

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
-LIGO-
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note LIGO-T980045- 00- D 4/16/98
Notes on the Pound-Drever-Hall technique
Eric Black

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: <http://www.ligo.caltech.edu/>

Contents

1	Introduction	2
1.1	Motivation	2
1.2	Background	2
2	A conceptual model	3
3	A quantitative model	5
3.1	The incident beam	5
3.2	The reflected beam and the error signal	5
3.3	Approximations near resonance	7
4	Fundamental limits	8
4.1	Shot noise limited resolution	8
4.2	The sensitivity theorem	9
5	Technical considerations	10
5.1	Getting close to the shot noise limit	10
5.2	Maximizing the slope of the error signal	10
5.3	Noise in other parameters	12
A	The reflection coefficient	12

1 Introduction

These are notes on the Pound-Drever-Hall technique for doing interferometry with a Fabry-Perot cavity. They are intended as introductory material to supplement the published literature for members of the Thermal Noise Interferometer group. Understanding the Pound-Drever-Hall technique is essential to understanding the workings of the Thermal Noise Interferometer (TNI), as well as LIGO. These notes include a review of the fundamentals of the technique as well as some quantitative predictions about sensitivity and noise.

1.1 Motivation

The technique was invented for stabilizing the frequency of a laser by locking it to a Fabry-Perot reference cavity, and that is still its primary use. It was invented by Ron Drever, based on an earlier microwave technique invented by R. V. Pound, and much of the implementation was worked out at JILA by Jan Hall's group. Although laser stabilization is its primary use, the technique can also be used the other way around: to lock a cavity to a laser. When you do this, you can measure small changes in the length of the cavity with extremely good precision, which is important if you want to do gravitational wave astronomy.

In our case, we want to measure the small changes in the length of a cavity caused by thermal noise to see how it will affect a gravitational wave antenna. We will use the Pound-Drever-Hall technique to measure this length noise. Since we need a very stable laser as a reference, we will also use the Pound-Drever-Hall technique for stabilizing the frequency of our laser.

1.2 Background

In writing this document I am assuming that the reader is familiar with Fabry-Perot cavities, but not necessarily with the Pound-Drever-Hall technique. I am also assuming an elementary knowledge of feedback control. Most physicists have no formal training in feedback control theory, but the subject is taught in the standard undergraduate electrical engineering curriculum. There are a number of very accessible textbooks on control theory, two of which are Friedland [1] and Franklin, Powell, and Emami-Naeni [2]. The first chapter in Friedland provides enough background to follow these notes.

Some excellent primary sources on the Pound-Drever-Hall technique are Tim Day's thesis [3], Ron Drever's original paper [4], and the standard analysis of the theory by Bjorklund, *et al.* [5]. A good collection of background material relevant to gravitational wave detection can be found in Peter Saulson's book [6]. Other references abound. The technique is also taught in the undergraduate optics lab here at Caltech, and the class notes for that lab are also a good source of information.




Figure 1: The basic layout for locking a cavity to a laser. Solid lines are optical paths, and dashed lines are signal paths. The signal going to the far mirror of the cavity controls its position.

2 A conceptual model

Fig. 1 shows a basic Pound-Drever-Hall setup. This arrangement is for locking a cavity to a laser, and it is the setup you would use to measure the length noise in the cavity.¹ You send the beam into the cavity; a photodetector looks at the reflected beam; and its output goes to an actuator that controls the length of the cavity. If you have set up the feedback correctly, the system will automatically adjust the length of the cavity until the light is resonant and then hold it there. The feedback circuit will compensate for any disturbance (within reason) that tries to bump the system out of resonance. If you keep a record of how much force the feedback circuit supplies, you have a measurement of the noise in the cavity.

Setting up the right kind of feedback is a little tricky. The system has to have some way of telling which way it should push to bring the system back on resonance. It can't tell just by looking at the

¹Locking the laser to the cavity follows essentially the same design, the only difference being that you would feed back to the laser, rather than the cavity.

intensity, since it is the same on either side of resonance. The system needs to be able to tell which way the intensity will go if it pushes one way or another. It has to be able to sample the derivative of the intensity. You could try moving the mirror back and forth to see which way makes the intensity go down, but it is usually easier to adjust the laser.

In Fig. 1, the Pockel's cell modulates the *phase* of the laser beam, and the reflected light is compared with the modulation signal (LO). For a conceptual understanding, it is more useful to imagine that you are modulating the laser's *frequency*.

If you are above resonance, increasing the laser's frequency increases the power in the reflected beam. Below resonance, increasing the laser's frequency decreases the reflected power. If you modulate the system, you can tell which side of resonance you are on by whether the reflected power is varying in phase or 180° out of phase with the modulation. (See Fig. 2.)




Figure 2: The reflected light intensity from a Fabry-Perot cavity as a function of laser frequency, near resonance. If you change the laser frequency, you can tell which side of resonance you are on by how the reflected power changes.

The mixer compares the modulation signal (generated by LO in Fig. 1) with the output of the photodetector, extracting the part that is at the same frequency as the modulation signal. (The mixer's output is just the product of its inputs.) The sign of the mixer's output is different on either side of resonance, and it is zero when the system is exactly on resonance. This is just what you want for a feedback control signal. In control theory, this feedback signal is called the *error signal* and is typically designated ε . Sending this error signal to the actuator on the mirror completes the feedback loop and locks the system on resonance.

You may want some amplification and possibly a change of sign (to make sure the feedback is negative), so you can put an amplifier between the mixer and the actuator. That's what the "Servo Amp" in Fig. 1 is for. You can tell how much noise is in the cavity by monitoring the error signal.

This conceptual model is really only good when the dithering the laser frequency slowly. If you dither the frequency too fast, the light resonating inside the cavity won't have time to completely build up or settle down, and the output will not follow the curve shown in Fig. 2. Also, although it is possible to dither the frequency of a laser, it is usually much easier and much more desirable to modulate its phase. In any practical implementation of this technique, you want to use a very fast

phase modulation, and you have to resort to a quantitative model to understand what is going on. The rest of these notes are devoted to this quantitative model.

3 A quantitative model

3.1 The incident beam

A simple laser beam will have an electric field given by $E_0 e^{i\omega t}$. I talked about dithering the frequency of this beam in the last section, but in practice it is easier to dither the phase. The results are essentially the same. The Pockels cell shown in Fig. 1 accomplishes this phase dither. After the beam has passed through this Pockels cell, its electric field is²

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

I have written it in this form to show that there are actually three different beams incident on the cavity: a carrier, with frequency $\omega/2\pi$, and two sidebands with frequencies $(\omega \pm \Omega)/2\pi$. Here, $\Omega/2\pi$ is the phase modulation frequency, and β is known as the modulation depth. If P_0 is the total power in the beam, then the power in the carrier is

$$P_c = J_0^2(\beta) P_0,$$

and the power in each sideband is

$$P_s = J_1^2(\beta) P_0.$$

When the modulation depth is small ($\beta \ll 2$), almost all of the power is in the carrier and the first order sidebands.

$$P_c + 2P_s \approx P_0$$

3.2 The reflected beam and the error signal

If you send a simple, monochromatic beam ($E_i = E_0 e^{i\omega t}$) into a Fabry-Perot cavity, some of it will get reflected. The properties of the reflected beam are straightforward: It must have the same frequency as the incident beam, and its amplitude must be proportional to the incident beam's amplitude. This suggests that the reflected beam would be given by $E_{ref} = F E_{inc}$, and if you do a careful calculation, that's exactly what you get. Most elementary optics books will get you started on this derivation, so I will not reproduce it here. (See Appendix A for the general formula for F and some useful results.)

²The reader who doesn't like Bessel functions will find that the small angle expansion, $E_{inc} \approx E_0 [1 + i\beta \sin \Omega t] e^{i\omega t} = E_0 \left[1 + \frac{\beta}{2} (e^{i\Omega t} - e^{-i\Omega t})\right] e^{i\omega t}$ works just about as well.

There may be some phase shift on reflection, but we can take that into account by letting the reflection coefficient F be complex.

The reflection coefficient depends on properties of both the beam and the cavity. For a symmetric cavity with no losses, it is

$$F = \frac{r (e^{i\phi} - 1)}{1 - r^2 e^{i\phi}}, \quad (1)$$

where r is the amplitude reflection coefficient of each mirror, and ϕ is the phase the light picks up in one round trip inside the cavity. The general formula for ϕ is $2\omega L/c$, where $\omega/2\pi$ is the laser frequency, and L is the cavity length. You can get a better feel for the physics by writing it in terms of the laser frequency in cavity free spectral ranges,

$$\begin{aligned} \phi &= \frac{\omega}{\Delta\nu_{fsr}} \\ &= 2\pi \frac{f}{\Delta\nu_{fsr}}, \end{aligned}$$

or the round trip optical path in the cavity in wavelengths of light,

$$\phi = 2\pi \frac{2L}{\lambda}.$$

The laser frequency and the cavity length are on equal footings. The reflection coefficient is periodic in both, following an Airy function, and it vanishes when either the laser frequency is an integer multiple of free spectral ranges, or the round trip optical path is an integer number of wavelengths. The two conditions are the same, of course, but sometimes you want to think in terms of laser frequency, and sometimes you want to think in terms of cavity length.

To calculate the reflected beam when there are several incident beams, you just multiply each of the incident beams by the reflection coefficient at the appropriate frequency. In the Pound-Drever-Hall setup, where we have a carrier and two sidebands, the reflected beam is

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

(I have written $F = F(\omega)$ in terms of the laser frequency for convenience.) What we really want is the power in the reflected beam, since that is what we measure with the photodetector. This is just $P_{ref} = |E_{ref}|^2$, or after some algebra

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

The mixer pulls out the term that is proportional to $\sin \Omega t$, so the error signal is

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




Figure 3: The Pound-Drever-Hall error signal, $\varepsilon/2\sqrt{P_c P_s}$ versus ϕ . Here, the modulation frequency is 4% of a free spectral range, and the cavity finesse is about 100.

Fig. 3 shows a plot of this error signal for a lossless cavity with identical mirrors. Note that the error signal crosses zero if any of the beams (the carrier or either of the sidebands) resonates in the cavity. However, the sign of the error signal's slope is opposite for the carrier and the sidebands. If you have made a mistake and wired your system up with positive feedback, then it is possible to lock to the sidebands instead of the carrier. This provides a good test to see if you've got the polarity of your servo right, since you can usually tell a sideband from a carrier just by looking.

3.3 Approximations near resonance

Let's consider a lossless cavity with identical, highly reflective mirrors. Then Eqn. 1 gives the reflection coefficient. If the carrier is near resonance, and the modulation frequency is high enough that the sidebands are not, we can assume that the sidebands are totally reflected, $F(\omega \pm \Omega) \approx -1$. This gives

$$P_{ref} \approx 2P_s - 4\sqrt{P_c P_s} \text{Im}[F(\omega)] \sin \Omega t.$$

I have neglected the reflected carrier, which will be proportional to $|F(\omega)|^2 \approx 0$. However, I do want to retain terms to first order in $F(\omega)$.

On resonance, ϕ is a multiple of 2π . Near resonance,

$$\phi \approx 2\pi N + 4\pi \frac{\delta L}{\lambda},$$

where δL is some small deviation in the cavity length from resonance, and N is an integer. If this deviation is only a small fraction of a wavelength (Yes, we really are sensitive to such small changes!),

we can write the reflection coefficient in terms of δL and approximate it by

$$F(\delta L) \approx \left[\frac{r}{1-r^2} \right] \left(i4\pi \frac{\delta L}{\lambda} \right) \\ \approx i \frac{4\mathcal{F}}{\lambda} \delta L.$$

Here \mathcal{F} is the finesse of the cavity. The reflected power near resonance is then

$$P_{ref} \approx 2P_s - 16 \sqrt{P_c P_s} \frac{\mathcal{F}}{\lambda} \delta L \sin \Omega t.$$

The error signal is now proportional to the change in the length of the cavity. We have entered the linear regime.

$$\varepsilon = -16 \sqrt{P_c P_s} \frac{\mathcal{F}}{\lambda} \delta L$$

We will need the slope of the error signal in the next section, so let's give it a name.

$$D \equiv 16 \sqrt{P_c P_s} \mathcal{F} / \lambda \tag{2}$$

4 Fundamental limits

4.1 Shot noise limited resolution

We use the error signal to measure the cavity noise by the relation $\delta L = -\varepsilon/D$. Expressed in terms of amplitude spectral densities, this is³

$$S_L = \frac{S_\varepsilon}{D}.$$

Any noise in the error signal itself will be indistinguishable from noise in the cavity. There is a fundamental limit to how quiet the error signal can be, even in the absence of length or frequency fluctuations, due to the quantum nature of light. The energy that falls on the photodiode comes in a pitter patter of photons, rather than a smooth flow, and the resulting noise is called shot noise.

On resonance, the reflected power, and hence the light power falling on the photodiode, will be $P_{ref} = 2P_s$. The shot noise in this signal has a flat spectrum with spectral density of [6]

$$S_\varepsilon = \sqrt{2 \frac{hc}{\lambda} (2P_s)}.$$

Dividing the error signal spectrum by D gives us the equivalent length noise,⁴

³The minus sign is essentially irrelevant for a spectrum.

⁴This is not the shot noise limit relevant to LIGO. There, they are concerned with the finite bandwidth of the cavity, which makes the shot noise limited length resolution turn up at high frequencies, instead of having a flat spectrum. The frequency at which this happens is so high for our shorter cavities that we can ignore it.

$$S_L = \frac{\sqrt{\hbar c}}{8} \frac{\sqrt{\lambda}}{\mathcal{F} \sqrt{P_c}}. \quad (3)$$

This is the limit to which you can resolve the length of the cavity. The same mechanism limits your ability to measure the frequency of the laser.

$$\begin{aligned} S_f &= \frac{f}{L} S_L \\ &= \frac{\sqrt{\hbar c^3}}{8} \frac{1}{\mathcal{F} L \sqrt{\lambda P_c}}. \end{aligned}$$

Since you can't resolve the frequency any better than this, you can't get it any more stable than this by using feedback to control the laser. Note that the shot noise limit does not explicitly depend on the power in the sidebands, as you might expect. It only depends on the power in the carrier.

It's worth putting in some numbers to get a feel for these limits. One cavity that we use in the TNI (the reference cavity for our pre-stabilized laser stage) is 20cm long and has a finesse of 10^4 . Our laser operates at 500mW with a wavelength of 1064nm . (Both of these are commercial items.) The best resolution we could possibly hope to achieve, were we to try and measure the length noise in this cavity with a perfect laser, would be

$$S_L = \left(8.1 \times 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}} \right) \frac{10^4}{\mathcal{F}} \sqrt{\frac{\lambda}{1064\text{nm}} \frac{500\text{mW}}{P_c}}.$$

If the cavity had no length noise and we were to lock the laser to it, the best frequency stability we could get would be

$$S_f = \left(1.2 \times 10^{-5} \frac{\text{Hz}}{\sqrt{\text{Hz}}} \right) \frac{10^4}{\mathcal{F}} \frac{20\text{cm}}{L} \sqrt{\frac{1064\text{nm}}{\lambda} \frac{500\text{mW}}{P_c}}.$$

4.2 The sensitivity theorem

Mizuno has estimated the quantum limited sensitivity of any interferometer from very general considerations, and his result is known as the sensitivity theorem [7]. Since shot noise is the quantum mechanism that limits the sensitivity of a Fabry-Perot cavity, our shot noise limit ought to give the same limit as Mizuno's sensitivity theorem.

The sensitivity theorem is usually stated in terms of strain, $S_h = S_L/L$. According to this theorem, quantum effects must limit your strain sensitivity to

$$S_h \geq \sqrt{\frac{2 \hbar \lambda \Delta f_{bw}}{\pi c \epsilon}},$$

where Δf_{bw} bandwidth of the interferometer, and ϵ is the energy stored in the light circulating in the interferometer. The bandwidth of a Fabry-Perot cavity is half its linewidth, $\Delta f_{bw} = c/(4\mathcal{F}L)$.

The energy stored in a Fabry-Perot cavity on resonance is the circulating power $P_{circ} \approx P_c \mathcal{F} / \pi$ [8] multiplied by the time it takes for light to circulate once, $2L/c$.

If we write the shot noise limit (Eqn. 3) in terms of these parameters, we get a strain resolution limit of

$$S_h = \sqrt{\frac{\hbar \lambda \Delta f_{bw}}{c \epsilon}},$$

which is consistent with the sensitivity theorem.

5 Technical considerations

5.1 Getting close to the shot noise limit

There are lots of other kinds of noise, of course, but most of these can be suppressed by feedback. Tim Day has done a nice analysis of technical noise sources in laser frequency stabilization [3], and the results are easy enough to adapt to cavity locking. The closed loop length noise in the cavity, in terms of the other noise sources, is

$$S_{L,cl} = \frac{\sqrt{S_{L,cavity}^2 + |K S_{v, servo}|^2 + |K G S_{v, pd}|^2}}{|1 + K G D_v|},$$

where $S_{L,cavity}$ is the free running length noise of the cavity, $S_{v, servo}$ is the voltage noise in the servo amplifier, and $S_{v, pd}$ is the voltage noise in the photodetector. The terms K , G , and D_v are the gains of the actuator, the servo amp, and the photodetector, respectively. The units of K are m/V , D_v has units of V/m , and G is unitless. The photodetector gain D_v is just the slope of the error signal D multiplied by the voltage response of the photodetector.

If you make the servo gain very large, $G \gg 1$, the noise spectrum limits out to

$$S_{L,cl} \approx \frac{S_{v, pd}}{D_v}.$$

The closed loop cavity noise is dominated by the noise in the photodetector. If the noise in the photodetector is mainly shot noise, then this reduces to the limit derived above.

If you want to get close to the shot noise limit, you need to concentrate on the photodetector (to reduce $S_{v, pd}$). High servo gain takes care of all the other noise sources. When you first look at it, it seems like magic that a high gain servo will suppress its own noise, but that's really the way it works!

5.2 Maximizing the slope of the error signal

If you want to minimize the closed loop noise, you want to maximize D . This won't really affect the shot noise limit, but it will help suppress technical noise sources. D depends on the cavity finesse, the laser wavelength, and the power in the sidebands and the carrier. Experimental details usually restrict your finesse and wavelength choices, but you often have quite a bit of freedom in adjusting

the sideband power. The question I want to address in this section is this: How does D (and hence the closed loop noise) depend on the sideband power?

D is proportional to the square root of the product of the sideband and carrier power. (See Eqn. 2.) This has a very simple form when $P_c + 2P_s \approx P_0$, *i.e.* when negligible power goes into the higher order sidebands.

$$D \propto \sqrt{P_c P_s}$$

$$\approx \sqrt{\frac{P_0}{2}} \sqrt{\left(1 - \frac{P_c}{P_0}\right) \frac{P_c}{P_0}}$$

A plot of D against P_c/P_0 traces out the top half of a circle, with a maximum at $P_c/P_0 = 1/2$. (See Fig. 4.) It is useful to express the power in the sidebands relative to the power in the carrier, and this gives

$$\frac{P_s}{P_c} = \frac{1}{2}.$$

D is maximized when the power in each sideband is half the power in the carrier, and this maximum is fairly broad.




Figure 4: An approximate plot of D , the slope of the error signal near resonance, versus P_s/P_c . The optimum value is at $P_s/P_c = 1/2$, and the maximum is very broad.

If you want to be more careful in this analysis, write D in terms of the Bessel functions of the modulation depth and maximize it. You'll find the optimum modulation depth to be $\beta = 1.08$, and you'll come up with essentially the same answer as with the simple estimate: $P_c/P_0 = 0.52$, $P_s/P_0 = 0.22$, and $P_s/P_c = 0.42$.

5.3 Noise in other parameters

The only two things the error signal is sensitive to are the length of the cavity and the frequency of the laser, and the sensitivity to each of these is comparable. It is not possible, in principle, to distinguish noise in the laser from noise in the cavity by just looking at the error signal. There are tricks to get around this, however. For example, you can split the beam and send it into two independent cavities. Any noise that is common to both will probably be in the laser, and you can eliminate it by doing a differential measurement on the two error signals. This is the sort of arrangement we will have in the TNI.

The system is not sensitive (to first order) to fluctuations in the power in the beam, the response of the photodiode, the modulation depth, and the relative phase of the driving and reference signals. The system will be sensitive to fluctuations in the sideband power near the modulation frequency, but most such noise is negligible at high frequencies. You don't have to worry about it if you choose your modulation frequency high enough. This insensitivity to a wide range of noise sources is characteristic of nulled lock-in detection, which is the general term for any modulated measurement operated around a zero output. If you want to get good at low noise measurements, the general theory of nulled lock-in detection is worth exploring!

A The reflection coefficient

The reflected beam is the coherent sum of two different beams: the promptly reflected beam, off the first mirror, and the leakage beam, from the cavity. It is not too hard, provided you keep track of your phases,⁵ to show that the total reflected beam is

$$F = E_{ref}/E_{inc} \\ = \frac{-r_1 + r_2 (r_1^2 + t_1^2) e^{i\Delta\nu_{fsr}}}{1 - r_1 r_2 e^{i\Delta\nu_{fsr}}},$$

where $\omega/2\pi$ is the laser frequency, and $\Delta\nu_{fsr}$ is the cavity free spectral range. The terms r_1 and r_2 are the amplitude reflection coefficients for the front and back mirrors, respectively, and t_1 is the transmission coefficient for the front mirror.

The amplitude of F follows an Airy function with respect to ϕ , and its properties are explored in any good undergraduate optics book. However, it is the phase of F that is important for understanding Pound-Drever-Hall locking, and that is not generally treated in textbooks. (The only treatment I am aware of right now is a partial one in Siegman [8].) Fortunately, there is a very simple way to think about the phase of F .

The reflection coefficient F traces out a circle in the complex plane with respect to the parameter ϕ . (See Fig. 5.) The amplitude of F is symmetric around resonance, but its phase is different depending on which side of resonance you are on. This phase is what the Pound-Drever-Hall technique samples

⁵There is some subtlety here! Physicists do it one way, and engineers apparently do it another. I am using the physicists' convention: real transmission and reflection coefficients, with a sign difference between reflecting off a surface from the right or the left.

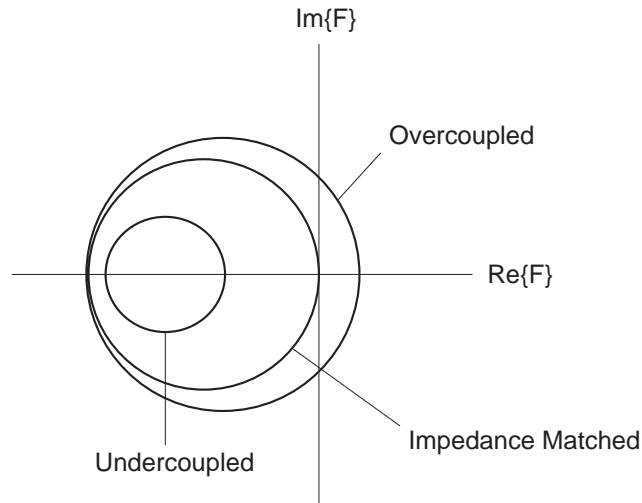


Figure 5: The reflection coefficient in the complex plane. As the laser frequency (or equivalently, the cavity length) increases, $F(\omega)$ traces out a circle (counterclockwise). Most of the time, F is near the real axis at the left edge of the circle. Only near resonance does the imaginary part of F become appreciable. Exactly on resonance, F is real, only this time it lies on the right edge of the circle.

by beating the reflected carrier with the sidebands. This is the real story behind the conceptual model illustrated in Fig. 2.

The leakage beam is generally 180° out of phase with the promptly reflected beam on resonance, so the interference between the two is destructive. If the leakage beam exactly cancels the promptly reflected beam, then the cavity is said to be impedance matched. If the leakage beam is too weak to entirely cancel the promptly reflected beam, the cavity is called undercoupled. If the leakage beam is much stronger than the promptly reflected beam, the cavity is said to be overcoupled. LIGO's cavities are strongly overcoupled, which allows them to get a reflected beam from their Fabry-Perot cavities even when they are on resonance. This is essential if you want to use the cavities in the arms of a Michelson interferometer.

References

- [1] Friedland, B., *Control System Design: An Introduction to State Space Methods*, (McGraw Hill, 1986).
- [2] Franklin, G. F., J. D. Powell, and A. Emami-Naeni, *Feedback Control of Dynamic Systems*, (Addison-Wesley, Reading, Massachusetts, 1987).
- [3] Day, T., *Frequency Stabilized Solid State Lasers for Coherent Optical Communications*, (Ph.D. thesis, Stanford University, 1990).
- [4] Drever, R. W. P., *et al.*, Appl. Phys. B 31, 97 (1983).
- [5] Bjorklund, G. C., *et al.*, Appl. Phys. B 32, 145 (1983).
- [6] Saulson, P. R., *Fundamentals of Interferometric Gravitational Wave Detectors*, World Scientific (1994).
- [7] Mizuno, J., *Comparison of optical configurations for laser-interferometric gravitational-wave detectors*, (Ph.D. thesis, Max-Planck-Institute für Quantenoptik, 1995).
- [8] Siegman, A. E., *Lasers*, (University Science Books, Mill Valley, California, 1986).