

Design and implementation of a narrow-line diode laser system for strontium laser cooling

Jacob VanArsdale,^{*} Toshihiko Shimasaki,[†] and David Weld[‡]
*Department of Physics, University of California,
 Santa Barbara, California 93106, USA*

(Dated: October 5, 2022)

We report a narrow-line diode laser system developed for increased efficiency of laser cooling on the 689 nm strontium 1S_0 - 3P_1 transition. This forbidden transition has a very narrow linewidth of 7.1 kHz, which requires a laser more stable than most commercial systems offer. To achieve this level of stability, we implement a Pound-Drever-Hall laser stabilization technique using an ultrastable cavity locked to a ^{84}Sr transition. We also outline practical techniques for reducing electronic and mechanical noise on a PDH apparatus. This allows for a decrease in the cooling time and an increase in our maximum Bose-Einstein condensate size for research in quantum many-body physics.

I. INTRODUCTION

The advancement of laser cooling techniques over the past 25 years has opened the door to countless ultracold atom experiments. In particular, strontium has been an element of active research in numerous fields, such as precision metrology, quantum information science, and quantum many-body phenomena [1,2,9]. Much of the interest in Sr relies on the features of the electronic level structure of alkaline-earth atoms, which make them ideal systems for the realization of Bose-Einstein condensates [3]. One of the most common cycling transitions used to cool strontium atoms is the forbidden 689 nm transition ($\Gamma/2\pi = 7.1$ kHz). This extremely narrow transition gives rise to a very low Doppler cooling limit, but also poses a technical challenge: commercial laser systems are only stable to ~ 20 kHz at best, yielding an inefficient cooling process on this transition (see Fig. 1 for a vi-

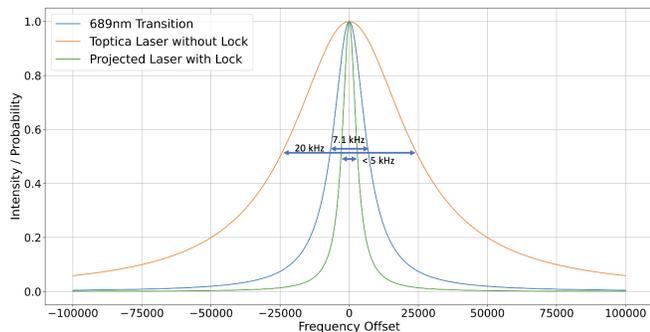


FIG. 1. A comparison of the linewidth of the 7.1 kHz transition in strontium with the stability of a Toptica DL Pro EC DL commercial laser system, and the projected reduction in linewidth with a Pound-Drever-Hall lock.

sual comparison of linewidths). Hence, a locking scheme implementing feedback is required to further stabilize a diode laser below 7.1 kHz and improve the strontium cooling efficiency.

Among the most widely used locking schemes for diode lasers is the Pound-Drever-Hall technique, in which the laser is locked to an ultrastable cavity, which itself can be locked to an atomic transition using saturation-absorption spectroscopy [4]. A key feature of this technique involves the use of an electro-optic modulator (EOM) to generate rf sidebands on the laser beam. However, all EOMs contain some degree of residual amplitude modulation (RAM), in which the birefringent crystal of the EOM acts like a Fabry-Perot etalon, causing unwanted modulation in the amplitude that can drift throughout the course of a given day [5-7]. This can be detrimental for the error signal of a laser lock, where any DC drift will cause the laser frequency to drift as well. In this paper, we detail the Pound-Drever-Hall (PDH) method with a dual temperature controlled regime (both cavity and EOM temperature stabilization) to eliminate DC drift in the error signal. In addition, we present techniques for acoustic and mechanical isolation for increased stability of the laser lock.

II. THEORY

A. The Pound-Drever-Hall Technique

A terrific feature of many diode lasers is that they are tunable, allowing the user to feed an electrical signal into the laser to control the output frequency within a few megahertz. If we desire to have a laser which has greater frequency stability than provided from the manufacturer, we can use this tunability feature to our advantage. A general procedure to do so is the following: 1. continuously measure the output frequency of the laser; 2. if the laser ever drifts down to a lower frequency, convert this information into an electrical signal to feed back into the laser to increase the output frequency; 3. if the laser

^{*} Also at Department of Physics, Colorado State University, Colorado 80521, USA; jacobvan@colostate.edu

[†] toshi@physics.ucsb.edu

[‡] weld@physics.ucsb.edu

ever drifts up to a higher frequency, convert this information into an electrical signal to feed back into the laser to decrease the output frequency. This will, in effect, hold the laser frequency approximately constant. It should be noted that the PDH technique can, in principle, compensate for noise over a wide range of frequencies and is typically ultimately limited by the electronic filters used in cleaning the error signal (which yields a lock bandwidth on the order of one megahertz). Hence, the drifts in frequency do not need to be slow in order for the laser to remain locked – fast jitters and even ‘hiccups’ due to sudden loud noises, for instance, can also be compensated, making the PDH lock a rather robust approach for laser stabilization.

In order to apply this technique, we must be able to measure the output frequency of the laser accurately. One way to do this is to couple the laser light to a Fabry-Perot cavity, and measure its transmission or reflection. A Fabry-Perot cavity acts as an ultra-narrow bandpass filter for light, only transmitting frequencies which are an integral multiple of the cavity’s free spectral range $\Delta\nu_{fsr} = c/2L$, where c is the speed of light and L is the length of the cavity. Figure 2 shows the transmission through a Fabry-Perot cavity as a function of the frequency of incident light. As the finesse of the cavity is increased, the filter becomes more narrow – the transmission peak becomes sharper. For this reason, it is ideal to have a high finesse cavity (>1000) for a PDH lock [8].

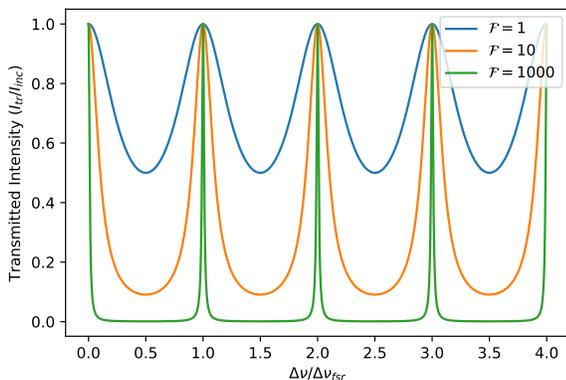


FIG. 2. Transmission through a Fabry-Perot cavity for finesse $\mathcal{F} = 1$, $\mathcal{F} = 10$, and $\mathcal{F} = 1000$. As the finesse increases, the peak becomes much more narrow.

If we measure the intensity of transmission or reflection from the cavity, changes in this intensity should correspond to a similar change in laser frequency. Since the laser is tunable, we have a choice of where on the resonance peak we want to lock. This choice is essential: if we choose to stabilize the laser at around 50% transmission and 50% reflection, (red dot in Fig. 3), then we have inadvertently coupled the frequency of the laser with its amplitude: changes in the measured intensity *might* be from changes in the laser frequency, but they also could

be from fluctuations in the intensity of the laser itself. In order to remove this issue, we can choose to measure the reflection of the laser from the cavity, and hold this to zero (blue dot in Fig. 2). Then, there is no dependence on the intensity of the laser, so any changes in reflected intensity are due to drifts in the laser frequency. However, there is one fatal flaw in this locking scheme: the reflected intensity of light is symmetric about the cavity resonance. This means that for a given measured intensity, the only information we acquire is the absolute value of the difference between the laser’s frequency and the cavity resonance – we do not know which side of the resonance we have drifted to. Instead, it turns out that the *phase* of the reflected light gives us the information to know what side of the resonance we are on – it is antisymmetric about the cavity resonance [8]. Let us diverge for a moment to provide a concrete mathematical treatment as a basis for this statement.

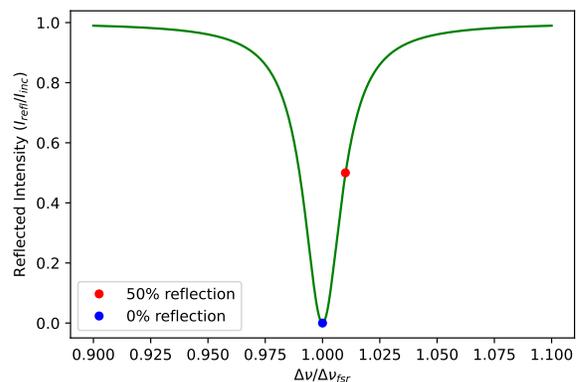


FIG. 3. Reflection from a Fabry-Perot cavity with a finesse $\mathcal{F} = 1000$. For reference, 50% reflection and 0% reflection points are identified as different locations on the resonance peak to lock the laser to.

The reflection coefficient of light from a lossless Fabry-Perot cavity is given by

$$F(\omega) = E_{\text{refl}}/E_{\text{inc}} = \frac{r(e^{i\frac{\omega}{\Delta\nu_{fsr}}} - 1)}{1 - r^2 e^{i\frac{\omega}{\Delta\nu_{fsr}}} \quad (1)$$

where the amplitude reflection coefficient is r , the frequency is ω , and the free spectral range of the cavity (length L) is $\Delta\nu_{fsr} = c/2L$ [10].

The phase of this complex-valued reflection coefficient is plotted in Figure 4. Above resonance, the phase of the reflected light is positive, and below resonance, the phase of the reflected light is negative. Since the phase of reflected light relative to the incident light provides information about how far away from resonance we are *and* which side of the cavity resonance we drift to, it makes for an ideal feedback signal to stabilize the laser.

Unfortunately, there do not exist electronics with a high enough bandwidth to directly measure the phase of

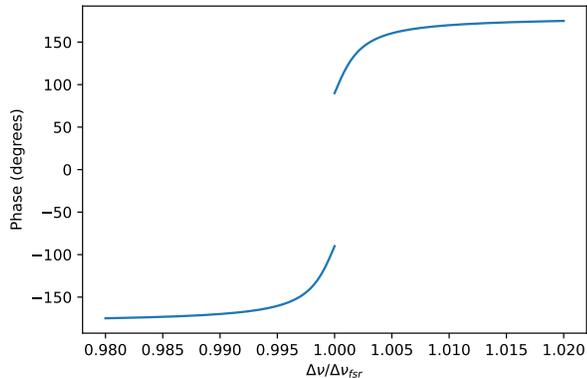


FIG. 4. Plot of the phase of the reflected light from the cavity as a function of laser frequency. The amplitude reflection coefficient of the mirrors is $r = 0.99911$. Note that there is a 180 degree discontinuity in the phase as the laser frequency crosses the cavity resonance.

visible light, so we must design another scheme to indirectly tease out this information. The approach we will focus on is known as the Pound-Drever-Hall method, in which the laser is modulated at an rf frequency (typically with an electro-optic modulator). Modulating the laser's frequency (or phase) will generate sidebands with a definite phase relationship to the incident and reflected beams. If we interfere these sidebands with the reflected beam, we get a beat-note at the modulation frequency, and we can measure the phase of this beat-note, since it is well within the capabilities of modern electronics. The phase of this beat-note will then tell us the phase of the reflected beam.

B. Creating the Error Signal: Measuring the Phase of Reflected Light

Let us consider the measurement of the phase in mathematical detail. An optical carrier at frequency ω which is phase modulated at frequency Ω , with modulation index β is given by

$$E_{\text{inc}} = E_0 e^{i(\omega t + \beta \sin(\Omega t))} \quad (2)$$

Here, the modulation index is the ratio of the power in each of the sidebands to the power of the carrier [9,10].

The following Bessel function expansion will help to elucidate the frequency components of the wave:

$$e^{i\beta \sin \theta} = \sum_{m=-\infty}^{\infty} J_m(\beta) e^{im\theta} \quad (3)$$

Using this expansion as an approximation to first order, the field becomes

$$\begin{aligned} E_{\text{inc}} &\approx [J_0(\beta) + 2iJ_1(\beta) \sin(\Omega t)] e^{i\omega t} \\ &= E_0 [J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t}] \end{aligned} \quad (4)$$

This is an excellent approximation as long as the modulation index β is small. As can be seen from Eq. (4), we have three beams emerging from the electro-optic modulator: a carrier with frequency ω , and two sidebands, with frequencies $\omega + \Omega$ and $\omega - \Omega$. In order to find the electric field of the light reflected from the cavity, we simply must multiply each of these components by the reflection transfer function F from Eq. (1):

$$\begin{aligned} E_{\text{refl}} &= E_0 [F(\omega) J_0(\beta) e^{i\omega t} + F(\omega + \Omega) J_1(\beta) e^{i(\omega+\Omega)t} \\ &\quad - F(\omega - \Omega) J_1(\beta) e^{i(\omega-\Omega)t}] \end{aligned} \quad (5)$$

The photodiodes used to detect this light actually measure the reflected power $P_{\text{refl}} = |E_{\text{refl}}|^2$, however. If we let $P_0 = |E_0|^2$ be the total power in the incident beam, then the power in the carrier is $P_c = J_0^2(\beta) P_0$ and the power in each of the sidebands is $P_s = J_1^2(\beta) P_0$. For small modulation depth β , the majority of the power is in the carrier and first order sidebands, such that $P_c + 2P_s \approx P_0$.

After some algebraic simplification, we find

$$\begin{aligned} P_{\text{refl}} &= P_c |F(\omega)|^2 + P_s \{ |F(\omega + \Omega)|^2 + |F(\omega - \Omega)|^2 \} \\ &\quad + 2\sqrt{P_c P_s} \{ \text{Re}[F(\omega) F^*(\omega + \Omega) \\ &\quad - F^*(\omega) F(\omega - \Omega)] \cos(\Omega t) + \text{Im}[F(\omega) F^*(\omega + \Omega) \\ &\quad - F^*(\omega) F(\omega - \Omega)] \sin(\Omega t) \} + (2\Omega \text{ terms}). \end{aligned} \quad (6)$$

The Ω terms arise from the interference between the carrier and the sidebands, and the 2Ω terms arise from the sidebands interfering with each other. The phase information we are after is contained within the Ω terms. In order to recover the phase, we need to demodulate this signal. In practice, this is accomplished with a mixer, a device that multiplies two electrical signals together (typically a double-balanced mixer is used to reduce leakage). Recall that the product of two sine waves, with one offset by an arbitrary phase ϕ , yields the superposition of a wave with the sum of the arguments and a wave with the difference of the arguments:

$$\sin(\omega t + \phi) \sin(\omega' t) = \frac{1}{2} [\cos((\omega - \omega')t + \phi) - \cos((\omega + \omega')t + \phi)] \quad (7)$$

If the waves have the same frequency ($\omega = \omega'$), this simplifies to

$$\sin(\omega t + \phi) \sin(\omega t) = \frac{1}{2} \cos(\phi) - \frac{1}{2} \cos(2\omega t + \phi) \quad (8)$$

Therefore, if we mix the reflected light with the modulating signal, and then use a lowpass filter to remove the 2ω component, we are left with a term that only depends on ϕ ; we have successfully measured the phase of

the reflected light. More specifically, in our case we use a sine wave with amplitude A from a signal generator with arbitrary (but constant) phase ϕ , and mix it with the photodetected power in Eq. (6):

$$\begin{aligned}\epsilon &= C_1 \sin(\Omega t + \phi) + 2C_2 \cos(\Omega t) \sin(\Omega t + \phi) \\ &\quad + 2C_3 \sin(\Omega t) \sin(\Omega t + \phi) \\ &= C_1 \sin(\Omega t + \phi) + C_2 [\sin(2\Omega t + \phi) + \sin(\phi)] \\ &\quad + C_3 [\cos(\phi) - \cos(2\Omega t + \phi)]\end{aligned}\quad (9)$$

where

$$\begin{aligned}C_1 &= A(P_c |F(\omega)|^2 + P_s \{|F(\omega + \Omega)|^2 + |F(\omega - \Omega)|^2\}), \\ C_2 &= A\sqrt{P_c P_s} \text{Re}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)], \\ C_3 &= A\sqrt{P_c P_s} \text{Im}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)].\end{aligned}$$

Low pass filtering removes the Ωt and $2\Omega t$ terms, so that we are left with

$$\begin{aligned}\epsilon &= C_2 \sin(\phi) + C_3 \cos(\phi) \\ &= A\sqrt{P_c P_s} \{\sin(\phi) \text{Re}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)] \\ &\quad + \cos(\phi) \text{Im}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)]\}\end{aligned}\quad (10)$$

Since A , P_c , P_s , ϕ , and Ω are all constants, we are left with an electrical error signal ϵ that only depends on fluctuations in the laser frequency ω . In practice, we need to also be able to manually adjust the relative phase ϕ between the modulation signal and photodetected signal to correct for unequal delays in the two signal paths. When the phases of these signals are not matched, there can be significant distortion of the error signal (see Fig. 5). When setting up a Pound–Drever–Hall lock, you can scan the laser frequency or the cavity length and empirically adjust the phase in one signal path until you get an error signal that looks like the $\phi = 0^\circ$ plot in Figure 5.

C. Residual Amplitude Modulation in EOMs

Due to the fact that the manufacturing process of birefringent crystals is not perfect, the phase modulation from electro-optic modulators can stray from the mathematical model. Imperfections on the crystal edges can cause some light to be reflected, leading to an effective etalon in which we have interference of reflected beams with transmitted beams [5-7]. Figure 6 shows this effect. The interference shows up as amplitude modulation on the photodetected light, which can severely degrade the quality of the PDH lock – and hence, the stability of the laser can be compromised. There are a few methods which have been implemented by others to combat this ‘residual amplitude modulation’ (RAM), including purposeful misalignment of the laser beam with the EOM transmission axis [6-7], careful polarization coupling with the EOM, active DC feedback to rotate the modulator axis [5], and temperature control of the EOM [6].

The simplest passive method to reduce RAM is to slightly misalign the laser with the EOM crystal, as

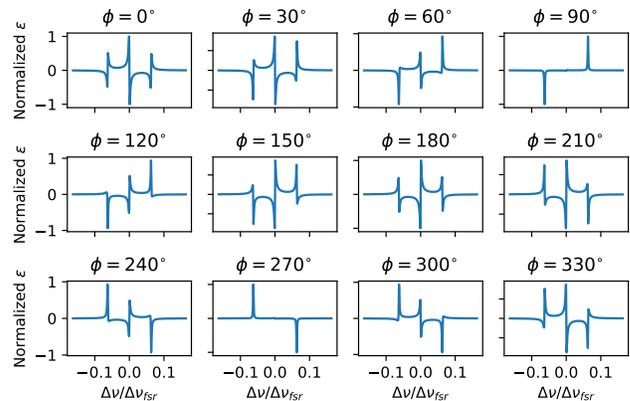


FIG. 5. Plots of the PDH error signal when scanning either the cavity or the laser, for different phases ϕ between the mixer inputs. The modulation frequency is $\Omega = 20$ MHz and the free spectral range is $\nu_{\text{fsr}} = 1.5$ GHz. The ideal signal that should be used for locking has $\phi = 0^\circ$, since it has the largest gradient near resonance, yielding a more effective linear region to correct for frequency drifts.

shown in Figure 7. Then, light reflecting off of the ends of the crystal will not be in the same beam path as the main beam, reducing the interference that leads to unwanted amplitude modulation.

An active method to remove residual amplitude modulation is to provide DC feedback to the crystal, effectively rotating the polarization axis to minimize the RAM. One can photodetect the output of the EOM and lock to the derivative of this signal. However, because crystals used for EOMs typically require high rf voltages (>100 V), it is common to use a resonant EOM in the locking scheme to remove the requirement of a high-voltage rf driver. (Resonant EOMs are simply electro-optic modulators with a resonant transformer inside, making them only operable at one rf frequency). Since a DC voltage cannot be applied to a transformer, this means that DC feedback is unfortunately not an option for users with resonant EOMs.

Birefringent crystals are also sensitive to temperature fluctuations. They can expand or contract as the temperature changes, which likewise changes the optical path length for light propagating through the crystal. This can cause a slow DC drift in the error signal throughout the course of a given day. As such, it can be useful to control the temperature of the EOM with a thermoelectric cooler (TEC), in tandem with a PID controller. More information on control theory and a guide on tuning a PID controller (to select the proportional, integral, and derivative gain coefficients) with the Zeigler-Nichols method is given in Appendix A.

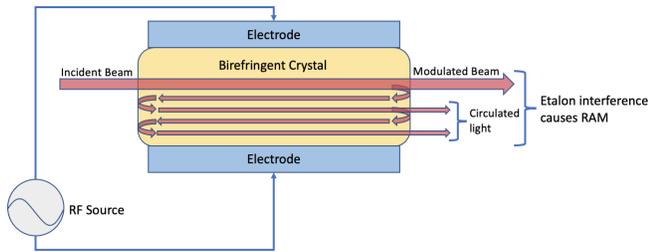


FIG. 6. Diagram of the etalon effect in an imperfect crystal, leading to residual amplitude modulation in the transmitted light from an electro-optic modulator.

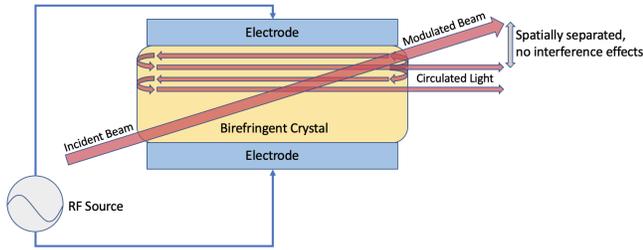


FIG. 7. Diagram of purposeful misalignment of laser light with crystal in order to reduce etalon effects.

III. EXPERIMENTAL SETUP

A. Laser System for Pound-Drever-Hall Lock

The typical setup for the Pound-Drever-Hall method is given in Figure 8. A 689 nm Toptica DL Pro ECDL (external cavity diode laser) is set incident on a resonant electro-optic phase modulator (Thorlabs EO-PM-R-20-C1), where it is phase modulated by a signal generator (Siglent 2210X) at an rf frequency of 20 MHz. The carrier (with sidebands) is then coupled to a high-finesse Fabry-Perot cavity ($\mathcal{F} \approx 3000$, fused silica plano-concave Layertec 102468 mirrors, 10 cm fused quartz spacer), such that the TEM₀₀ mode is maximized. For information detailing the process of coupling light to a cavity, see Appendix B. The cavity is placed in an ultra-high vacuum chamber to provide further isolation from its environment, and is temperature stabilized with a resistive tape and an Arroyo 5300 controller. The reflected light is photodetected with a high-bandwidth photodiode (Thorlabs PDA100A), where it is filtered with a 20 MHz bandpass filter (Mini-Circuits BBP-20R5+), attenuated, and amplified with an rf amplifier (Mini-Circuits ZFL-500LN+). Here, it is mixed with the modulating signal using a doubly-balanced mixer (Mini-Circuits ZRPD-1+), whose relative phase can be controlled with the rf generator. (The Siglent 2210X has two output channels which can be internally locked with a phase-locked-loop. One output channel is sent to the EOM for phase modulation, and the other is sent the mixer for demodulation). This feature is essential because the phase can significantly

distort the error signal, as detailed in the Theory section (see Fig. 4). The output of the mixer is then lowpass filtered at 1 MHz (Thorlabs EF508), and a PID controller servo adjusts the gain of this error signal to achieve maximal stability. The signal is then fed back into the Toptica laser to control its output frequency.

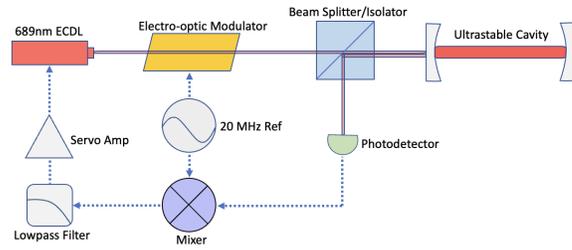


FIG. 8. Diagram of the main components for the PDH lock.

In order to provide active stabilization of the cavity itself (which contains a piezo-electric transducer for cavity length adjustment), we utilize a technique known as saturation-absorption spectroscopy. Here, we lock to a ^{84}Sr transition as an absolute frequency reference. More information on this scheme can be found in [9].

B. Electronic, Acoustic and Mechanical Isolation Considerations

There are, of course, many other sources of noise that impede on a clean error signal. The most prominent types can be broken up into three categories: electronic noise, acoustic noise, and mechanical/vibrational noise. The latter two can be taken care of with damping material (we use a 0.5" polyurethane sheet), placed between the PDH setup and the optical table it rests on. In particular, we place the cavity and optical components on a small optical breadboard (Thorlabs Nexus B1824F), and set this breadboard on the polyurethane sheet. Additionally, one can build an enclosure for the lock with damping material inside, such as high-density acoustic foam, to limit noise caused by people talking or components being moved around in the lab – or even 120 Hz flicker from fluorescent overhead light.

The electronic noise can be much more subtle and challenging to deal with. The best way we have found to remove this noise is through the use of an isolation transformer. This is a 1:1 transformer with grounded Faraday shields between the windings to remove high frequency noise on the power source. It is particularly useful if there is a high-current power supply nearby, which can create severe 60 Hz noise on ground loops that otherwise cause horrific problems for the Pound-Drever-Hall error signal. It is important to look for all possible locations for ground loops to exist within the setup. For instance, some possible locations are in the power cords of photodiodes, the high voltage power of a vacuum ion pump

for the Fabry-Perot cavity, or even the USB connector on a CMOS camera. We have experienced all of these as sources of noise in practice.

The error signal can be viewed by scanning the cavity length with an internal piezo-electric transducer over a small range, using a triangle wave. Assuming the change in the cavity length is approximately directly proportional to changes in applied voltage on the piezo, viewing the error signal on an oscilloscope will show the same signal shape that we would see in frequency space. In Fig. 9 our actual error signal is pictured, before and after electronic, acoustic, and mechanical noise reduction is employed.

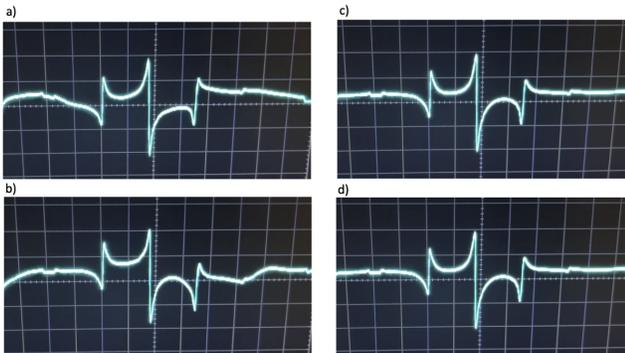


FIG. 9. (a) and (b) show the PDH error signal before noise correction, with (b) taken a few seconds after (a), to show the timescale of the noise. (c) and (d) show the error signal after noise correction, also taken a few seconds apart. There is no longer a visible change on short time scales after the electrical and mechanical isolation is implemented.

IV. CONCLUSION

In this paper, we have detailed a method to optimize the Pound-Drever-Hall (PDH) method with temperature control of both the cavity and EOM to eliminate DC drift in the error signal. This lock is used for laser cooling on the strontium 689 nm intercombination transition (colloquially called the ‘red MOT transition’). In addition, we highlighted techniques for electronic and mechanical isolation of the PDH system for increased stability. This lock should allow for a reduced cooling time on strontium atoms and increase the number of atoms we can cool to a Bose-Einstein condensate, yielding an improved signal-to-noise ratio on all future experiments with the strontium machine in the Weld lab.

V. ACKNOWLEDGEMENTS

I would like to thank my postdoctoral advisor Toshi Shimasaki for all his advice throughout the summer, teaching me countless ‘tricks of the trade’ and instilling

the wise words in me: only change one thing at a time. I would also like to thank my advisor David Weld for his kind support, providing motivation and creative ideas for this project, including the (unfortunately unsuccessful) RF choke trick (which still stands at an impressive record of 2-101). Thanks go to all of the incredible members of the Weld Lab, including Yifei Bai for dealing with my plethora of single-item orders while maintaining the happiest attitude I have ever seen, to Jeremy Tanlimco for his helpful thermodynamics knowledge introducing me to the tastiest food on earth: madeleines, to Roshan Sajjad for his insistence to help out in any situation and ensure we never under-utilize the chop saw, to Quinn Simmons for his computer expertise and his excellent taste in music, and to Peter Dotti, who taught me how to be a fantastic kindergarten leader! Finally, I would like to thank UCSB REU site director Sathya Guruswamy, for her endless advice and support, ensuring everything ran smoothly throughout the summer. This work was generously supported by NSF REU grant PHY-1852574.

Appendix A: Feedback and Control Theory

There are many situations in which it is desired to control or stabilize a parameter of a dynamical system. The canonical example is the thermostat, used to control the temperature in a room. One way in which one might consider accomplishing this task is to apply a simple ‘on-off’ feedback mechanism. This can be described as follows:

$$u = \begin{cases} u_{\max}, & \text{if } e > 0 \\ u_{\min}, & \text{if } e < 0 \end{cases} \quad (\text{A1})$$

where the error $e = s - c$ is the difference between the set point s and the current value c , and u is the actuation command (i.e. turning on/off an air conditioner or heater) [11]. This system will generally oscillate about the set point, and is only an acceptable option if the oscillation is sufficiently small. For many systems, however, something more stable is required which can eliminate this oscillation. In temperature control of the Fabry-Perot cavity in the PDH lock, for instance, this is essential: the fused quartz spacer used has a coefficient of thermal expansion of $5.5 \times 10^{-7}/^\circ\text{C}$, which corresponds to an 825 kHz change in the frequency of the lowest order supported cavity mode for a 1°C change in temperature. Of course, this will be detrimental when we are looking to stabilize a laser to below 5 kHz. To control a dynamical system where stability is paramount, the most common method is to use proportional-integral-derivative (PID) feedback.

The first term in PID control is proportional feedback, which takes the form $K_p e$, where K_p is the gain of this term. The further away the current value is from the set point, the larger this feedback term will be, and as the current value nears the set point, this feedback will become smaller. Proportional feedback offers a vast improvement over the on-off method, but still suffers the

drawback that it has a propensity for ‘steady-state error’, where the system will stabilize to a point below or above the set point. This is common when there are external influences consistently driving the system in a certain direction; for instance, when setting the temperature of an electro-optic modulator to a few degrees above room temperature, if it is not well insulated, heat will constantly be dissipated as the room attempts to bring the EOM to thermal equilibrium. This can cause the controller to stabilize below the set point, never reaching the desired temperature. To counteract this issue, we can implement the second term: integral feedback, which takes the form $K_i \int_0^t e(\tau) d\tau$. In practice, this is essentially treated like a running sum which adds the error terms over the full run time. When utilizing integral feedback, the steady-state error will always be zero (the longer the control variable remains below the set point, the larger this feedback term will become).

A third parameter that can optionally be added is a derivative term, which is of the form $K_d \frac{de}{dt}$. Derivative feedback is used to predict future error and help the controller remain closer to the set point when disturbances enter the system. This can be seen clearly in the first order linear extrapolation of the error to a time t_d in the future: $e(t + t_d) \approx e(t) + K_d \frac{de}{dt}$. In total, then, we have the following control function:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de}{dt} \quad (\text{A2})$$

In essence, the PID controller implements a feedback function $u(t)$ which uses as much information on the time domain as possible: a proportional term which incorporates information about the present, an integral term which incorporates information about the past, and a derivative term which incorporates information about the future.

1. PID Tuning with the Zeigler-Nichols Method

As the user of a PID controller, the primary task is to optimize the gain coefficients K_p , K_i , and K_d for the given system at hand. Many people have developed methods to select values for these coefficients, but the most popular (possibly for its simplicity) is known as the Zeigler-Nichols method. It should be noted that this is not always the most robust method, as it does not use a large amount of process information, but it is sufficient in many circumstances, at the very least to give the user a starting point from which small manual adjustments can be made. The procedure is as follows:

1. Allow the process to reach a steady state without feedback.
2. Create a ‘step’ function by increasing the input by a suitable (constant) amount.

TABLE I. Zeigler-Nichols coefficients for proportional, integral, and derivative gain of a PID controller [12].

Type	K_p	K_i	K_d
P	$1/a$		
PI	$0.9/a$	$0.3K_p/\tau$	
PID	$1.2/a$	$0.5K_p/\tau$	$0.5K_p\tau$

3. Measure the output (control variable) as it comes to a new steady state. In conjunction with Fig. 10, determine the values of τ and a by finding the x and y intercepts, respectively, of the steepest tangent to the curve.
4. Use your values for τ and a in Table 1 to determine the optimal gain coefficients K_p , K_i , and K_d for your system.

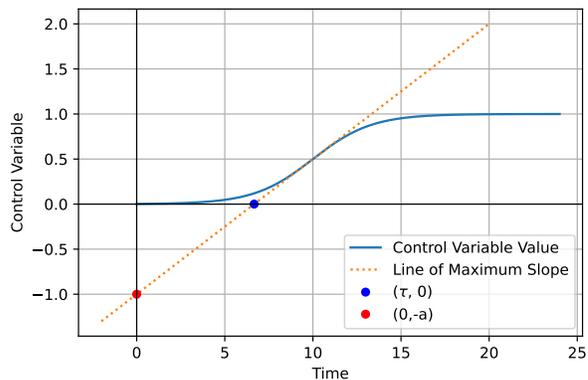


FIG. 10. Zeigler-Nichols tuning method. Take the steepest tangent of the unit step response curve: the x -intercept is τ and the y -intercept is $-a$.

Appendix B: Coupling Light to a Fabry-Perot Cavity

Coupling light to a Fabry-Perot cavity may seem challenging at first, but by taking a systematic approach, the process can be made efficient and rather painless. Before beginning to couple, it is highly recommended to place the lens used for cavity mode matching in the beam path, since it can deflect the beam if inserted after all your hard work has already been done. Additionally, if possible, use linearly polarized light into the cavity. This can be accomplished with a half-wave plate and quarter-wave plate. If while coupling, you find yourself unable to find the circular modes, try rotating the linear polarization. This can help improve the transmission, optimizing the polarization for your particular cavity. A recommended procedure to couple to the TEM₀₀ mode is as follows:

1. Scan the cavity with the piezo-electric transducer. Use a large triangle wave ($\sim 5 V_{pp}$) applied to the

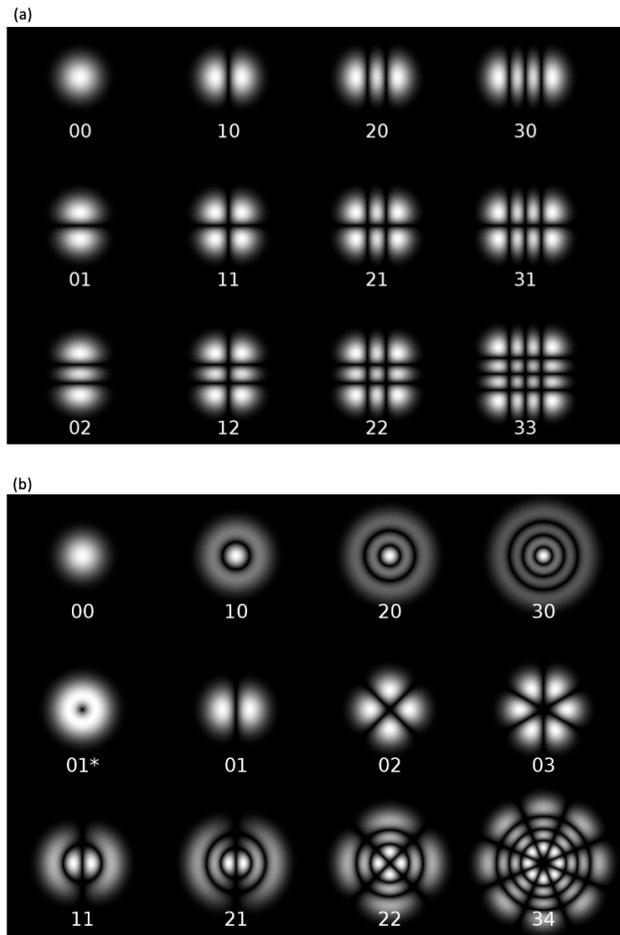


FIG. 11. (a) Rectangular TEM cavity modes. (b) Cylindrical TEM cavity modes. (Used from [13]).

piezo driver to scan over as much of a free spectral range as possible.

2. Set up a CMOS camera (such as a Basler ac1920-25um) on the output of the cavity to be able to see transmitted light. Use a long exposure time ($>200000 \mu s$) and a high gain (>12 dB) since the light will initially be very faint.
3. Walk the beam with two mirrors until you are able to see light on the CMOS camera. Once you see fringes, you are very close: these are higher order cavity modes.
4. At this point, your goal is to find circular/ring modes. These are cylindrical TEM₁₀, TEM₂₀, TEM₃₀, etc. Most likely, your fringes are in a straight line (like a rectangular TEM₀₈ mode, for instance). See Fig. 11 for a reference of what the different modes look like. Experiment with the mirrors to see which nob(s) bring the fringes closer together, towards their center. Continue to bring them closer together until you see circular modes.
5. Now that you have circular modes, decrease the scanning range (i.e., the voltage of the triangle wave) to ~ 20 mV_{pp}. Manually adjust the piezo driver voltage slowly as you scan across the full free spectral range. Continue until you find the TEM₀₀ mode.
6. Increase the scanning range to ~ 200 mV_{pp}. Replace the CMOS camera with a photodiode (we use the PDA100A set at 70dB gain), connected to an oscilloscope. It can be useful to split the signal used for scanning the cavity, and send it to the oscilloscope to use for triggering. Make small adjustments to the mirrors to maximize the peak height on the scope.
7. Give yourself a pat on the back and crack a big smile. You have successfully coupled the laser light to the cavity!

-
- [1] C. Adams and E. Riis, Laser cooling and trapping of neutral atoms, *Progress in quantum electronics* **21**, 1 (1997).
 - [2] C. E. Wieman, D. E. Pritchard, and D. J. Wineland, Atom cooling, trapping, and quantum manipulation, *Reviews of Modern Physics* **71**, S253 (1999).
 - [3] F. Sorrentino, G. Ferrari, N. Poli, R. Drullinger, and G. Tino, A strontium sample for ultracold atomic physics, high-precision spectroscopy and quantum sensors, in *EUROPHYSICS CONFERENCE ABSTRACTS ECA*, Vol. 30 (Citeseer, 2006) p. 35.
 - [4] R. Drever, J. L. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, Laser phase and frequency stabilization using an optical resonator, *Applied Physics B* **31**, 97 (1983).
 - [5] W. Zhang, M. Martin, C. Benko, J. Hall, J. Ye, C. Hagemann, T. Legero, U. Sterr, F. Riehle, G. Cole, *et al.*, Reduction of residual amplitude modulation to 1×10^{-6} for frequency modulation and laser stabilization, *Optics letters* **39**, 1980 (2014).
 - [6] Q. A. Duong, T. D. Nguyen, T. T. Vu, M. Higuchi, D. Wei, and M. Aketagawa, Suppression of residual amplitude modulation appeared in commercial electro-optic modulator to improve iodine-frequency-stabilized laser diode using frequency modulation spectroscopy, *Journal of the European Optical Society-Rapid Publications* **14**, 1 (2018).
 - [7] A. E. Domínguez, W. E. O. Larcher, and C. N. Kozameh, Fundamental residual amplitude modulation in electro-optic modulators, *arXiv preprint arXiv:1710.10719* (2017).
 - [8] E. D. Black, An introduction to pound-drever-hall laser frequency stabilization, *American journal of physics* **69**,

- 79 (2001).
- [9] Y. Li, T. Ido, T. Eichler, and H. Katori, Narrow-line diode laser system for laser cooling of strontium atoms on the intercombination transition, *Applied Physics B* **78**, 315 (2004).
- [10] W. Nagourney, *Quantum electronics for atomic physics and telecommunication* (OUP Oxford, 2014).
- [11] K. J. Åström, *Introduction to stochastic control theory* (Courier Corporation, 2012).
- [12] J. G. Ziegler, N. B. Nichols, *et al.*, Optimum settings for automatic controllers, *trans. ASME* **64** (1942).
- [13] Dr. Bob, Transverse modes — Wikipedia, the free encyclopedia (2021), [Online; accessed 11-September-2021].