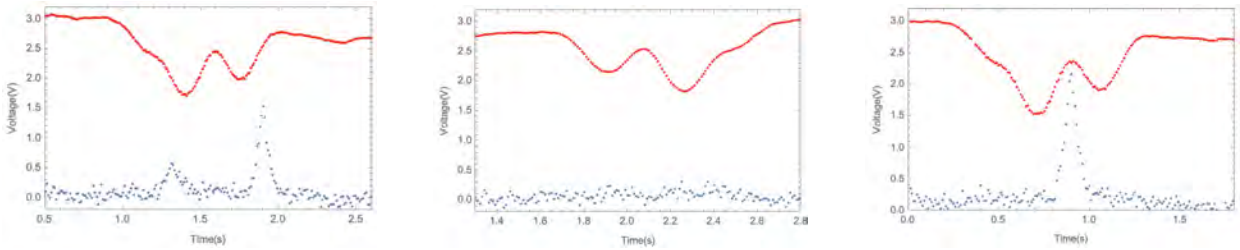

(a) ...14.7GHz (b) ...14.95GHz (c) ...15.95GHz

Figure 16: Sample Scans Locked at Different Frequencies With a 206/207 Cell. The frequency corresponds to the last three digits on the 1279 nm wavemeter to the first decimal place. Each graph shows the raw 368 nm spectrum (red) and the lock-in output (blue), using the AOM modulation as the reference signal. Fig. a demonstrates the hyperfine splitting of the 207 isotope by locking the 1279 nm laser to the right of the F = 1/2 level (refer to 13 to place the frequency on natural-abundance dispersion curve). Fig. b and c shows no evidence of hyperfine splitting when locked the left and right of the central feature, but Fig. c shows the Doppler-narrowed 206 feature.

3.4 Concluding Remarks and Next Steps

The two locking schemes presented in this report allow us to lock the 1279 nm laser within 1 MHz RMS noise using a function generator, a PID controller, and an ordinary lock-in amplifier. Both schemes reliably pump a specific velocity class of atoms to the $^3P_1$ level without relying on an EOM and high-frequency lock-in, as the lab used in earlier setups.

There are a few more steps that could improve the current locking scheme. We recently quantified the amplitude of the frequency modulation from 0.04 V_{pp} to be around 40 MHz peak-to-peak, which is far too high. John put in a variable voltage divider between the function generator and the 1279 nm laser controller, allowing us to decrease the amplitude of this modulation. A next step would be to explore locking quality using different levels of attenuation. It would also be helpful to quantify the amount of light we need to devote to locking in the separate-furnace setup. If we identified the minimum percentage of the 1279 nm beam needed for a stable lock, we could use a polarizing beam splitter (rather than a slide) to pick off that amount of light, leaving the rest for the experiment. These improvements, in addition to previous tests described in this report, lay a solid foundation for conducting measurements using the complete spectroscopy system.

Bibliography

[1] A. C. Kinney, Bachelor's thesis, Williams College (2024).

[2] C. Yang, Bachelor's thesis, Williams College (2024).

[3] D. J. Griffiths and D. F. Schroeter, Introduction to Quantum Mechanics (Cambridge University Press, 2018), 3rd ed.

[4] S. C. Doret, Rev. Sci. Instrum. 89, 023102 (2018).

[5] RP Photonics, Acousto-Optic Modulators, RP Photonics Encyclopedia (accessed 2025). https://www.rpphotonics.com/acoustoopticmodulators.html

[6] R. Wang, Bachelor's thesis, Williams College (2024).

[7] M. Gunawardena, P. W. Hess, J. Strait, and P. K. Majumder, Rev. Sci. Instrum.79, 103110 (2008).

[8] The Editors of Encyclopaedia Britannica, lead, Encyclopaedia Britannica (2025). https://www.britannica.com/science/lead-chemical-element