

Ph 77 ADVANCED PHYSICS LABORATORY

– ATOMIC AND OPTICAL PHYSICS –

Expt. 72 – Laser Frequency Stabilization

I. BACKGROUND

In many precision optical measurements, it is desirable to have a laser with a well-defined frequency. For example, atomic physics experiments often require lasers with frequencies fixed at or near the frequencies of atomic resonance lines. For tunable lasers, it is therefore necessary to have a means of controlling the laser's operating frequency, stabilizing it and often "locking" it at a desired value. In this lab we will investigate methods for achieving this, and describe one particularly powerful technique known as the Pound-Drever-Hall method (named after R. V. Pound, who first used the technique to stabilize microwave oscillators in the 1940's, and Caltech professor R. W. P. Drever, who extended the ideas to the optical domain in the early 1980's, and John Hall, who added important pieces of RF technology) (Drever *et al.* 1983).

Laser frequency locking is actually quite a complex topic, with a number of interesting subtleties. For example, one can both "frequency lock" and "phase lock" a laser. The former, frequency locking, means that the laser's average frequency is fixed, but its linewidth remains equal to the laser's intrinsic linewidth (for our diode lasers the linewidth is around 1 MHz). With phase-locking techniques, one can take a laser, such as one of our diode lasers, and precisely control the phase ϕ of the electric field, keeping $\phi(t)$ close to $\omega_0 t$, where ω_0 is the laser's desired angular frequency. If the phase is precisely controlled, this serves to *reduce* the laser's linewidth. Laser linewidths have been reduced to sub-Hz levels (so $\Delta\omega/\omega < 10^{-15}$) for periods of several hours using these techniques. In this lab we will not worry about phase locking, and look only at frequency locking. (An introduction to phase locking can be found in Hall and Zhu 1992.)

The topic of laser frequency locking is itself embedded in the much larger field of control theory, which is a branch of engineering that deals with methods for controlling physical systems. You may have heard some of the terminology in control theory, which includes words like "servosystems" and "servomechanisms," "feedback control systems," "feedback loops," and the like. The generic servosystem contains several parts: 1) the *system*, or *plant*, which is something we want to control; 2) an *actuator*, which changes the state of the system; 3) a *sensor*, which detects the state of the system, (usually relative to some *reference*), and 4) a *controller*. The basic idea is that the controller reads the sensor, then drives the actuator in such a way to send the system to the desired state. If done well, the system will go to the desired state quickly, and then deviate from the desired state as little as possible. For the problem at hand, laser stabilization, the system will be the laser (which has an output frequency we want to control), the actuator is the laser PZT (which changes the laser's frequency), the sensor is some frequency-selective optical element (an atomic resonance cell, or an optical cavity), and the controller is some electronics box that connects the sensor signal with the actuator via a *control loop*.

Side Locking. One of the simplest control techniques is the side-locking method. One starts with some

frequency-selective optical element that produces a voltage signal, $V(\omega)$, as a function of laser frequency. If one wishes to lock the laser frequency at a frequency ω_0 , and $dV/d\omega(\omega_0) \neq 0$, then one subtracts a reference voltage to make an *error signal* $\varepsilon(\omega) = V(\omega) - V(\omega_0)$. This error signal then serves as input to a feedback loop, which adjusts the laser's frequency to make $\varepsilon = 0$. The side-locking method is useful if one wishes, for example, to lock the laser to the side of a peaked resonance feature. It only works, however, if $dV/d\omega(\omega_0) \neq 0$, which means one can lock to the side of a resonance feature, but not to the peak – hence, the name, side-locking.

One often wants to lock the laser to the *peak* of the resonance feature, so what then? The side-locking method would not work in this case, since $dV/d\omega(\omega_0) = 0$; a non-zero error signal $\varepsilon(\omega) = V(\omega) - V(\omega_0)$ would tell you that the laser frequency was not equal to ω_0 , but it would be insufficient to say whether the frequency was too high or too low.

Dither Locking. One technique that does work in this circumstance is called dither locking. The idea is to sinusoidally modulate, or dither, the laser frequency at some frequency Ω , producing a voltage signal $V(t) = V(\omega(t)) \simeq V[\omega_{center} + \Delta\omega \cos(\Omega t)]$. If $\beta = \Delta\omega/\Omega \gg 1$, then the voltage signal behaves as if the laser frequency were slowly oscillating back and forth (see the RF modulation discussion in the previous handout).

A lock-in amplifier with reference frequency Ω can then produce an error signal $\varepsilon(\omega)$ that is the Fourier component of $V(t)$ at frequency Ω . Expanding $V(\omega) \approx V_0 - A(\omega - \omega_0)^2$ for ω near the peak frequency ω_0 , we have

$$\begin{aligned} V(t) &= V[\omega_{center} + \Delta\omega \cos(\Omega t)] \\ &\simeq V(\omega_{center}) + \frac{dV}{d\omega}(\omega_{center}) \cdot \Delta\omega \cos(\Omega t) + \dots \end{aligned}$$

and so the Fourier component is

$$\begin{aligned} \varepsilon(\omega) &\sim \frac{dV}{d\omega}(\omega) \cdot \Delta\omega \\ &\sim 2A \Delta\omega (\omega - \omega_0). \end{aligned}$$

This error signal has the desired properties that $\varepsilon(\omega_0) = 0$ and $d\varepsilon/d\omega(\omega_0) \neq 0$; thus it can be used in a feedback loop to lock the laser frequency at ω_0 . (If this isn't clear, draw a picture of $V(\omega)$ near ω_0 and think about it. No matter what the form of $V(\omega)$, selecting the Fourier component of a dithered signal produces an error signal proportional to $dV/d\omega$.)

In the lab, you will set up a dither locking experiment that will lock the laser frequency to a resonance of the confocal cavity, and we'll go over the details of this experiment in the next section.

One problem with the dither-locking method is that the servo bandwidth (how fast the servo can control the system) is limited to a frequency much less than the dither frequency Ω , which in turn must be much less than the frequency scan $\Delta\omega$, and thus much less than the linewidth of the resonance feature on which one wishes to lock. For a very narrow resonance feature, dither locking may not work well. In general, if we want to control something well, we need to control it quickly – in the language of control theory, we need a servo with a large bandwidth.

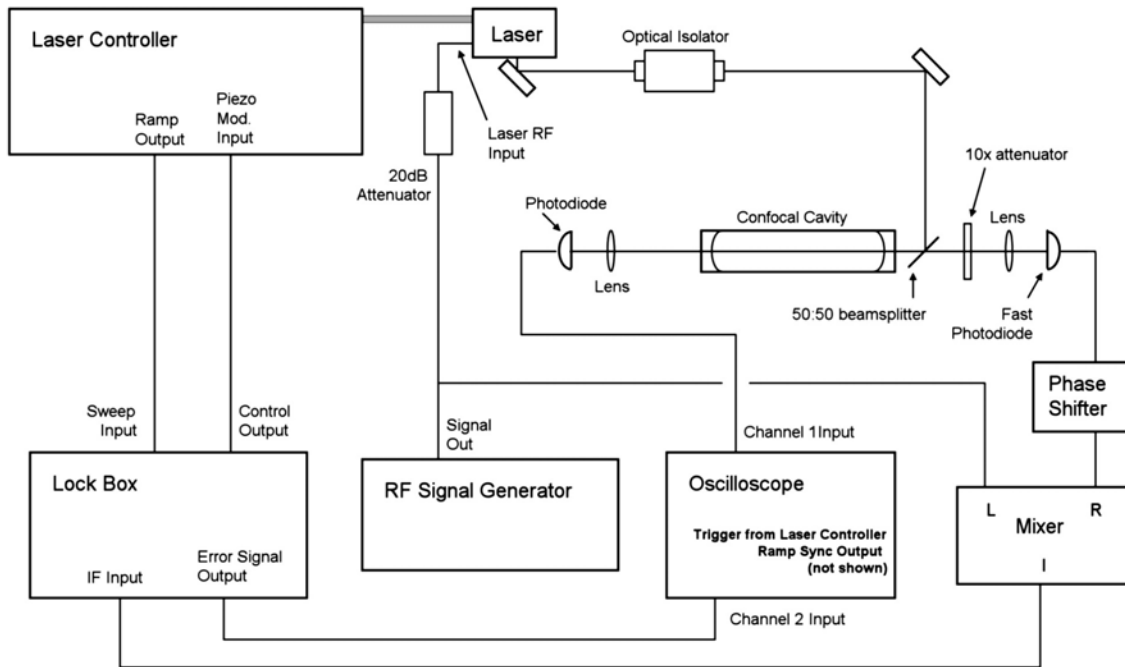


Figure 1. Set up for Pound-Drever-Hall laser locking.

The Pound-Drever-Hall method. We can increase the bandwidth of our dither-locking servo by increasing Ω , but we soon find ourselves in the regime where $\beta \ll 1$. When this happens, we are no longer dithering the frequency, but are modulating the phase of the electric field (again, see the RF modulation discussion in the previous handout). We are adding sidebands to the laser, and our overall picture of what's happening changes, just as it did with RF modulation. The new picture leads us to the Pound-Drever-Hall method.

The P-D-H layout you will use in the lab is shown in Figure 1. It may look complicated at first, but it's fairly simple once you break it down and examine the different pieces:

1. The optics are straightforward, essentially what you set up before. The transmitted signal shows the usual cavity transmission peaks when the laser frequency is scanned. Your goal is to lock the laser to one of the cavity modes, so the transmitted signal stays at the peak value. The transmitted signal is only used as a diagnostic in this lab. When the laser is scanned, it shows that the cavity is properly aligned, and it shows the RF sidebands on the laser. When the laser is locked, you should see a high transmitted signal. This signal is not used as part of the control loop.

2. The RF Signal Generator is used to modulate the laser frequency and thus add sidebands. This is the same as in the previous lab, except the signal is attenuated 100x by a 20dB attenuator. (This is because the mixer needs a large RF amplitude while the laser needs a much smaller RF amplitude.) To set the

notation, we assume two weak sidebands at frequencies $\omega \pm \Omega$ around the laser's central carrier frequency at ω . The modulation frequency Ω is chosen to be slightly greater than the width of the cavity transmission peaks.

3. The Lock Box is some electronics to feed the error signal back to the laser. We'll get to that below, after we look first at how the error signal is produced.

4. The beam reflected off the cavity is sent to a fast photodiode that has a 50 MHz response. For the P-D-H method to work, the bandwidth of the fast photodiode must be higher than the sideband frequency, which will be around 20 MHz. (The other photodiodes only go up to about 1 MHz.) The output from this photodiode will be used to generate the error signal.

5. The Phase Shifter is nothing more than a box of selectable cables of different lengths. Sending the signal through a cable adds a delay, which is the same as adding a phase shift. The longer the cable, the larger the phase shift.

6. The Mixer (or Double-Balanced Mixer) is another passive element, consisting of a set of transformers and diodes. How it actually works is not trivial, so we're going to skip over that part (you can look it up if you're curious). In a nutshell, the mixer takes two inputs – the Local Oscillator (L) and the RF signal (R) – and essentially multiplies them together. The L and R signals are both RF signals with the same frequency but different phases. The output (I) is then proportional to

$$\cos(\Omega t) \cos(\Omega t + \eta) = \frac{1}{2} \cos(\delta) + \frac{1}{2} \cos(2\Omega t + \eta)$$

which has terms at DC and at 2Ω . The mixer output is sent to the Lock Box, where it is smoothed to remove the high-frequency term, and the smoothed output becomes the error signal.

Putting the pieces together, the error signal can be calculated with the following analysis, adapted from Bjorklund *et al* (1983). The reflection amplitude from a Fabry-Perot cavity is given by (see the previous handout):

$$E_r = \frac{(1 - e^{i\delta})\sqrt{R}}{1 - R e^{i\delta}} E_i \equiv A[\delta, R] E_i$$

where R is the mirror reflectivity and $\delta = 2\pi\Delta\omega/\Delta\omega_{FSR}$. You've already seen this in the previous lab, so it should look familiar; a plot of the reflected intensity $|E_r|^2$ shows dips at the cavity resonance (the transmitted light shows peaks, so therefore the reflected light has dips).

For $\beta \ll 1$ we can write the input electric field as a carrier and two weak sidebands

$$E_i = E_0 \left\{ -M e^{i(\omega-\Omega)t} + e^{i\omega t} + M e^{i(\omega+\Omega)t} \right\}$$

which gives the reflected amplitude

$$E_r = E_0 e^{i\omega t} \left\{ -M A_- e^{-i\Omega t} + A_0 + M A_+ e^{i\Omega t} \right\}$$

where

$$A_0 = A[2\pi\Delta\omega/\Delta\omega_{FSR}, R]$$

$$\text{and } A_{\pm} = A[2\pi(\Delta\omega \pm \Omega)/\Delta\omega_{FSR}, R].$$

The reflected signal at the photodiode is then $I_{phot} \propto |E_r|^2$, which you can verify is a real function containing DC terms and terms proportional to $\cos(\Omega t)$ and $\sin(\Omega t)$. Because I_{phot} is a real function, the

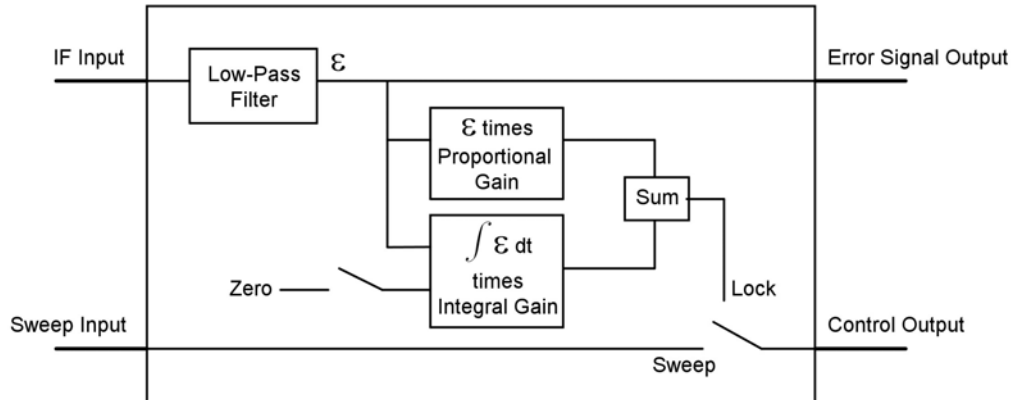


Figure 2. Schematic of the Lock Box.

effect of the mixer is essentially to multiply the photodiode signal by $\exp(i\Omega t + \zeta)$ and take the real part, where ζ is a constant phase factor depending on the relative phase of the photodiode and local oscillator signals at the mixer. Low-pass filtering gives the error signal

$$\varepsilon \propto \text{Re} \{ e^{i\zeta} (A_0^* A_+ - A_0 A_-^*) \}.$$

Problem 1. Calculate and plot $\varepsilon(\Delta\nu)$, where $\Delta\nu = \nu - \nu_0$ in MHz, for a variety of values for the phase factor ζ , assuming a sideband frequency $\Omega/2\pi = 18$ MHz, $\Delta\nu_{FSR} = 375$ MHz, and a cavity reflectivity of $R = 0.95$. (Hint: see the error signal in Figure 4 below.) In the lab, you will generate $\varepsilon(\Delta\nu)$ as the laser frequency is slowly scanned around a cavity resonance and display it on the oscilloscope. If all goes well, the display should match your calculations. You can change ζ using the Phase Shifter, and you should be able to reproduce your calculations at different ζ .

The Lock Box. Figure 2 shows the essential elements of the Lock Box, which is used to feed the error signal back to the laser piezo modulation input, thus controlling the laser frequency and keeping it locked on a cavity resonance. Again, there's quite a bit going on here, so we need to break it down:

1) The IF input is where the error signal enters the Lock Box. Since the mixer output has a substantial 2Ω component, this signal is run through a low-pass filter and an amplifier to clean it up. The result is the smoothed error signal $\varepsilon(t)$. This goes to the Error Signal Output so it can be monitored on the oscilloscope.

2) A ramp signal goes to the Sweep Input (see Figure 1). If the switch on the Lock Box is set to Sweep, then the Sweep Input goes directly to the Control Output. Assuming a triangle-wave input, this simply sweeps the laser frequency back and forth as usual. With the switch set to Sweep, you can view the transmitted light $I_{trans}(\Delta\nu)$ and error signal $\varepsilon(\nu)$ on the oscilloscope to make sure everything looks okay.

3) When the switch is set to Lock, the sweep signal is disconnected. Then the laser frequency is not

sweeping, but the control circuitry will try to lock the laser.

4) Start with the Zero switch closed, which essentially removes the Integral Gain part of the control loop. Then the error signal is simply multiplied by the Proportional Gain and sent to the Control Output; i.e. $C = G_{prop}\varepsilon$, where C is the Control Output signal and G_{prop} is the proportional gain.

If ε is nonzero, then the control signal is nonzero also, and this changes the laser frequency. If the signs are right, then the control will tend to push the laser frequency back toward the cavity resonance. If the gain G_{prop} is very large, then even a small ε will produce a large C , and the laser will stay close to the locked position, as desired. If G_{prop} is not so large (which will be our case in the lab), then the servo will bring the laser closer to the correct frequency, but it will not be able to compensate for large deviations.

One problem with this is that one cannot make G_{prop} very large or the servo will have problems. There are always some phase lags in the electronics and in the mechanical elements, and if the lags are large enough they will result in an instability, causing the servo to oscillate. The theory behind these instabilities can be quite complicated, and we cannot go through it all here. Suffice it to say that problems arise if you increase G_{prop} too much and try to push the servo too hard.

But if G_{prop} is not so large, then the servo will not lock very well. In general, the laser will need a nonzero C to keep the frequency on resonance, and this means ε must be nonzero also, because $C = G_{prop}\varepsilon$. We want to have $\varepsilon = 0$, and in general this cannot happen using proportional gain alone.

5) The next step is to switch in the integral gain, which gives

$$C(t) = G_{prop}\varepsilon(t) + G_{int} \int_{t_0}^t \varepsilon(t') dt'$$

where G_{int} is the integral gain and t_0 is the time the zero switch is opened. Now things get interesting, as the servo has some memory. As long as ε is nonzero, the second term in the above will grow (or diminish, if $\varepsilon < 0$). This means C grows also with time. If the signs are right, this will make C grow until ε goes smoothly to zero. And from that point C will remain unchanged as long as ε remains equal to zero. This effective memory means that the Lock Box can produce a nonzero C even when $\varepsilon = 0$. The result, which you will see in the lab, is a better lock.

Finally, as a disclaimer, we note that servo control theory is a very rich field, vastly more than the above quick explanation might suggest. There is a whole branch of engineering devoted to the topic (at Caltech, this is CDS – Control and Dynamical Systems). Control loops are often quite complex, and their analysis is certainly nontrivial. The above discussion only scratches the surface of this field slightly, but it's enough to demonstrate some of the basics in the lab.

II. LABORATORY EXERCISES

The Thermal Expansion of Aluminum. As an amusing warm-up exercise to this lab, use the extreme sensitivity of these optical techniques to measure the thermal expansion coefficient of aluminum.. You can do this by heating the cavity slightly and watching the cavity expand using the optical spectrum analyzer. You can heat the cavity nicely by touching it, and a thermistor is mounted to the cavity tube to

monitor its temperature.

First set up the cavity as shown in Figure 3 below, with the cavity transmission spectrum displayed on the oscilloscope as the laser frequency is swept. Once the system has stabilized, you should not see much drift of the cavity transmission peaks on the 'scope. Then use a digital ohmmeter to measure the resistance of the thermistor on the cavity. The two wires coming from the body of the cavity are the thermistor leads, and the resistance should be about 2000 Ohms at room temperature. The thermistor resistance changes by about 4.5 percent for each degree of temperature change, with the resistance decreasing as the temperature increases.

The basic idea of this measurement is to count the cavity peaks that drift by on the 'scope as the cavity temperature is changed. One obvious approach to this would be to stabilize the system at one cavity temperature, then change the temperature to some new stable value while counting peaks. The problem with this direct approach is that it takes a long time for the cavity temperature to stabilize.

[Show that the stabilization time scale for the cavity temperature is approximately $\tau \approx CL^2\rho/\kappa$, where C is the specific heat of the aluminum cavity tube, ρ is the density of aluminum, κ is the thermal conductivity, and L is the approximate size scale of the cavity. With $C \approx 900$ J/kg-K, $\kappa \approx 250$ W/m-K, $\rho \approx 2700$ kg/m³, and $L \approx 0.1$ m, we have $\tau \approx 100$ seconds. And of course one would have to wait several times τ before the temperature is stable.]

A better approach is to simultaneously monitor the time derivatives of the cavity length and the cavity temperature. Starting from a fairly stable temperature, heat the cavity slowly by touching it gently. Try to heat it uniformly along its length (because of the L^2 in the above expression). Record the thermistor resistance while you count cavity peaks going by on the 'scope. This is best done with two people. Make a measurement with the cavity heating, and another with the cavity cooling. From this you can extract the cavity length as a function of cavity temperature, and thus the thermal expansion coefficient. Report your measurement in relative expansion per degree, which should be around 20×10^{-6} K⁻¹.

Laser Locking I - Dither Locking.

The next step is to lock the laser frequency to the confocal cavity using the dither locking technique described in Section I. The optics and electronics set-up is shown in Figure 3. This is quite different from the P-D-H set-up, so let's look at the different pieces:

1. The optics are again familiar, and somewhat simpler than the P-D-H set-up, since you will be using the transmitted signal to generate an error signal. To make life easier, adjust the cavity length and/or the input beam alignment so the cavity has an effective finesse of about 10. The low finesse makes the lock work much better. With a higher finesse, the frequency jitter in the cavity and/or laser makes the cavity output very jumpy. The feedback control won't be fast enough to take out the high-frequency jitter (which comes from acoustic and/or seismic effects), and the resulting noise will quickly throw the servo out of lock. With a low finesse, the jitter doesn't show up much on the cavity output, so the servo only has to compensate for low-frequency drift. To make the servo work with a high cavity finesse, we would probably have to reduce the jitter passively first (with acoustic damping, seismic isolation, etc.).

2. You sweep the laser frequency using the Ramp Generator and Piezo Controller on the Laser Controller,




Figure 3. Set-up for dither locking.

basically the same as before. (You should know what this is doing by now. If not, ask your TA to remind you.) The only difference is that the signal is diverted into the Lock Box. When the switch on the Lock Box is set to Sweep, then the Sweep Input goes directly to the Control Output. By setting up these two cables, the laser frequency will scan, as you have done before. Things seem to work best if you turn the Ramp Generator to full amplitude, then turn the Piezo Modulation gain down (both on the Laser Controller) so you are only scanning over a couple of transmission peaks.

3. The photodiode output goes directly to the oscilloscope, so you can monitor the light transmitted by the cavity.

4. Use a BNC tee to also connect the photodiode output to the Detector Input on the Laser Controller. This way you have a separate detector gain control, which will be useful.

5. Configure the Signal Generator to produce a sine wave at 50 kHz; turn the gain down low and make sure the output is set to the 0-2V position. This goes to the Current Modulation Input, which will dither the laser frequency. If you turn the amplitude up, you should be able to see the dither directly on the photodiode output. If this doesn't make sense, or you're not sure what to look for, consult your TA.

6. Now for the error signal. First send the Sync Output from the Signal Generator to the Reference Input of the Lock-In Amplifier. The Lock-In will automatically sense the reference signal and lock onto it. You should see the reference frequency on the Lock-In display. (If not, ask your TA.)

7. Next send the Detector Monitor Output to the Lock-In Signal Input. To begin at least, set the

Lock-In time constant to 100 μsec and the sensitivity to 1 V, and set the input filter to 12dB. With this, the Lock-In will do its thing (ask your TA if you don't now how a lock-in amplifier works) and detect the Fourier component of the input signal at the reference frequency. That is, it multiplies the input signal by $\sin(\omega t)$, integrates for the selected time constant, and sends the result to its X output. (It also multiplies by $\cos(\omega t)$ and sends the result to the Y output, but you won't be using that today.)

8. Then send the X output to the IF input of the Lock Box, and send the Error Signal Output to the Oscilloscope as shown. All the Lock Box does at this point is smooth the signal a bit, so basically you're sending X to the oscilloscope. With this, you should see a sensible error signal on the 'scope. Its shape will be essentially the derivative of the cavity output signal, because that's what the dither-lock set-up does. (Why? See Section I and think about it.) Show your TA the signals before proceeding. It usually takes some tweaking (maybe a lot of tweaking) to get everything looking good before you attempt to lock the laser frequency to the cavity. The goal is to produce a clean, smooth error signal that is about a one volt peak-to-peak.

Check the Phase of the Error Signal. Once the error signal looks good, change the reference phase using the Lock-In controls (ask your TA for details.) Changing the phase from zero means the Lock-In multiplies by $\sin(\omega t + \delta)$ instead of $\sin(\omega t)$, where δ is the Lock-In phase. If you change the phase from 0 to 180 degrees, the error signal should change sign. Check this.

The Lock Box looks at the error signal and feeds a signal back (the Control Output) to the Laser Controller. When you move the switch to Lock, it will try to control the laser frequency and zero the error signal. For this to work, the error signal has to have the correct sign, and you change this using the Lock-In phase. Which sign is the right sign? You just have to try it to see.

Flip the switch on the Lock Box to Lock, while the Integral Gain is off. This means you are only using Proportional Gain. Adjust the Piezo Controller DC Offset by hand to see if you can get the servo to lock right, remembering that the servo has only a small dynamic range (that is, the Control Output signal is fairly small). When you're close to the cavity resonance, you should see the error signal go to zero while the transmitted light signal is near its maximum value. If the sign is right, the error signal will dwell around zero. If the sign is wrong, the error signal will tend to avoid zero. The various nuances of the behavior are hard to explain, so consult your TA at this point.

Try this exercise using both signs of the error signal. You should find that the cavity achieves a weak lock with one sign, but not with the other.

Lock the Laser. Once you know which sign is the right sign, lock the laser again using only Proportional Gain. Keep one hand on the DC offset knob to keep the lock centered so the error signal remains close to zero (while the cavity transmission is near maximum). You are then part of a human servo as you watch the 'scope and feed back to the DC offset knob. If all is well, you can make this work, sort of. Then have your partner place a hand a few centimeters away from the cavity, heating it slightly. (Don't touch the cavity; that's too much heating.) You should see that this induced cavity drift makes it much more difficult for the human servo to keep the laser on resonance.




Figure 4. The top trace shows a typical error signal; the bottom trace shows the cavity transmission, taken at the same time.

Take the hand away, get the lock working again, and then switch on the Integral Gain. The cavity should remain locked, and the human servo should no longer be necessary to keep the lock centered. Try heating the cavity again. The lock should stay centered. How does this magic work? Read the explanation in Section I again, and think about it. **Print a copy of the error signal and transmitted signal with the laser locked, and again with the laser scanning.** With the full lock on, clap your hands or tap lightly on the table to test the lock stability. What happens if you change the DC offset on the piezo controller (with the lock on)?

Laser Locking II - The Pound-Drever Hall Method.

Once you are satisfied with your dither lock, you can move on to the Pound-Drever-Hall lock, shown in Figure 1. The basic idea is the same as with dither locking, except the method for generating the error signal is quite different. To begin, use 18 MHz sidebands, the same as you used for your calculations. When you have everything set up, you should see an error signal that looks something like that shown in Figure 4. If not, consult your TA if you need help debugging the set-up. One common problem is saturating the fast photodiode, and you will probably have to attenuate the light hitting the sensor, as shown in Figure 1.

Once you see something reasonable, change the phase using the Phase Shifter, and try changing the RF frequency. **Print out several error signal traces at different phases and frequencies. In particular, try to generate signals that match your calculations.**

When you set out to lock the laser, be sure to trim the DC level of the error signal. Far from resonance, the error signal should be at zero volts. **Print out traces with the laser locked and unlocked.**

III. REFERENCES

Bjorklund, G. C. *et al.* 1983, "Frequency-modulation (FM) Spectroscopy—Theory of Line Shapes and Signal-to-noise Analysis," *Appl. Phys. B* 31, 145.

Drever, R. W. P. *et al.* 1983, "Laser Phase and Frequency Stabilization Using an Optical Resonator," *Appl. Phys. B* 31, 97.

Hall, J. L. and Zhu, M. 1992, "An Introduction to Phase-Stable Optical Sources," *Proceedings of the International School of Physics – Enrico Fermi*, edited by E. Arimondo, W. D. Phillips, and F. Strumia (North-Holland).

Yariv, A. 1991, *Optical Electronics*, Saunders College Publishing, 4th ed.