

A review of Pound-Drever-Hall laser frequency locking

M Nickerson

JILA, University of Colorado and NIST, Boulder, CO 80309-0440, USA

Email: nickermj@jila.colorado.edu

Abstract. This paper reviews the Pound-Drever-Hall technique of frequency stabilization, primarily as applied to locking a laser source to a resonant mode of an optical cavity. Presented are both a high-level conceptual overview, discussing the general principles of operation, and a mathematical derivation of the linear, offset-dependent error signal used to stabilize the laser. Briefly covered are details of servo loops used to enact frequency locking, and various methods by which laser frequencies may be adjusted.

1. Introduction

A common necessity in labs utilizing lasers is a highly stable laser source. One possible method of creating a stable laser is to use a Fabry-Perot laser cavity of known stability as a reference. If the frequency of a laser can be dynamically adjusted to match the resonant mode of a cavity, the laser will then have the same stability as the cavity. An excellent technique for achieving this is the “Pound-Drever-Hall” locking method, invented in 1983. It was conceived by Ron Drever, based on similar microwave techniques used by R. V. Pound. The practical implementation was largely developed by Jan Hall at JILA [1]. The most popular laser-locking method currently in use, the Pound-Drever-Hall method originally delivered better than 100 Hz stability [2]. Modern implementations with careful vibration isolation have managed to achieve sub-Hz stability [3]. Also of note, this method can be adapted to lock a cavity to a stable laser via piezo length actuation, which is useful in some circumstances, such as precisely measuring small changes in cavity length [4]. In this use, the laser is unmodified, while the resonant mode of the cavity is altered by changing the cavity’s length. The rest of this review will assume that a laser’s frequency is being adjusted to match a stable cavity.

2. Conceptual overview

The objective of laser locking is fairly simple. One has a laser of unknown (unstable) frequency, and wishes to dynamically alter this to match the known frequency of a specific cavity mode. Thus, the difference between the laser's frequency and the cavity's resonant frequency, δf , must be continually measured, and used to appropriately alter the laser frequency. The Pound-Drever-Hall technique is one method to accomplish this, and has the advantage of being insensitive to laser intensity fluctuations.

There is currently no electronics which can directly measure δf . To find information about the sign of δf , it is possible to use the reflective properties of a Fabry-Perot cavity. When the laser frequency is far from cavity resonance, all light is reflected from the cavity. When it is near resonance, however, the reflection coefficient, F , depends on δf , vanishing when δf is null. Older laser locking methods simply used this dependence of F to stabilize the frequency. However, since F is symmetric around resonance, it was not possible to lock a laser's frequency directly on the cavity mode; one would not be able to determine which side of it the laser frequency was on. Thus, reflection-only locking must use some small offset, such that the value of reflected intensity can indicate which side of the desired value one is on. Unfortunately, this method suffers from dependence on laser intensity; if the intensity changes, it mimics the effect of a frequency change.

The Pound-Drever-Hall method avoids this: while the reflection coefficient is symmetric, the derivative of F with respect to δf is antisymmetric across resonance. This provides a signal proportional to the frequency offset, called an error signal, indicating which side of resonance the laser is on. To determine the derivative at any point, one can simply vary the frequency slightly and see how the reflected intensity responds, as indicated in figure 1. Above resonance, with a positive F derivative, the intensity will respond in phase with frequency variations. Below resonance, the intensity will respond with opposite phase, and on resonance, the intensity will not change

significantly. Thus, with slight frequency variations, one can determine which side of the cavity resonance the laser frequency is on, and adjust appropriately. This is the basic operational concept of the Pound-Drever-Hall method.

Figure 2 shows a basic setup that is capable of implementing this. The light from a laser passes through an electro-optic modulator, which modulates its phase as given by an oscillator (in practice, this is easier than modulating the frequency). It then passes through the cavity, and the reflected beam is monitored by a photodetector. Mixing the photodetector signal and the output from the oscillator isolates the part of the signal measuring the derivative of F , which is then fed into a servo amplifier to appropriately adjust the laser frequency, locking the laser to the cavity's resonance.




Figure 1. The reflection intensity near resonance as a function of frequency.

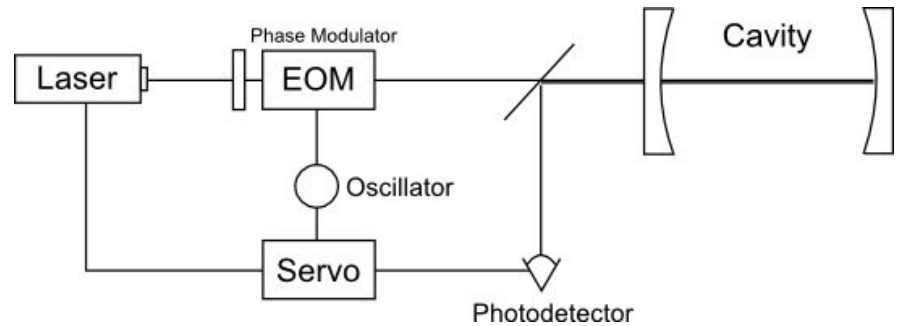


Figure 2. A basic setup capable of implementing Pound-Drever-Hall locking.

3. Detailed treatment: extracting an error signal

To understand the Pound-Drever-Hall locking technique in detail, it is necessary to explore a mathematical treatment of the process. The transition from frequency to phase modulation and extraction of the error signal is not necessarily intuitive, but is obvious in the mathematics.

We start with the properties of the Fabry-Perot resonator cavity. If we assume that the incident and reflected beams have an approximately constant frequency, then we can write the magnitude of the incident electric field as $E_{inc} = E_0 e^{i\omega t}$. The reflected light consists of both the promptly reflected

beam, directly off the mirror, and the leakage beam, which is the part of the standing wave inside the cavity that leaks out through the mirror. Examining the phases of each component, the prompt reflection has a phase shift of π . The leakage beam consists of multiple phase components, the first round-trip component having a phase shift of $-2L\omega/c$ from its travel along the cavity of length L , where $\omega = 2\pi f_L$ is the angular laser frequency. For a cavity with symmetric mirrors of reflectivity r and transmittivity $t = \sqrt{1-r^2}$, the reflected beam will be

$$E_{refl} = E_0 \left(r e^{i(\omega t + \pi)} + t r t e^{i(\omega t - 2L/c)} + t r^3 t e^{i(\omega t - 4L/c)} + \dots \right)$$

which simplifies to

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

where $\Delta\nu_{FSR} = c/2L$ is the free spectral range of the cavity, the frequency distance between resonances. See figure 3 for plots of the amplitude and phase of $F(\omega)$. It is evident that the phase is similar to the derivative of $F(\omega)$ discussed earlier in section 2, in that it is antisymmetric across the cavity resonance frequency. Thus something proportional to this would make an ideal error signal.




Figure 3. Plots of the amplitude and phase of the reflection coefficient $F(\omega)$, assuming r of 0.999.

Now we introduce the phase modulation into the laser beam. An electro-optic modulator, or EOM, is usually used for this. An EOM could be considered similar to a variable waveplate; it is a crystal which can rapidly change its index of refraction in response to electrical stimuli. This induces a small phase change in light passing through it, which can be rapidly varied. A laser beam passing through an EOM modulated at frequency $\Omega/2\pi$ will have an electric field of

$$\begin{aligned} E_{inc} &= E_0 e^{i(\omega t + \beta \sin \Omega t)} \\ &\approx E_0 \left(J_0(\beta) + 2i J_1(\beta) \cdot \sin \Omega t \right) e^{i\omega t} \\ &= E_0 \left(J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega + \Omega)t} - J_1(\beta) e^{i(\omega - \Omega)t} \right) \end{aligned}$$

valid for small β , where the J 's are Bessel functions. This expansion explicitly shows that after phase modulation, there are really three different frequency components in the beam. The two J_1 components are termed 'sidebands', being offset to either side of the primary frequency component by the modulation frequency $\Omega/2\pi$. Note that a first-order Taylor expansion will give similar results.

Once this phase-modulated beam reaches the cavity and reflects, each frequency component will simply be transformed by $F(\omega)$, resulting in a reflected electric field of

$$E_{refl} = E_0 \left(F(\omega) J_0(\beta) e^{i\omega t} + F(\omega + \Omega) J_1(\beta) e^{i(\omega + \Omega)t} - F(\omega - \Omega) J_1(\beta) e^{i(\omega - \Omega)t} \right)$$

However, the reflected electric field is not measured directly. Rather, the power (intensity),

$P_{refl} = |E_{refl}|^2$, is the quantity detected by the photodetector. Expanding out the reflected power,

$$\begin{aligned} P_{inc} &= P_c |F(\omega)|^2 + P_s \left(|F(\omega + \Omega)|^2 + |F(\omega - \Omega)|^2 \right) \\ &+ 2\sqrt{P_c P_s} \operatorname{Re} \left[F(\omega) F^*(\omega + \Omega) - F^*(\omega) F(\omega - \Omega) \right] \cdot \cos \Omega t \\ &+ 2\sqrt{P_c P_s} \operatorname{Im} \left[F(\omega) F^*(\omega + \Omega) - F^*(\omega) F(\omega - \Omega) \right] \cdot \sin \Omega t \\ &+ (O[2\Omega]) \end{aligned}$$

where P_c and P_s are the power of the carrier (primary) and sideband components, respectively. It is evident that the reflected power contains several frequency components: a DC power from the carrier,

two components oscillating at the modulation frequency from the sidebands, and higher-order components from the interactions between the sidebands.

To retrieve the phase information, we need to separate one of the oscillating terms. For a reasonably fast modulation frequency, the $\sin \Omega t$ term with the imaginary component of $F(\omega)$ will be dominant. Since this is simply proportional to the original oscillation signal passed to the EOM, one can use a “mixer” to extract it. A mixer is an electronic device that essentially multiplies two signals together; multiplying $P_{inc} \propto \sin \Omega t$ from the photodetector with $\sin \Omega t$ from the oscillator will result in a DC component and a $\cos 2\Omega t$ term. Isolating the DC component, usually with a low-pass electronic filter, results in the Pound-Drever-Hall error signal,

$$\varepsilon = 2\sqrt{P_c P_s} \operatorname{Im}\left[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)\right]$$

Figure 4 shows a plot of the PDH error signal against laser frequency offset from resonance, where the modulation frequency is 5% of the cavity’s free spectral range. Note that the error signal is antisymmetric, as expected, and there is a large frequency over which the sign corresponds to which side of cavity resonance the laser is on. Additionally, the error signal has a high slope near resonance, which is helpful for the servo loop.




Figure 4. Plot of the Pound-Drever-Hall error signal, when the modulation frequency is high (here, 5% FSR)

If the modulation frequency is sufficiently high, as assumed, then the sidebands will be far from resonance. Thus they will be totally reflected, and the error signal becomes proportional to $\text{Im}[F(\omega)]$. If the laser frequency is close to the cavity resonance, then $F(\omega)$ can be approximated as

$$\text{Im}[F(\omega)] \approx \frac{1}{\pi} \frac{\delta\omega}{\delta\nu}$$

where $\delta\omega$ is the offset from resonance, and $\delta\nu$ is the cavity linewidth. Then the error signal can be approximated as

$$\varepsilon = \frac{4}{\pi} \sqrt{P_c P_s} \frac{\delta\omega}{\delta\nu} = D \delta f$$

$$D \equiv \frac{8\sqrt{P_c P_s}}{\delta\nu}$$

where δf is the frequency offset from resonance. It is apparent that the error signal is thus linear with offset frequency near resonance, an ideal condition for a servo loop to operate. This region corresponds to the highly sloped central area about zero in figure 4.

4. Feedback loop and frequency adjustment

Now that we have a clear error signal that is linear with frequency offset, the rest of the Pound-Drever-Hall technique consists of using this to adjust the laser frequency onto the cavity resonance. Methods for achieving this vary, but they all include some sort of servo loop.

A servo loop is a general term for any system, usually electronic, that takes some input error signal, and adjusts the output until the error is zero. A common class of servos are “PID” devices. These simply transform an input by a Proportional term, an Integral term, and a Derivative term, then output the result. The proportional term provides immediate negative feedback, attempting to zero the input. The integral term is used to eliminate any residual offsets that persist over time; one can see the necessity of this by considering the final value of a proportional-only system. If the error input is

given by $\varepsilon = O + A$, and the P system has an output $O = -k \varepsilon$, then the final value of the error signal will end up as a constant $\varepsilon = A/(1+k)$. By adding an integral term, $O \propto -k_I \int \varepsilon$, the error signal is forced to zero after some time. A derivative term is usually used to avoid overcorrecting slowly-reacting systems; it primarily slows down the reaction.

For laser stabilization, generally only a PI loop is used. Lasers respond fast enough to frequency adjustments that a derivative term is unnecessary. At JILA, the primary servo used for frequency stabilization is the “JILA loop filter”. This is an analog electronic PI servo with a very fast response time of 10s of nanoseconds. The response speed is generally rated by bandwidth, for instance, a 100 MHz bandwidth device has a response time of 10 ns. Responding at this rate in a Pound-Drever-Hall locking system, it will successfully eliminate any frequency fluctuations of 100 MHz or less, if properly adjusted.

To successfully stabilize a laser, a servo loop must have some means of adjusting the laser’s output frequency. There are several common methods to do this, although the details often vary by laser type.

Many lasers respond to temperature adjustment of the lasing medium. In a YAG laser, for example, adjusting the temperature of the YAG crystal will change its refractive index. This changes the optical path length of the lasing cavity, altering the lasing frequency. Temperature adjustment is not used as a primary means of laser locking, however, being very slow relative to the frequency noise. This is primarily due to thermal dissipation times in the lasing medium, which can require up to a second or more to come to thermal equilibrium. Due to the large frequency range available through thermal adjustment, it is often used as an initial coarse frequency tuning, before any servo loops are implemented.

A common frequency modulation method in free-space lasers is the insertion of a secondary, piezo-actuated, optical cavity inside the lasing cavity. Figure 5 shows a schematic example of what this might look like. One or two mirrors, mounted on piezoelectric actuators, are inserted into the standard laser cavity. When voltage is applied to the piezos, they slightly change their volume, moving the mirror. This alters the length of the cavity, changing the frequency at which a standing wave can form, and thus at which the gain medium can continuously lase. The modulation speed of piezo-actuated laser cavities is restricted only by the piezoelectric material, and the structure of the cavity. Many commercial lasers have rated speeds in the dozens of megahertz. Due to this very high speed, piezo-actuated cavities are a common method for servo loop frequency adjustment. In general, the frequency range of piezo actuators is restricted to a few hundred megahertz. This is limited both by the very high voltages required for significant piezoelectric displacement, and the free spectral range of the laser cavities. Another limitation is the integral nature of this tuning method. The piezoelectric actuators must be built into the laser cavity, and cannot easily be added to sealed units after manufacture.

A third method of frequency tuning is the use of Acoustic-Optic Modules. These are crystals with piezoelectric devices mounted on one edge. Modulating the piezos at high frequencies sets up acoustic waves within the crystal, and light passing through the AOM experiences Bragg diffraction off of the acoustic waves. The diffracted light is Doppler shifted by the piezo modulation frequency, which provides the method of adjusting laser frequency. The tuning range is limited by the maximum piezo frequency, generally under 300 MHz. The modulation bandwidth is also restricted to the few-megahertz range by the sound velocity within the crystal. AOMs can be used with any laser, however, in any position within the optical setup. They are a method of frequency tuning that is completely independent of the laser source, and thus can be a convenient method of adding cavity-locking capability to any laser. Unfortunately, the diffraction angle of the emerging laser depends on the modulation frequency. Using an AOM in a servo loop, therefore, causes constantly changing beam

geometry. A solution to this problem is to use a radial mirror, and double-pass the light through the AOM, such as is shown in figure 6.




Figure 5. Schematic example of a piezo-actuated lasing cavity.

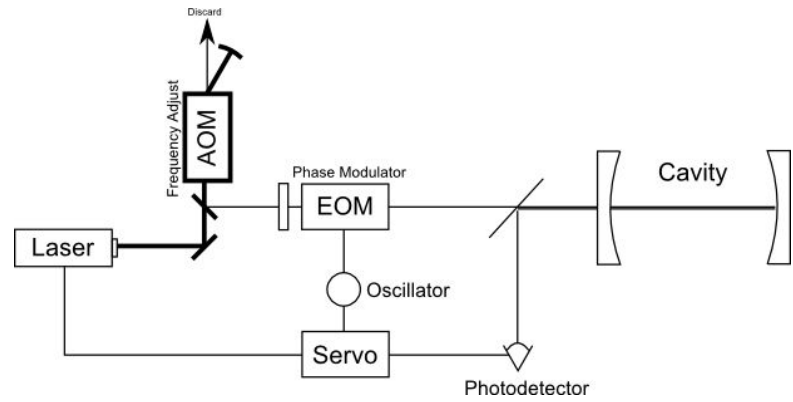


Figure 6. Possible use of an AOM for frequency tuning in a PDH setup.

5. Conclusion

The Pound-Drever-Hall laser stabilization method allows a laser to be frequency-locked to a stable cavity. The basic operating principle is conceptually simple: read the laser's offset from cavity resonance, and feed this into a servo loop which adjusts the frequency. The Pound-Drever-Hall error signal is linear with offset frequency near resonance, making the servo loop straightforward to manage. This linear signal is generated by measuring the phase difference between the incoming laser light, and the light resonating in the cavity. Sidebands are introduced into the primary laser signal by phase modulation, providing a reference by which to measure this phase difference. Once an error signal is obtained, a servo loop interfaces with a frequency-modulating component in or near the laser cavity, and thus adjusts the laser frequency until the error signal is minimized. Using the Pound-Drever-Hall method, stable lasers with as little as a half-hertz linewidth have been created.

References

- [1] Eric D Black, "An introduction to Pound-Drever-Hall laser frequency stabilization", *American Journal of Physics* 69, 79 (2001).
- [2] R W P Drever, J L Hall, F V Kowalski, J Hough, G M Ford, A J Munley, and H Ward, "Laser phase and frequency stabilization using an optical resonator", *Appl. Phys. B* 31, 97 (1983).
- [3] J Alnis et. al., "Sub-Hz line width diode lasers by stabilization to vibrationally and thermally compensated ULE Fabry-Perot cavities", *Physical Review A* 77, 053809 (2008).
- [4] Eric D Black, "Notes on the Pound-Drever-Hall technique", LIGO Technical Note (1998).
- [5] V Leonhardt, "Displacement Measurements on Suspended Mirrors for Off-resonant Thermal Noise Detection", PhD thesis, University of Hannover (2003).
- [6] N P Schmitt, P Peuser, S Heinemann, A Mehnert, "A model describing the single and multiple line spectra of tunable microcrystal lasers", *Optical and Quantum Electronics* 25, 527 (1993).