

1. STABILIZING DIODE LASERS TO HIGH-FINESSE CAVITIES

Richard W. Fox, Chris W. Oates, Leo W. Hollberg

National Institute of Standards and Technology, Boulder, Colorado, USA

1.1 Introduction

This chapter is written for the chemist, physicist, or engineer who is interested in locking a diode laser to an optical cavity. There already exist many good references on the locking of lasers [1, 2], including some on diode lasers in particular [3, 4]. In this chapter we aim to provide a practical hands-on guide, with most or all of the theory left to cited references.

The motivations for locking lasers to high-finesse optical cavities include applications as varied as laser cooling, length metrology, and analytical and precision spectroscopy. In some cases it may be the stabilization of the laser's frequency to a narrow cavity resonance that is of interest, while in others it may be the huge build up of intracavity power or the long effective path length between the mirrors that is important. The case of locking diode lasers is of particular interest because of their unique characteristics, relatively low cost and widespread use in many applications.

Before proceeding further, we should offer a note of caution that the term "diode laser" now encompasses many devices whose electrical and optical characteristics differ greatly. Included in this category are the Fabry-Pérot laser, vertical-cavity surface-emitting lasers, and distributed feedback lasers. While the basic approach for locking these different types of lasers is the same, the performance of the lock and the details in the electronic feedback loops will be quite different. Rather than overwhelm the reader with details for all these different types, we instead focus on one type in particular, the widely used extended-cavity Fabry-Pérot laser. For those interested in locking other types of diodes to cavities, this chapter should still be a good place to start.

This chapter consists of the following. We begin with a general discussion of the issues involved in diode laser locking and introduce the reader to some of the terminology. We then describe in detail the various steps needed to lock the laser to a cavity resonance: (1) derivation of the error (locking) signal, (2) design of the electronic feedback circuitry, (3) initial locking of the laser, (4) adjustment of the feedback design, and (5) evaluation of the lock performance. We illustrate this discussion by frequency locking an

extended-cavity diode laser, reducing the linewidth to a few hertz relative to the cavity. We conclude with an example in which we modified the locking apparatus for a cavity ring-down demonstration. Included are results showing the laser repetitively locking and unlocking to the cavity.

1.2 Introduction to Diode Laser Locking

Regardless of the application, the basic goal of locking the frequency of a laser to a cavity is to reduce the frequency fluctuations between the laser and cavity. The noise spectrum of the laser's frequency fluctuations leads to an effective "linewidth" of the laser, which conceptually describes the broadening of the laser's spectrum around its central frequency. The concept of laser linewidth can be a little confusing but is valuable, so it is worth a brief discussion here. The basic confusion arises from the fact that the laser linewidth is dependent on the timescale over which it is evaluated, as a laser's noise spectrum typically contains very different fast and slow components. A useful intuitive picture is that the fast fluctuations give rise to what is called the "fast linewidth" (or sometimes just "linewidth"); by "fast" we usually mean in a time less than the effective spectroscopic interaction time (typically tens or hundreds of microseconds). With this definition, the linewidth defines the narrowest feature that the laser is capable of resolving. The low-frequency fluctuations (usually of larger magnitude) then cause jitter of this narrow spectral line in frequency space. Even slower changes often occur because of thermal effects, which can cause the laser's central frequency to drift.

In order to reduce the laser linewidth, one needs a stable frequency reference suitable for measuring the laser's frequency fluctuations. One can then construct a feedback loop that attempts to compensate these fluctuations. A Fabry-Pérot cavity, which usually consists of two high-reflectivity mirrors separated by a spacer, is a convenient choice for a frequency reference. With mirror reflectivities $> 99\%$ over a given range of wavelengths, a cavity has a series of evenly spaced, sharp resonances (typically with linewidths < 10 MHz). The resonance or "fringe" separation is $c/2L$, where c is the speed of light and L the length of the optical cavity (i.e. the distance between the mirrors). Radio-frequency modulation techniques can be used to derive an electronic error signal that represents the deviations of the laser frequency from a given cavity reference fringe. One then uses electronic feedback (primarily to the laser current in the case of diode lasers) to control the laser frequency and minimize its deviations relative to the cavity fringe.

If one succeeds in making the electronic feedback bandwidth (i.e. range of frequencies over which corrections are effectively imposed) wider than the dominant noise spectrum of the laser, the laser's fluctuations can be controlled such that the laser appears tightly locked to the peak of the cavity mode. In this case, the linewidth relative to the cavity's center frequency can be a small fraction ($\approx 10^{-4}$) of the cavity's linewidth. (In this case, the actual laser linewidth will almost certainly be determined by the mechanical stability of the cavity itself.) Likewise, even though the power is not actively stabilized, the power transmitted from the cavity can be stable to better than 1%. On the other hand, if the electronic feedback stabilizes only the low-frequency portion of the laser's fluctuations, the center frequency will be fixed, but no linewidth narrowing will occur and relatively little of the laser power will be coupled into the mode. Thus we see that there are two basic determining factors in the final performance of the lock: the initial noise spectrum of the laser and the electronic bandwidth achievable in the feedback or "servo" system.

The short cavity lengths and low facet reflectivity found in most diode structures result in "low- Q " cavities, and generally the optical output will have a minimum spectral linewidth of 10 MHz, and often much larger. Thus, if we expect to narrow the linewidths of these lasers significantly, we will need feedback loops with sufficient gain up to 10 MHz and beyond. However in most cases (for most commercial diodes), a feedback bandwidth of 10 MHz is quite difficult to achieve because of the characteristics of the diode laser itself. At low modulation frequencies, the frequency shift with injection current is primarily caused by thermal effects. That is, because the injection current causes heating and the index of refraction is temperature dependent, we find that the subsequent frequency modulation (df/dI) is negative. This effect diminishes with modulation frequency, and at times on the order of 1 μ s, another smaller effect, caused by the electronic charge carriers, starts to dominate. For this case we have the opposite dependence, i.e., df/dI is positive. The resulting net phase response of the diode is troublesome for the servo designer, making electronic feedback bandwidths beyond a few megahertz increasingly difficult to achieve.

To achieve tight locking, it is therefore often necessary to reduce the frequency fluctuations of the laser by other means before locking it to the cavity. A common technique is to increase the optical cavity's Q by adding optics external to the laser chip in order to form a longer cavity. A laser cavity built in this manner is often called an ECDL, "extended-cavity diode laser," or "external-cavity diode laser." The use of frequency-dependence reflection for the external cavity (e.g. with a diffraction grating) can offer the additional benefits of frequency pulling and reduced mode competition. There are many extended-cavity lasers available commercially, and several

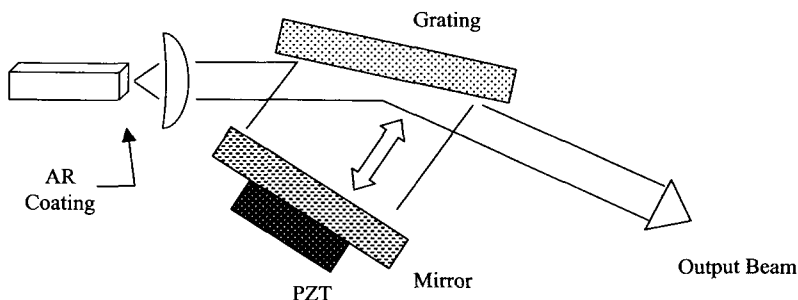
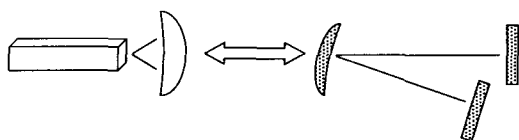


FIG. 1.1. A diode laser arranged in a Littman-Metcalf configuration. The beam is incident at grazing incidence on the diffraction grating, and the first order is reflected back by a mirror. The zero order serves as the laser output. Coarse wavelength tuning can be done by adjusting the mirror angle, and fine-tuning by using a PZT to adjust the cavity length. However, true single-mode scanning over a large range requires that the laser cavity length be changed synchronously with the mirror angle.

good references for their design and construction [5, 6]. The laser used as an example here is an off-the-shelf commercial diode that was antireflection coated and built into a Littman-type extended cavity about 8 cm long (see Fig. 1.1). The extended cavity is formed with a diffraction grating (for wavelength-selective optical feedback) and a high-reflectivity mirror mounted on a piezo-electric transducer (PZT). The laser operates around 830 nm and yields an output power of 6 mW at an injection current of 70 mA. We can tune the laser frequency quickly over small excursions with the laser current, or slowly over larger excursions with the voltage on the PZT. The resulting fast linewidth for this laser is roughly 50 kHz, while there are low-frequency vibrations of the mechanical structure (caused by room acoustics), which cause the laser output to jitter in frequency on the order of a megahertz. We note that in many cases, one must use an off-the-shelf commercial diode (without an additional antireflection coating), which usually requires that a shorter extended cavity be used to reduce laser mode competition. Unfortunately, shorter extended cavities yield less frequency noise reduction, leaving more work to be done by the electronic feedback system. Nonetheless, the techniques described in this chapter should still be directly applicable to this case as well and provide good locking performance [7].

We should also add that while the use of electronic feedback is convenient, it is not the only option for stabilizing the diode laser's frequency to a cavity. The use of resonant optical feedback directly from a cavity is a well-established technique that can be used to stabilize even a simple Fabry-Pérot diode to certain cavities [8–10]. A requirement is that the

(a)



(b)

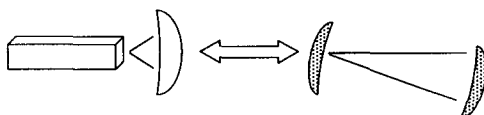


FIG. 1.2. Two configurations to optically lock a diode laser to an optical cavity. The cavity must return light to the laser when on resonance, while reflecting away the off-resonance beam. Diagram (a) is a three-mirror V-shaped cavity, while diagram (b) is a confocal optical cavity aligned in a V-configuration. For stable operation, the path length from the laser to the cavity must be servo controlled. Not shown are mirrors and a coupling lens to align the laser to the cavity.

optical alignment be such that light is reflected or returned to the laser diode only when the cavity is in resonance. Two possible configurations include a three-mirror V-shaped cavity, and a confocal cavity aligned in the $c/4L$ configuration (see Fig. 1.2). Additionally, it is necessary to electronically control the phase of the light returning to the diode laser, which can be accomplished with a mirror mounted on a PZT [11].

1.3 Generating the Error Signal: The Pound–Drever–Hall Method

The first step in the locking procedure is to generate an electronic error signal that can be further processed for locking the laser to the cavity. Perhaps the two most convenient approaches for generating this signal are *side-locking* and the *Pound–Drever–Hall (PDH) method* [1]. With the simpler technique, side-locking, the laser is stabilized to the side of the cavity fringe without the use of modulation techniques. While this technique is straightforward to implement and has some advantages (e.g. simplicity, modulation-free), it suffers from several serious disadvantages. First, because it is modulation-free, one necessarily detects the error signal

at DC, where there can be significant amplitude noise. Second, side-locking has a much smaller acquisition range, which means the laser and cavity frequencies need to be nearly coincident before the system will lock. Third, the lock is less robust because perturbations (e.g. vibrations) that drive the laser to the opposite side of the resonance will cause the feedback loop to push the laser further from resonance. Fourth, because one locks to the side of the resonance rather than the top, there is reduced buildup of optical power in the cavity, and there may be increased noise on the transmitted intensity.

The PDH technique circumvents most of these drawbacks by modulating the frequency of the light, enabling detection of the error signal at a high frequency where the technical noise is near the shot-noise limit. The resulting demodulated error signal has a high signal-to-noise ratio and a large acquisition range, which can produce robust locks. Furthermore, this error signal has odd symmetry about the line center that enables locking to the top of a cavity fringe. For these reasons, we prefer the PDH method and focus on its implementation throughout the remainder of this chapter.

The PDH technique is described in detail in many references [1–4], but we will review the general approach here. Consider a purely frequency modulated (FM) laser beam impinging on the input mirror of an optical cavity and reflecting back to a detector. For low modulation index (as we typically use), one can view the frequency spectrum of the modulated light as consisting of a carrier with two sidebands: one at higher frequency with a phase relative to the carrier that is in phase with the modulation, and one at lower frequency that is out of phase by 180° . As long as there is no absorption or phase shift of the laser carrier or modulation sidebands with respect to one another, the detector photocurrent will not have a signal at the modulation frequency. (Only the optical phase is modulated, not the optical power.) A simple view of this fact is that the beating between the carrier and the upper frequency sidebands creates a photocurrent modulation that is exactly canceled by the out-of-phase modulation from the lower frequency side. If a sideband is attenuated or phase shifted, or the carrier's phase is shifted, the photocurrents will not cancel and RF power at the modulation frequency will appear on the detector signal.

Near a cavity resonance, the resultant optical reflection of the carrier from the cavity is phase shifted with respect to the sideband components that are further away from the cavity resonance. Consequently, the detector photocurrent will show power at the modulation frequency. The laser's frequency noise will then appear as noise sidebands centered around the modulation frequency. When this signal is mixed to base-band (using phase-sensitive detection with the appropriately chosen phase), the result is

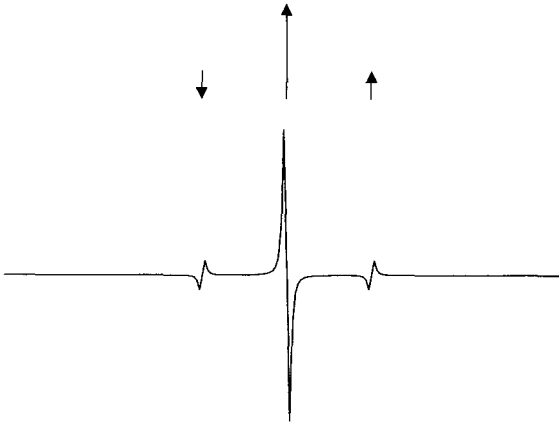


FIG. 1.3. A schematic representation of the Pound-Drever-Hall error signal as if the laser was scanned over a cavity resonance. The cavity reflection is detected and demodulated by a mixer using a phase-delayed local oscillator. Proper adjustment of the phase delay results in a sharp central frequency discriminator, with secondary peaks spaced by the modulation frequency.

a frequency discriminator with odd symmetry that may be used to correct the frequency of the laser as shown schematically in Figure 1.3. Here we note that the light seen by the detector actually consists of two components: the fraction of the input beam that is reflected, plus the fraction of the internal cavity wave that is transmitted back out of the input coupler. The detected photocurrent represents the interference of these two components. For Fourier components of the laser's frequency noise below the cavity linewidth, this system acts like a frequency discriminator as described above. At frequencies above the cavity linewidth, the input field is essentially heterodyned with the cavity wave. Thus, for these Fourier components, the system acts as a phase discriminator, which causes the system response to faster frequency fluctuations to decrease as $1/f$. In electronic terms, the PDH technique gives us a frequency error signal with a sensitivity that can be measured in volts per hertz of optical frequency. At the Fourier frequency corresponding to the cavity linewidth, the sensitivity starts to decrease, and continues to decrease as $1/f$. At some higher frequency, the error signal will cease to be useful as the magnitude decreases to the level of the background noise, although this limit is usually well above the attainable servo bandwidth.

For those who are perhaps puzzled by some of terminology above, recall that phase-sensitive detection is the process that a simple lock-in amplifier uses, although most commercial lock-in devices do not operate at the

frequencies required here (10–50 MHz). In this chapter we show how to build the necessary circuitry using RF amplifiers and filters, a simple mixer, and subsequent operational amplifier (op-amp) filtering and amplification.

Our experimental layout is shown in Figure 1.4. The light goes from our ECDL through an optical isolator and electro-optic modulator to the optics used for coupling to the optical cavity. While this setup would be suitable for most types of lasers, let us emphasize a couple of aspects particular to diode lasers. Because diode lasers are extremely sensitive to optical feedback, good optical isolation (> 50 dB) is necessary to attenuate the light reflected from the cavity. Additionally, the laser's spatial mode is not round, but is instead oval because the laser's output aperture is asymmetric. Although a spherical collimating lens is often used at the diode output, correcting the beam's aspect ratio with anamorphic prisms or cylindrical lenses would result in better spatial matching to the cavity mode. This would in turn increase the power coupled to the cavity. In the examples here, a single lens and an adjustable aperture are used for spatial mode-matching to the cavity, and two alignment mirrors are provided between the lens and the cavity. The cavity reflection is returned to the photodetector using a quarter-wave plate along with a polarizing beam splitter. The detector signal is amplified and filtered to pass the modulation frequency, and phase-sensitively demodulated using a balanced mixer. The resulting signal is then

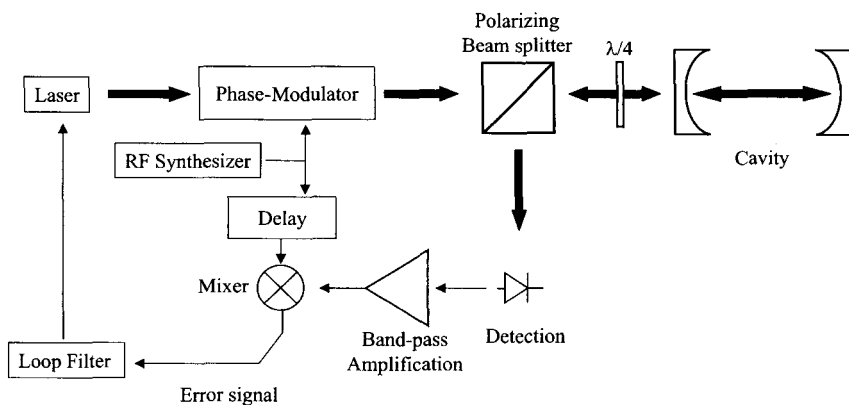


FIG. 1.4. A basic Pound–Drever–Hall locking scheme. A modulation frequency in the range 10–40 MHz is typically used. The phase of the RF local oscillator input to the mixer is adjusted to give a symmetrical frequency discriminator at the mixer output. In the absence of residual amplitude modulation or RF pickup, the center of the cavity resonance will correspond to zero at the output. The error signal must be properly filtered and amplified to control the laser.

amplified and filtered as desired for the feedback loop, and finally sent to the diode laser frequency control elements.

Let us now see how to design and construct these various pieces, starting first with the optics before moving on to the electronics.

1.3.1 Coupling Light into the Cavity

In this section, we describe how to choose and align the optics used for coupling the light from the diode laser into the optical cavity. The optimization of this procedure requires consideration of two distinct and important concepts, *mode-matching* and *optical impedance-matching*. Mode-matching refers to adjustment of the input beam's size, shape, and wavefront curvature to match the cavity mode. Impedance-matching refers to the adjustment of the cavity parameters (e.g. mirror reflectivities) to maximize coupling into the cavity. This occurs when the transmission through the input mirror equals all the other losses of the resonator. For instance, a two-mirror cavity would be impedance-matched if the loss from the input mirror equaled the sum of the second mirror's loss plus the loss resulting from scatter and absorption on the mirrors and anything else in the cavity. If this ideal condition is met, the reflection of a perfectly spatially matched input beam will destructively interfere completely with the cavity wave transmitted back through the input coupler. Thus, no net power will be reflected off the input mirror on resonance in steady state. Usually this condition is not met, and even a perfectly spatially matched beam will not couple all the power into the cavity. However, coupling efficiencies of 80% are not uncommon. Instead, the reflected signal from the cavity on resonance will show a dip to a minimum, but not a dip to zero. In some applications, however, such as ring-down spectroscopy, impedance-matching may be less important than some other consideration such as maximizing the ring-down time. In such a case, one might minimize the loss at each mirror, for example, even if it means that less power is coupled.

Matching the spatial mode of a diode laser's beam to that of a spherical mirror cavity is a challenge because the laser's output beam is neither round nor "diffraction limited." Conceptually, think of a beam emanating from a waist in the cavity towards the laser, and let us call this the "cavity wave." The beam waist and radius of curvature of the cavity wave may be calculated from the mirror geometry [12]. The mode-matching task is to shape and focus the diode beam such that it approaches the conjugate of this cavity wave, i.e., a wave that is exactly like it except moving towards the cavity. Typically we start with a beam that is collimated with an additional lens, if necessary, after it emerges from the optical isolator or electro-optic

modulator. A first step is often then to make the collimated diode beam's cross section somewhat round by using a pair of anamorphic prisms. Next, at the fictitious point where the diode beam and the cavity wave are the same size, a plano-convex lens of the proper power will focus the diode beam to a waist at the same position as the cavity's waist. If the collimated diode beam is larger than the cavity wave, use either a collimating lens of shorter focal length on the diode, or a multiple lens solution to reduce the beam's cross section. For instance, a positive lens can be used to create a converging diode beam that at some point will have the same cross-sectional size as the cavity wave. A negative lens of the proper focal length at this point will reduce the wavefront curvature, and cause the diode beam to reach a waist at the same position as the cavity waist. See Reference [12] for more details on mode-matching.

A brief description on exactly how to align the diode beam to a high-finesse, two-mirror Fabry-Pérot cavity is warranted here. Much of this discussion applies to ring cavities as well. An adjustable diaphragm prior to the lens is useful, both for alignment and for reduction of the power coupled to higher order modes. It is easiest to align the diode beam to the cavity if there are two adjustable turning mirrors of good quality between the cavity and the focusing optics. Adjust the two mirrors such that the input beam is incident at the center of the cavity mirror, and also such that the cavity reflection hits the turning mirrors at the same points as the input diode beam. Observe the cavity reflection on the diaphragm; if the beam size is approximately the same size as the input beam at the diaphragm, then input wavefront curvature is well chosen. If not, either the input diode beam's radius of curvature is incorrect, or the position of the waist is not near the cavity wave's waist. Careful calculation or some experimentation should enable one to get close.

Now, with a detector placed just after the cavity, one can monitor the cavity transmission on an oscilloscope, while sweeping the laser frequency slowly (≈ 30 Hz) with a triangle drive signal fed to the PZT-controlled mirror. Alternatively, the cavity length can be swept if one of the cavity mirrors is mounted on a PZT. Once modes are observed, the laser's wavelength sweep should be adjusted so that about two cavity free-spectral-ranges are covered. Fine adjustment of the alignment (and mode-matching) should then lead to good coupling to the fundamental TEM_{00} mode. One common problem that arises is to figure out which peak actually corresponds to the TEM_{00} mode. To this end we offer several suggestions. Closing the adjustable aperture to a small size will attenuate the coupling to other modes more than the TEM_{00} mode, because it is spatially symmetric and the smallest mode. Also, all the mode peaks except for the TEM_{00} mode are actually two or more modes with approximately the same resonance

frequencies. This degeneracy is broken in high-finesse cavities by any departure of the mirror surfaces from a spherical shape [13]. This leads to a fine splitting of all the higher order transverse modes that is easily seen by zooming in on an unknown mode by reducing the laser sweep and increasing the oscilloscope gain. The fundamental mode does not have this degeneracy and will appear as a single peak. Observation of the transmitted spatial mode profile by eye or with a video camera after the cavity can help to optimize alignment into the TEM_{00} mode as well. With some patience, one can usually get to the point where the TEM_{00} peak is a factor of five or ten times larger than any of the competing modes.

When working with very high-finesse cavities, one may notice that the peak heights vary drastically from sweep to sweep. Sweeping more slowly may increase the heights of the peaks, but may or may not reduce the variation. These effects result from the laser frequency not staying in resonance long enough (i.e. the cavity decay time) to reach equilibrium (recall that the cavity buildup time is long for high-finesse cavities). Alignment under these conditions requires patience, although we will see later on that when the servo system is operating, one can easily optimize the alignment by working with the laser locked to the cavity fringe.

We also note that mirror contamination can affect the loss of the fundamental mode relative to higher order modes, especially the first odd transverse mode. It is well known that the mirrors of high-power dye laser systems will become contaminated by material deposited on the mirror surface just where the beam power is maximum. With good low-loss mirrors, the ring-down cavity power can be 10 W continuous-wave, even for milliwatt input powers, so one can expect this problem occasionally if the cavity is operated with high power. A characteristic sign of this occurring is a change in the spatial coupling such that the TEM_{00} mode power is degraded in relation to the transverse modes.

1.3.2 Modulating the Laser Frequency

There are several competing factors to consider when choosing the modulation frequency for the PDH lock. The frequency needs to be sufficiently high so that the process of filtering and demodulating to base-band does not yield a significant phase shift within the desired feedback bandwidth. We find that choosing a modulation frequency an order of magnitude higher than the required bandwidth is more than sufficient. Although modulating at even higher frequencies (hundreds of megahertz) is possible, this requires more skill to avoid distorting the DC baseline because of RF pickup. It also makes the local-oscillator phase more

susceptible to variations through changes in temperature of the cables, and more susceptible to residual amplitude modulation (AM) offset errors associated with variations in the optical path length. In practice, we find that modulation frequencies ranging from 15 MHz to 40 MHz work well.

To modulate the laser frequency, we either modulate the laser's injection current directly or use an electro-optic modulator (EOM) placed between the laser and the cavity. Of these two schemes, using an EOM is preferable because it yields much less residual AM. Residual AM can be quite troublesome for the locking process, as it leads to DC offsets after demodulation, which in turn shift the lock point away from the peak of the cavity transmission. In principle such offsets can be compensated, but often the factors responsible for the residual AM are unstable, and lead to a drifting offset that is more difficult to null. Thus, we use an EOM when frequency drift relative to the cavity is important, but employ direct-current modulation for less critical locks.

Commercial EOM units are generally configured either as resonant circuits or as broadband modulators. In order to generate sidebands of sufficient amplitude (5–10% of the carrier peak intensity), a resonant configuration makes sense because it reduces the required RF power by about a factor of ten. The addition of a low loss inductor and capacitor external to a broadband EOM (as shown in Fig. 1.5) works well and Q 's > 10 are easily achieved. In order to know what fraction of the light is coupled into the sidebands by the modulator, one can couple light into the cavity and monitor the transmission spectrum. The sidebands can then be observed by sweeping the laser's frequency over a cavity resonance, or alternatively by sweeping the cavity if its length is adjustable with a PZT. This signal can then be used to adjust the RF power sent to the EOM in order to generate sidebands of the desired amplitude.

Of course, an EOM adds considerable cost to the system, and it is often possible to lock by modulating the laser directly. The success of this approach depends on the magnitude of the amplitude modulation that accompanies the desired frequency modulation. This in turn depends on the laser design, specifically on parameters such as the length of the extended cavity and the output coupling of the laser. Unfortunately the same process that serves to reduce the laser's frequency noise and emission linewidth relative to the bare diode laser also reduces the laser's frequency deviation for a given modulation current. Consequently, modulating the injection current of a laser with a long extended cavity and a high reflectivity grating will produce large AM sidebands, with very little associated frequency modulation. One can check the modulation characteristics by aligning the laser to an optical cavity and observing the amplitudes of the induced sidebands. If direct modulation produces obviously asymmetric sidebands,

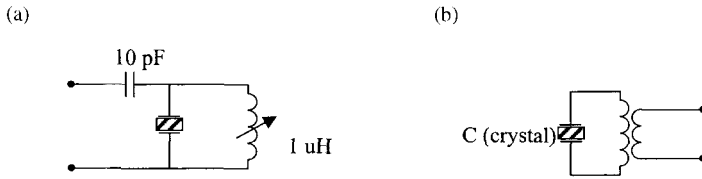


FIG. 1.5. A broadband nonresonant electro-optic modulator can produce sidebands adequate for locking if it is driven resonantly. With an inductor mounted externally but close to the modulator, the tank circuit shown will resonate at $2\pi f \approx (LC)^{-1/2}$. In configuration (a), an additional external capacitor is necessary for coupling (an air-spaced adjustable capacitor works well). Tuning the inductor will affect primarily the resonance frequency, while tuning the coupling capacitor will change the strength of coupling and hence the input impedance. Stronger sidebands will be obtained with configuration (b), however, this requires that neither side of the modulator crystal be grounded. The highest Q will be obtained by using nonadjustable air-core inductors and tuning the modulation frequency into resonance.

this is an indication that the relative amounts of amplitude and frequency modulation are comparable, and that the degree of frequency modulation is sufficient to create an error signal for locking. The asymmetry of the sidebands is caused by the AM and FM components being in-phase on one side of the carrier and out-of-phase on the other side. For an ECDL, if the direct modulation sidebands are symmetrical, it is possible that they are primarily AM sidebands, and the mixer output will have a DC component that is large with respect to the error signal. The amount of AM is also easily detected by a high-speed photodiode.

Direct modulation can be accomplished by coupling an RF signal to the injection current with a small capacitor (≤ 50 pF), as shown in Figure 1.6. The standard electrostatic precautions and warnings apply to this current input terminal. For instance, one needs to ensure that no current pulse is delivered to the laser when connecting a cable or turning on the RF function generator. When in doubt, the safe practice would be to connect the function generator while it is on, but with the output at zero.

Regardless of the modulation technique, the amplitude of the sidebands imposed on the laser is not critical, and it is possible to trade some modulation depth for gain elsewhere in the loop. A good starting point is to set the sideband amplitude at approximately 5–10% of the carrier. If the modulation sidebands are too weak, then the noise floor of the resulting system will be higher than necessary, and the resulting laser linewidth relative to the cavity will be higher than necessary. Also, the acquisition range will be poor, as the wings of the PDH error signal will be degraded to the noise level. If the modulation sidebands are too large, the power in the carrier and the subsequent cavity power buildup and transmittance will suffer.

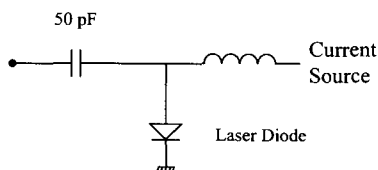


FIG. 1.6. A simple circuit to apply RF modulation to the laser through a coupling capacitor.

1.3.3 Detecting the Reflected Light

Light reflected off the input mirror of the high-finesse cavity can be picked off with a simple beam splitter, or a polarizing beam splitter and either a quarter-wave plate or a Faraday rotator. The latter choices use the light more efficiently, while the former is cheaper and reduces the number of optics in the beam path (which can cause residual AM). Attention must be paid to multiple reflections in the path between the cavity and the detector, because they can lead to étalon effects that are all too familiar to spectroscopists. Étalon effects usually manifest themselves as ripples in the mixer output as the laser is swept in wavelength. These baseline ripples can be quite large and seriously compromise the locking performance. Thus, it is important to take some care with the optics to avoid normal reflections that can lead to unintentional étalons. Slight tilting of the detector will avoid sending the reflection from the detector surface back towards the cavity. Likewise, defocusing or using a longer focal length lens to capture the light on the detector will result in less scattered light being reflected back towards the cavity. The detector monitoring the cavity reflection should be a low-noise detector with a reasonable bandwidth and also have a low-frequency DC output. Although a simple photodiode and a load resistor may be used to capture the reflected signal, the signal size will be limited because the load resistor needs to be relatively small in order for the detector to have sufficient bandwidth. The signal then must be routed through an RF amplifier to reach levels optimum for the mixer input (see Fig. 1.4). Serious consideration should be given to using a transimpedance amplifier to convert the photocurrent output to a voltage. There are wide-band transimpedance amplifiers available with built-in feedback resistors and impressively low input noise. Although not necessary for locking purposes, the DC output is useful for monitoring the laser power and to diagnose laser-cavity coupling problems.

The low-noise aspect of the detector is important if the reflected light power from the cavity is expected to be low. To have optimum signal-to-noise for the error signal, the photocurrent fluctuations must be larger than the detector noise. The amount of light power needed for this to be the case

can be calculated assuming the photocurrent fluctuations are attributable to shot noise. Of course if the photocurrent noise is actually higher than the shot level, the detector noise level will be exceeded with even less light than this calculation indicates. With no light incident, the detector–amplifier combination will exhibit some output noise density at the modulation frequency, measured in $V/\sqrt{\text{Hz}}$ which can easily be measured by an RF spectrum analyzer. One can calculate the power P_0 that would be necessary for the photocurrent noise to exceed the detector noise by setting this background noise level equal to the shot noise,

$$G\sqrt{2qP_0R}$$

In this expression G is the transimpedance gain in V/A , q is the electronic charge, 1.6×10^{-19} C, and R is the detector's responsivity in A/W . We note that there are commercial photoreceivers sold as “low noise” that will require more than 1 mW of power on the detector before the photocurrent shot level is larger than the detector noise. However, there are also transimpedance amplifiers with more than 100 MHz of bandwidth that require less than 50 μ W of power before the light noise is larger than the detector's noise.

1.3.4 Demodulation

Once a satisfactory signal is obtained from the detector, it is time to demodulate the error signal so that it will be ready to be used in the servo system. The effect of frequency modulation was to encode the laser frequency fluctuations around the modulation frequency (rather than DC); now it is time to “decode” this error information by mixing it down to DC. This demodulation process is easily accomplished by combining our error signal from the detector with a phase-shifted version of the local-oscillator signal used to generate the modulation in a doubly balanced mixer. We note that it can be beneficial to add an RF band-pass filter between the detector and the mixer, as signals outside of the anticipated servo bandwidth only add excess noise at the mixer input. However, care needs to be taken that the band-pass filter is not too narrow, as this will lead to phase shifts that will limit the servo bandwidth. Unless there is significant noise outside the servo bandwidth, a good starting point is to use no filter prior to the mixer.

For good mixer performance, the local-oscillator (LO) signal power should be enough to fully turn on the mixer diodes. (While mixers are available for a number of different power levels, we find that those requiring an LO power of +7 dBm to be a convenient choice.) The size of the resonance signal from the detector into the mixer's RF port should be within

a few decibels of the maximum given in the manufacturer's specifications (typically the RF port level will be a maximum of +1 dBm for a +7 dBm LO). If the output of the detector is much less than this, a low-noise RF amplifier can be inserted before the mixer's RF port. The RF signal, which is centered around the modulation frequency, is down-converted by the mixer to base-band (i.e. DC), with the resultant output accessible at the mixer IF port. This is done in a coherent manner, and with some conversion efficiency, typically about -6 dB. The IF port output also contains the same information at twice the modulation frequency, as a doubly balanced mixer converts the input signal at f_{RF} to $f_{\text{LO}} - f_{\text{RF}}$ and $f_{\text{LO}} + f_{\text{RF}}$. In addition, if the mixer IF port output is observed with a spectrum analyzer, a smaller but often significant narrowband component at f_{LO} will be found. These high-frequency outputs of the mixer can be sufficiently attenuated by the loop filter, which as we shall see is actually a low-pass filter. However, it is possible that these residual RF signals can generate a small DC signal through rectification by any nonlinearity in the circuit. This can wreak havoc with the frequency stability of a system that was designed to closely track the laser to the center of the cavity. It thus makes sense to attenuate the strong high-frequency components with resonant LC notch filters at the mixer output. The resonances should be of sufficient Q to avoid unnecessary phase shifts within the desired servo bandwidth [14]. In Figure 1.7, we show circuit components to reduce f and $2f$ (30 MHz and 60 MHz) while introducing only $1^\circ/\text{MHz}$ of phase shift in the region from 1 MHz to 10 MHz.

The lineshapes one detects at the output of the mixer IF port depend on the relative phase between the RF signal and the LO. In order to generate a frequency discriminator with odd symmetry and maximum amplitude, it is essential to set this relative phase to be dispersion-sensitive (i.e. 90° out-of-phase). The phase of the LO fed to the mixer is easily adjusted with a phase-shifting IC, a delay box, or simple lengths of cable. One could instead adjust the phase of the RF signal, but this may lead to undesired delay or phase shifts in the servo loop. To set the LO phase to the ideal value, it is easiest to start by using an oscilloscope to monitor the signal from the IF port while sweeping the laser or cavity over a cavity resonance. A low-pass filter (e.g. with a corner at ≈ 100 kHz) will attenuate the $2f$ signal and help to see the features more clearly. One then adjusts the LO phase and looks for the appropriate discriminator line shape (see Fig. 1.3). There will in fact be two such "ideal" settings; these yield mirror images of the desired discriminator signal (switching between these two settings gives an easy way to change the overall sign of the feedback signal). If the laser frequency jumps around too much to find the phase setting that optimizes the discriminator amplitude, an alternative approach is to first adjust the phase to minimize the central feature. Then shift the phase by 90° by adding or

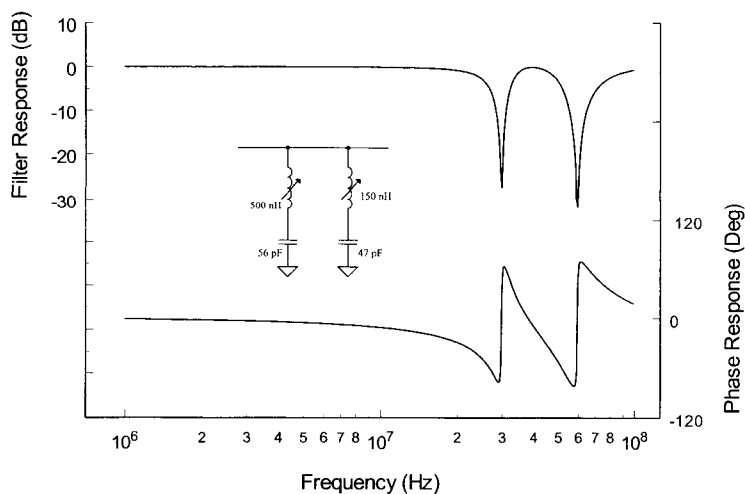


FIG. 1.7. Series LC traps at the mixer IF port can reduce the high-frequency components while adding very little phase shift within the servo bandwidth. The curves show measured data for a two-notch filter that adds only $1^\circ/\text{MHz}$ of phase shift in the region from 1–10 MHz. Once the filter is installed in the circuit, tunable elements are used to set the frequency.

subtracting the appropriate time delay, which equals one-quarter of the modulation period. The signal will appear larger as the laser is swept past the cavity resonance more slowly, because there is more time for the light to build up in the cavity mode. The error signals for a system using EOM modulation and one using injection-current modulation are shown in Figure 1.8. Note that the error signal observed when current modulation is employed contains a large DC component. This is attributable to the strong intensity modulation (residual AM) that must be compensated at the integrator input. Subsequent changes of the optical system (alignment for instance) will change this DC level, and require some adjustment of the offset compensation.

The optimum locking will be obtained when the highest possible signal-to-noise (S/N) exists on the error signal. The error signal's discriminator will have a linear slope measured in units of volts per laser frequency, for instance V/kHz. The factors that determine this slope include the cavity's linewidth, the sideband amplitude, the power on the detector, detector gain (V/A), postdetector amplification, and the mixing process. The noise on this discriminator should be attributable to the laser frequency noise and ultimately at a much lower level by shot noise and detector noise. The best performance comes when the system is shot-noise-limited, i.e., the S/N is

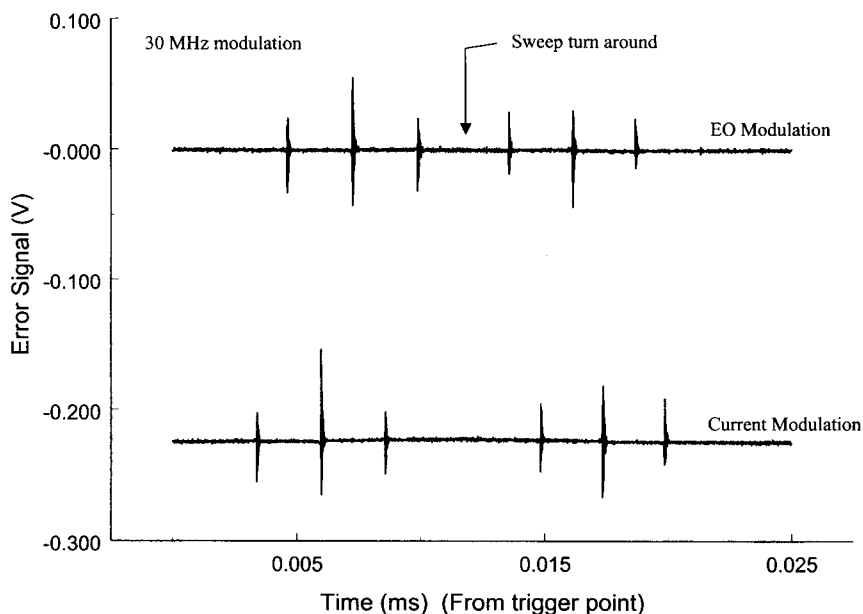


FIG. 1.8. The low-pass filtered output from the mixer is the error signal. The local oscillator phase is set to maximize the amplitude, slope, and symmetry of the central frequency discriminator. Intensity modulation in the lower example results in a large DC offset, which must be compensated by adding a DC current to the integrator input.

limited by the intrinsic noise on the photocurrent. However this depends on whether the laser intensity fluctuations at the modulation frequency are at the shot-noise limit.

If the laser exhibits technical noise in excess of the shot-noise level, noise reduction is possible by using a subtracting detector configuration. In this case, a fraction of the laser beam is sampled by a photodiode prior to the cavity, and the photocurrent is subtracted from the reflection signal. With proper attention to the phase of the two signals, much of the excess common-mode noise can be canceled. We will not dwell further on this approach, because ECDLs are usually within 10 dB of the shot level at practical modulation frequencies. Furthermore, even though a little tighter frequency control is to be gained, there is less light available for the cavity ring-down signal.

A good approach is to ensure that the noise on the signal reaching the mixer is indeed caused by the light. This is most easily checked by sending the RF input to a spectrum analyzer and monitoring the baseline noise levels

near the modulation frequency. With the laser off-resonance (and the light level set to the approximate value that exists when the system is locked), one simply compares the noise levels with the light blocked and unblocked. If the noise does not increase when the light is unblocked, then the system is likely to be limited by the detector–amplifier noise and either the load resistance or the light level should be increased.

1.4 The Loop Filter

The loop filter connects the error signal to the laser, thus completing the feedback loop. The goal of the feedback electronics is to supply enough gain to drive the laser’s frequency fluctuations to the noise floor over as much of the servo bandwidth as possible. With proper design, the electronics should neither limit the residual frequency noise level, when the laser is locked, nor limit the maximum servo bandwidth. The residual frequency noise should be limited by the light’s amplitude fluctuations (technical or shot noise), and the correction speed of the loop should be limited primarily by the laser chip itself. In this section we describe the construction of a loop filter that attempts to achieve these objectives. We start with a very simple introduction to servo system theory and then proceed to a concrete example, namely, a loop-filter design suitable for locking the frequency of a diode laser to a high-finesse cavity.

1.4.1 Introduction to Feedback Systems

A short introduction to the stability of negative feedback systems may be useful at this point. Our objective here is not to treat general control systems [15], but simply to provide some basic understanding and a few rules-of-thumb to aid the construction of a feedback loop to lock a diode laser. With that in mind, we note that a feedback loop supplies corrections to some device (e.g. the laser) within some electronic bandwidth. By “loop” we refer to a signal path that goes from the laser to the cavity, to a detector, to an electronics gain stage, and back to the laser. At low frequencies, the loop will have very high gain, so that small deviations of the error signal are greatly amplified and the device is forced to move in the proper direction, effectively suppressing the deviation. In general, the perturbations will be reduced by the factor of $[1+A(f)]^{-1}$, where $A(f)$ refers to the system gain at the Fourier frequency f . However, consider a small sine-wave perturbation of the laser frequency occurring at a somewhat higher rate, say 100 kHz. As this signal traverses the loop,

various components will shift its phase. For instance, a simple low-pass filter made from a resistor and capacitor will begin to introduce a noticeable phase shift at one-tenth of the filter -3 dB frequency, $(2\pi RC)^{-1}$. As the signal frequency increases, the filter starts attenuating the signal, and the resultant phase shift increases. This phase shift stops accumulating at a frequency about ten times $(2\pi RC)^{-1}$, with the filter output lagging the input by 90° . Primarily as a consequence of the various loop components that act as low-pass filters, the correction signal will be phase shifted from the original perturbation. (Any additional time delay, such as lengths of cable, will further increase the phase lag.) We can safely predict that at a high enough frequency, the correction signal will in fact be 180° away from the negative feedback; in other words, positive feedback! In this situation, if the correction signal causes the laser to move more than the original perturbation, an oscillation will of course quickly develop. To avoid this scenario, one must ensure that the loop gain is less than 0 dB , or unity gain, in the frequency range where the phase shift of the loop is approaching 180° . The *phase margin* of the loop refers to exactly how far (in degrees) from 180° the phase shift is when the loop gain falls below unity. If the phase margin is less than about 45° , the system will tend to oscillate. In the time domain, insufficient phase margin is associated with an under-damped system, and a ringing response to transients will be observed.

The task of the servo design is to ensure that there is sufficient gain at low Fourier frequencies while keeping the unity-gain frequency low enough such that the loop is stable. This naturally leads to a gain versus frequency dependence with negative slope. In fact, a plot of this dependence is a standard part of the servo designer's toolbox. These plots, called Bode plots, are drawn with log-log axes to simplify the interpretation of the various shapes. An example of such a plot is shown in Figure 1.9. A simple picture of the open loop transfer function is the gain as a function of frequency that a signal would experience if it were injected into the diode laser and detected after it travels through the whole system. The bandwidth of the loop refers to the frequency at which the gain falls to 1 (or 0 dB). This is commonly called the unity-gain point, and is often in the range of several megahertz for a servo controlling a diode laser via the injection current.

A rule-of-thumb is that at the unity-gain point, the slope of the loop transfer function may not be decreasing by much more than 20 dB per decade of frequency, which corresponds to a factor of 10 per decade or an f^{-1} slope on the log-log plot. At lower frequencies the slope may be steeper, as shown in Figure 1.9, but the transition from a steeper slope to the f^{-1} slope may not occur near the unity-gain frequency. This model

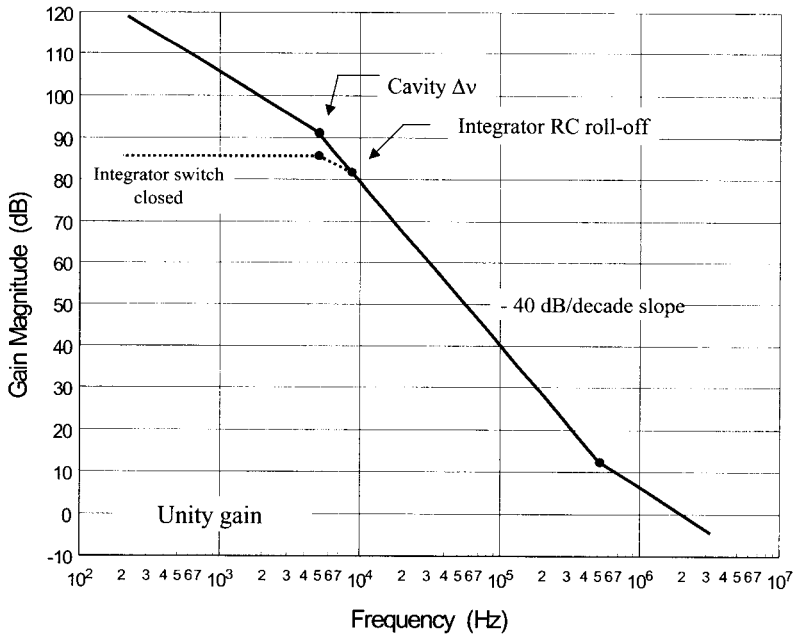


FIG. 1.9. The loop gain transfer function, or Bode Plot, shows loop gain as a function of frequency. Transitions to a steeper slope are caused by RC roll-off elements and the optical cavity. The transition from -40 dB/decade to -20 dB/decade near 500 kHz is caused by a differentiating stage. Note that the loop crosses the 0 dB axis with a slope of 20 dB/decade.

transfer function fulfills the requirements discussed earlier, namely, large gain at low frequencies while maintaining sufficient phase margin at unity gain to protect against loop oscillation.

1.4.2 Characterizing the Loop Transfer Function

We attempt to achieve this stable transfer function by compensating the response inherent in the actual components such as the laser, cavity, and detector. First, the transfer function of the elements in the loop (except for the loop filter itself) needs to be estimated and/or measured. Then a loop filter is designed using op-amps and filters to provide gain and adjust the shape of the actual loop to approach the desired form. In this section we will investigate the loop transfer function *before* the addition of the loop filter, while the next section focuses on the loop filter itself.

The two elements that should have the greatest effect upon the shape of the transfer function are the cavity and the diode laser. Other elements such as the detector have been designed to have minimal effect on the loop transfer function. (We also assume that the cables in the loop path are as short as possible to minimize loop time delay.) As mentioned previously, the reference cavity acts as a frequency discriminator at frequencies below the cavity linewidth and as an integrator (f^{-1} slope) for frequencies above the cavity linewidth. The transfer function of the diode laser itself is considerably more complicated and varies significantly among devices. However, most common low-cost diode lasers are similar in that the phase response rolls off sharply near 1 MHz. As we will see, to compensate this phase roll-off we use phase-lead circuits.

While we often determine the component values for the compensation of the laser empirically, it is certainly possible (and probably preferable when time and equipment are available) to measure the complete phase and amplitude transfer function of the laser's frequency modulation response. This can be accomplished with a frequency discriminator such as that provided by the side of an optical cavity resonance. A plot of the laser's transfer function can be obtained by modulating the current of the laser and recording the amplitude and relative phase of the reflected or transmitted intensity modulation. A number of issues must be addressed in order to make sure the measured phase shift actually represents the laser's phase response. The detector's response must be known, along with that of the frequency discriminator [16]. The laser frequency must loosely be controlled such that it remains on a linear portion of the fringe during the measurement. Also, the laser's amplitude modulation must be removed from the measurement, which can be done with a subtracting-detector arrangement. Careful adjustment and matching of the DC intensity of the beams can be avoided by using negative feedback to match the photocurrents at low frequencies [17].

Although we have mentioned the *shape* of the servo loop's response curve, another quantity that we need to determine is the amount of gain required from the loop filter. Insufficient gain could easily bring the 0 dB point (the unity-gain point) up to the steep portion of the gain curve (i.e., the f^{-2} portion of the curve), rendering the system unstable. Conversely, too much gain could push the unity-gain point up to a frequency with insufficient phase margin, again rendering the system unstable.

The gain of the loop depends on the slope of the discriminator and on the modulation characteristics of the laser itself. For Fourier frequencies less than the cavity linewidth, the discriminator slope is independent of frequency, and can be calculated by slowly sweeping over the cavity line while observing the error signal. If you know the cavity linewidth, simply

divide the peak-to-peak error-signal magnitude by the cavity linewidth to obtain an estimate of the slope in volts per hertz.¹ The modulation function of the laser can be determined at a single frequency (we suggest 10 kHz) by driving the laser current with a sine-wave function generator applied through a coupling resistor. (Our circuit, shown later, is coupled to the laser via a $2\text{ k}\Omega$ resistor.) A measurement of the deviation of cavity resonance will then result in a hertz per volt measurement of the laser modulation characteristic at the 10 kHz rate.²

When multiplied together, these two measurements yield a dimensionless number, call it A' , that indicates the loop gain at 10 kHz without the loop filter. Expressed in decibels, this would be $20\log[A']$. One can then use this to scale the overall gain needed to achieve the desired transfer function shown in Figure 1.9. As we will see in the next section, there are many ways to adjust the overall gain in the circuit.

1.4.3 Loop-Filter Electronics

The loop transfer function described in the previous section is realized by building the proper electronic loop-filter circuit. Fortunately, a fairly simple circuit based on op-amp electronics is sufficient to provide the gain and shaping required for locking a diode laser. The circuit diagram for the loop filter we have implemented for our test laser is shown schematically in Figure 1.10. The basic feedback circuit consists of a first stage op-amp that filters and amplifies the error signal coming from the mixer. The output from the first stage is split into two channels, a fast channel that goes to the laser's injection current, and a slower channel that goes to the PZT controlling the length of the extended cavity. The circuit also contains resistors and capacitors used to shape the transfer function and control the overall gain, as well as six switches (these can be controlled manually or by TTL signals) that are used to enhance the performance of the servo system.

¹If the approximate cavity linewidth is unknown, an oscilloscope with cursors can be used to measure the time between peaks of the central discriminator. Then the measurement can be calibrated using the known frequency offset of the sidebands as a frequency scale.

²Care is necessary to avoid electrically damaging the laser. It is safe to connect the function generator while the power is on, but the output should be adjusted to its lowest setting. With the laser frequency sweeping slowly so that the cavity resonance is visible, one can turn up the drive voltage until a measurement of the deviation can be made with calibrated oscilloscope cursors.

Let us consider this circuit in more detail, with the initial assumption that all the switches are closed. The first stage consists of an op-amp with parallel RC feedback channels. At DC, the effective op-amp gain is $220\text{K}/1\text{K}$ or 46.8 dB , but at the frequency $(2\pi RC)^{-1}$ or 8 kHz , the gain starts decreasing and falls as f^{-1} . This stage is responsible for one half of the -40 dB/decade slope shown in our model transfer function in Figure 1.9. The remaining half is attributable to the cavity. (As mentioned earlier, the cavity roll-off that occurs beyond the $[5\text{ kHz}]$ cavity linewidth contributes a f^{-1} or -20 dB/decade response.) Additional shaping of the transfer function with the first stage is achieved by including a capacitor in parallel with the input resistor. The result is a decreasing input impedance starting at the frequency $(2\pi RC)^{-1}$ or 480 kHz for the values shown. The purpose is to bring the system gain slope from a f^{-2} or -40 dB/decade slope to a f^{-1} or -20 dB/decade slope, prior to the unity-gain frequency. Note this inflection point is shown in Figure 1.9. A much smaller resistor is included in series with the input capacitor in order to limit the very high frequency gain. Depending on the op-amp used, this resistor may not be necessary, and in any event should be as small as possible.

An extra boost in low-frequency gain (below Fourier frequencies of 8 kHz) can be achieved by opening switch $S3$, because then the op-amp functions as an integrator. In fact, the gain at lower frequencies increases as f^{-1} until the amplifier “open-loop” gain is achieved. This level will depend on the particular op-amp and supply voltage, but could easily be 100 dB . Thus, along with speed and noise performance, the open-loop gain should be considered when selecting an op-amp. Because of the increased gain at low frequencies, opening the integrator switch $S3$ forces the error signal ever closer to zero. However, with $S3$ open, if the laser is not locked to a cavity mode then amplifier “a” will integrate to a voltage supply rail and relocking will not necessarily occur. With $S3$ closed, the gain is moderate and the system can respond if the laser frequency is tuned near a cavity mode. Therefore, a good strategy is to acquire the locked condition with $S3$ closed; one can then open it after the lock is established. In Figure 1.9 we show (with a dashed line), the system gain through the injection current channel with $S3$ closed. We emphasize however, that this transfer function does *not* include the low-frequency gain supplied by the PZT channel. As we will describe, for Fourier frequencies under 1 kHz the PZT channel gain becomes significant and is in fact dominant near DC.

Another feature of the first stage is the external offset current that is fed in via switch $S2$. The op-amp itself should be balanced using the standard offset compensation circuit that can be found in the amplifier datasheet. We have chosen to include an additional offset current channel because the circuit may be used with injection current modulation instead of an

electro-optic crystal. In this case, substantial current offsets are necessary to compensate for the laser's amplitude modulation. We stress once again that current modulation is not recommended if drift of the locked laser relative to the exact center of the cavity is critical, because the amount of AM from the laser typically changes with time. The voltages that the offset current is derived from (shown in Fig. 1.10 as $\pm V_o$) should be well filtered and no larger than necessary to compensate for the mixer offset.

The output from the first-stage amplifier is then sent through separate paths to the two correction elements. First let us consider the fast path, which passes through a variable attenuator and goes to op-amp "b" that feeds current to the laser diode. If gain beyond that provided by the first stage is needed, one can appropriately set the feedback resistance on op-amp "b." Otherwise, one can simply omit op-amp "b," which will enable a slight increase in the attainable servo bandwidth. Regardless, the resulting "fast" correction signal is sent to the diode laser through switch S6 and a 2 k Ω resistor. This resistance together with any capacitance inherent in the laser will form a low-pass filter, possibly contributing phase lag that would affect the system bandwidth. Another possible coupling configuration would be to use a smaller coupling resistor and limit the voltage of the amplifier output, but this is not recommended as it could expose the laser to dangerously large current spikes. We emphasize here that caution is required to avoid damaging the laser, because the feedback is coupled directly through the resistor to the laser anode or cathode (depending on the diode's polarity). Note that, although Figure 1.10 is drawn with a coupling to the anode of a positively biased laser, no changes are necessary to couple to the cathode of a negatively biased laser. At the end of this section, we describe how to use S6 to protect the laser against potentially fatal current spikes from the loop-filter electronics.

In order to provide phase-lead to compensate the diode laser, we have added parallel capacitance on the second stage input resistance and on the coupling resistor to the laser (see Fig. 1.10). While the optimum values for these components will depend on the particular laser, the values shown are probably good starting points. In addition, in both of these places and on the first-stage input there is also a second resistor in series with the capacitor. These resistors will limit the circuit's high-frequency gain in the region above 30 MHz. The purpose is to prevent parasitic loop oscillations that can develop at very high frequencies. Again the values shown are probably good starting points, although decreasing the resistances may allow a wider bandwidth as long as the loop remains stable.

Also included on the input to op-amp "b" is an offset control. Its purpose is to bias the second amplifier's output such that under steady-state

conditions (or for the unlocked state) there is very little current through the coupling resistor. In other words, we adjust the pot so that the amplifier's output is close to the laser's junction voltage. This gives the circuit the advantage that it can be connected to the laser by closing switch S6 without appreciably perturbing the laser frequency.

The second path, or "slow" feedback, goes to op-amp "c" that controls the voltage sent to the PZT driver for the laser cavity end mirror (see Figure 1.1). The advantage of this configuration is that the low-frequency corrections are dominated by the large gain in the PZT channel so that the laser current DC level remains constant (large changes in this current can lead to laser mode hops). In fact, under the locked condition the DC level after the first stage should be equal to zero when this stage is configured as an integrator. This highlights one of the aspects of using integrators in these feedback loops—a nonzero correction can be applied (after amplifier "c" in our circuit) while a zero-error signal can be maintained. We also benefit from the increased low-frequency gain in the overall loop transfer function. For added flexibility in this stage we have included a switch, S5, across the feedback capacitor, so we can switch between flat gain and a full integrator (as we did with S3 on the first stage). We have assumed that the PZT driver has a polarity switch, so one can find the right sign for the feedback. If the PZT driver does not have a polarity switch, a simple inverting amplifier following the integrator may be required. We note also that for more accurate locking (a smaller offset from the center of the cavity resonance), amplifier "c" should have its inherent offset adjusted prior to any input. This can be done with a trim pot wired as shown in the amplifier datasheet. This adjustment is not important for the second-stage amplifier "b," because the gain of this stage is small relative to that of the first stage.

We control the overall loop gain with a potentiometer that varies the attenuation between the first and second amplifier stages. A one-turn pot is used to minimize stray capacitance that might limit the feedback bandwidth. Care should be taken to minimize the length of the signal path by positioning the pot close to the amplifiers, preferably by mounting it on the board.

The choice of amplifiers for the various stages is based on bandwidth, noise density, and offset performance. For the two amplifiers in the current correction path ("a" and "b"), we want to have amplifiers with little time delay (recall that time delay leads to unwanted phase shifts) and good noise performance. Thus we usually work with op-amps with gain-bandwidth products > 30 MHz and a noise level substantially less than that of the input signal. Of course, the level of the input noise most likely depends upon how much light is incident on the detector (recall the

discussion in Section 1.3.3). For light power above 100 μW , the noise requirement on the op-amp is not severe and should not present a problem. The op-amp in the PZT channel ("c") does not need to be fast, so a standard field-effect transistor op-amp with good DC characteristics should suffice.

As we have already seen, there are several analog switches in the circuit, which we use to increase the low-frequency gain, and to allow quick, remote unlocking and relocking via TTL signals. Most of these switches are optional, with the exceptions being the PZT integrator switch S5 and the output protection switch S6. However, these may be simple manual toggle switches. The TTL-controlled switches are available from many manufacturers with many different features. Because speed is relatively important, we implement switches that can respond in about 100 ns. Low-leakage current is also important in the integrator configurations; the switches we chose are specified at 100 pA. Minimizing switching transients is another consideration, with the pertinent specification showing up on many datasheets as "charge injection."

In order to maximize performance and flexibility, it is important to consider the layout of the electronics. The electronics board should be located close to the laser to avoid unnecessary time delays and cable capacitance. To achieve higher loop bandwidths we typically place the loop filter in a small box adjacent to the laser on the optical table. We have also chosen to place the coupling resistor and capacitor to the laser on the electronics board instead of physically placing them next to the laser. This allows an additional TTL-controlled switch to be placed on the laser-side of the coupling resistor, for the purpose of rapidly shunting a few milliamperes of current away from the laser to quickly shift the laser frequency. Such switching may be useful in cavity ring-down spectroscopy, and is discussed in the following sections.

We conclude this section with some words of caution. The output switch S6 is necessary in the present implementation in order to protect the laser against current spikes from the loop-filter electronics during cable connect/disconnect episodes or power up/down. When connecting or disconnecting the cable between the circuit and the laser, one first opens S6. Similarly, before turning the circuit power off, it is wise first to open this switch and disconnect the laser. While there are certainly more elegant solutions that might be included to protect the laser without contributing unwanted phase shift in the signal channel, we have found that this approach works fine in the laboratory setting. A more robust solution may be necessary, however, if the circuit is to be handled by less-experienced users.

1.5 Locking the Laser and Loop Optimization

1.5.1 Initial Locking

With the loop filter designed, all the pieces are in place to lock the laser to a cavity fringe. It is easiest to start by sending a triangle wave to the laser's or cavity's PZT driver to sweep slowly over a cavity free-spectral range and view the transmitted signal on an oscilloscope. A sweep rate of ≈ 30 Hz is a good compromise between sweeping too quickly which reduces the power buildup in each mode, and too slowly which makes alignment difficult. First we need to compensate any DC offsets in the signal coming from the mixer. With the laser nominally aligned as described in Section 1.3 and the output switch S6 open, the signal at the output of op-amp "a" should yield a low-pass filtered version of the error signal, centered at zero volts DC. Any offset from zero can be adjusted with the offset current through switch S2.

We usually check for locking action by closing switch S6 and turning up the gain knob (with the PZT feedback gain turned off for the moment). An increase in transmission, widening of the resonance, or even a flat section on top of the fringe are all indications that the polarity of the system is correct and the feedback loop is attempting to keep the laser locked to the fringe (see Fig. 1.11). If the overall sign of the servo is wrong, the fringe will be suppressed and the system will try to lock to the sidebands instead. This situation can be corrected by choosing the opposite slope for the discriminator (e.g. by changing the demodulation phase by 180°). Adjusting the overall gain by turning the potentiometer will give some clues as to whether more gain is required or not. If the servo seems to be trying to lock, one can then reduce the sweep to zero and the laser should stay in lock for increasing fractions of the sweep. If this is indeed the case, now is a good time to optimize the alignment to the optical cavity, because the sweep-to-sweep amplitude variations will be greatly reduced.

If the servo clearly perturbs the laser but no locking action is observed, even on the sidebands, it is possible that the gain is too low and the servo is unstable. Remember that a portion of the servo loop's gain curve is decreasing as f^{-2} , or -40 dB per decade. If the overall loop gain is such that the servo's unity-gain point occurs on this steep slope then no locking will occur. Determining the system gain as detailed in Section 1.4.2 will offer a clue as to whether this is the problem.

When the laser is locked to the cavity, increasing the gain past the optimum point will cause a servo oscillation that will be evident on the error signal. Another indication that this is occurring is a reduction of

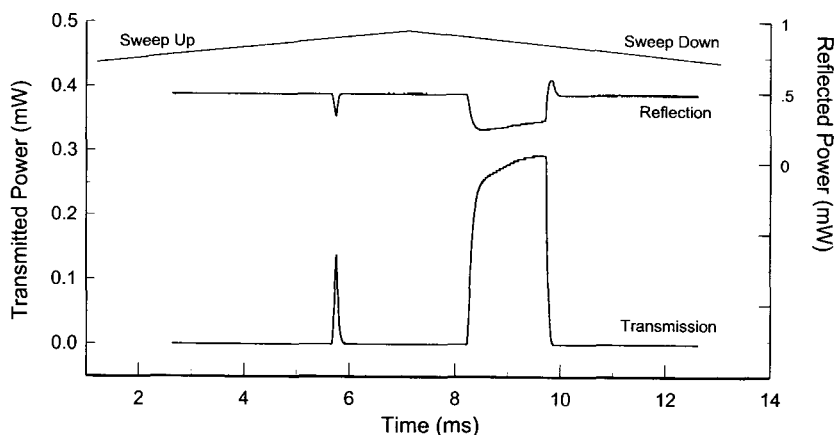


FIG. 1.11. Traces of the laser power transmitted through, and reflected from a high-finesse cavity as the laser frequency is swept up and down. The injection-current portion of the electronic servo is on, and briefly locks the laser to the TEM_{00} mode as the laser passes the resonance from each direction. Here the asymmetry is attributable to the offset adjustment of the high-gain stage in the feedback loop, but similar behavior is possible from other causes (see text).

the power transmitted through the cavity as power is taken out of the carrier and put into oscillation sidebands. Reducing the gain slightly (just below oscillation) provides a good initial setting. In a similar way, the PZT gain can be increased until instability ensues and then reduced slightly. If turning on the PZT gain drives the laser out of lock, the sign of the PZT feedback may need to be reversed. With the PZT gain turned up, the signal observed after the first amplifier stage ("a" in Fig. 1.10) should be driven to zero.

As an aside, we note that with high-finesse cavities there is often an asymmetry in the locking performance in regard to the direction of the frequency sweep (see Fig. 1.11). This usually is the result of a nonzero offset in the locking electronics. Such asymmetries can also be caused by localized heating of the high-finesse cavity mirrors, caused by absorbed light as the mode power builds up. In this case, the thin films will expand and the cavity resonance will shift slightly toward higher frequencies. This effect is similar to the nonlinearity observed with gas-filled cavities [18]. Regardless of the cause, the net result is that the effective sweep rate of the laser relative to the cavity is effectively higher in one direction than the other and/or the locking range is reduced. Consequently, the laser spends less time locked to the cavity, possibly to the extent that maximum transmission is never achieved.

1.5.2 Adjusting the Loop Parameters

In the previous section we identified the lock condition simply by observing the light level transmitted through the cavity. In order to optimize the feedback system, it is necessary to have a better diagnostic. Fortunately, much of this information can be obtained from the same error signal one uses for the locking. A directional coupler placed just before the mixer can be used to extract a small fraction of the error signal without significantly changing the overall loop gain. One can then use an RF spectrum analyzer to view the noise sidebands, which are centered at the modulation frequency. The spectrum analyzer display should resemble the trace shown in Figure 1.12, which consists of a center peak at the modulation frequency with noise sidebands on either side. If the gain is increased, the sidebands will increase in amplitude as the servo starts to oscillate; if the gain is then reduced slightly so that the system is well behaved, the width of the noise spectrum is an indication of the servo bandwidth.

One can now attempt to increase the bandwidth and reduce the level of the low-frequency noise by changing component values in the loop-filter circuit. We recall that there are three resistors in the circuit of Figure 1.10 that serve as convenient points to place a parallel capacitor, thus creating a high-pass element with its corresponding phase advance. These resistors include the first- and second-stage amplifiers' input resistors, and the coupling resistor to the laser diode. One can use the spectrum analyzer display to monitor the error signal's noise spectrum and adjust these component values in an iterative fashion to obtain the best performance.

It is worth spending some time on the process by changing components and values to arrive empirically at a wider system bandwidth. A simple technique that often helps is to place a finger on one of the phase-lead capacitors, thereby increasing the actual capacitance slightly. This may offer a clue as to whether more or less capacitance is desired. Not only can the capacitor values be adjusted, but different op-amps can also be inserted to determine whether amplifier speed is playing a role in the system bandwidth. Similarly, the length of the cable between the detector and mixer can be adjusted to determine whether time delay is an issue.

A spectrum analyzer monitoring the error signal can be useful in improving the performance of the servo loop as illustrated using our demonstration system. Figures 1.12a and 1.12b show a noise spectrum for the locked ECDL using EOM modulation. These data are referenced to the mixer input, and are measured by using a -10 dB coupler prior to the mixer and subsequently correcting the data plots by 10 dB. The lowest level depicted in Figure 1.12a is the laser's amplitude noise level. This represents the lowest level to which we can reduce the frequency noise—further gain

only writes amplitude noise onto the frequency of the laser. The AM noise was measured by unlocking the system and attenuating the beam reflected from the cavity such that the power incident on the detector was the same as when the system was locked. The signal when no light was incident on the detector is not shown, but is approximately 10 dB below the light level. A number of technical noise spikes are evident in the laser's AM spectrum, along with a peak at the modulation frequency, 29 MHz.

The top trace shows the error signal, which contains several important features. One can see at ± 4 MHz from the center, large "servo bumps," which are characteristic of the servo system. Here the phase for the correction signal is approaching 180° (positive feedback), so the noise is increased rather than suppressed. If one turns up the loop gain, the system will oscillate (displaying tall, sharp peaks) in this region. The oscillation frequency indicates roughly the bandwidth of the servo—if the bandwidth is much less than anticipated, then further evaluation of the loop filter and laser/detection system is needed. Not surprisingly, we have observed that the bandwidth can change with the laser's operating parameters. This is to be expected, because we know that the frequency noise spectrum of the laser

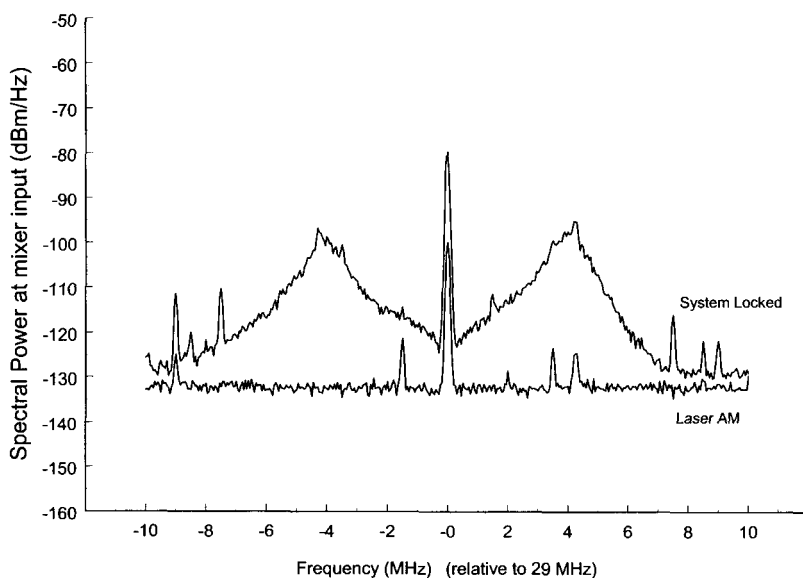


FIG. 1.12a. The noise spectrum at the mixer input, measured with a RF spectrum analyzer using a 100 kHz resolution bandwidth. The noise sidebands are peaking at 4 MHz, indicating the approximate servo bandwidth nearly this wide. The noise with the laser unlocked and shifted away from the resonance is shown in the lower trace. An expanded scale is shown in Figure 12b.

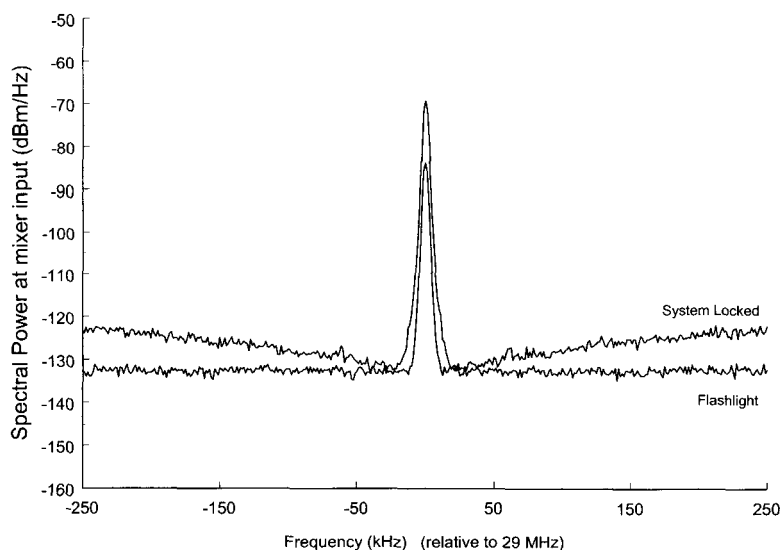


FIG. 1.12b. The noise spectrum at the mixer input, under the same conditions as Figure 1.12a, but measured with a 3 kHz resolution bandwidth. The lower trace, generated by a flashlight, shows the noise fluctuations of a shot-noise-limited photocurrent of the same DC level as when the system is locked. The laser's amplitude noise shown in the preceding figure is essentially at the same level. Note that the error signal approaches the intensity noise at low frequencies, indicating that the system is suppressing frequency fluctuations to the shot-noise limit in the kilohertz range.

changes as the laser sees optical feedback, or approaches a mode hop, for instance. An increase in the magnitude of the FM noise spectrum would be accompanied by an increase of the laser's FM modulation coefficient, which in turn will affect the servo loop's gain. Similarly, optical feedback characteristics (e.g. alignment of the feedback mirror) of the laser cavity will change the operating parameters, so tweaking up the alignment can sometimes restore the locking performance.

The error-signal noise at low frequencies appears to be well above the amplitude noise in Figure 1.12a, however this is a measurement artifact caused by the wide spectrum analyzer bandwidth. Taking a closer look, as in Figure 1.12b, the error-signal noise when locked is approaching the amplitude noise at approximately 20 kHz from the carrier. At even lower frequencies where the servo has higher gain, the error-signal noise will appear to be less than the amplitude noise level, however this is simply an indication that the laser is being frequency modulated by the amplitude noise. Note that we have calibrated our noise measurements to the shot-noise-limit with the aid of a flashlight to supply a shot-noise-limited photocurrent. As in the case

of the laser amplitude noise measurement in Figure 1.12a, the shot-noise measurement was accomplished by adjusting the DC photocurrent to the same level as exists when the system is locked. The data traces show that the laser's AM noise is essentially at the shot-noise level in the region near 30 MHz.

We noted the power peak at the modulation frequency in Figure 1.12a. There is almost always some residual power detected at the modulation frequency because of RF pickup by the laser cables or housing. Another possible source of AM power at the modulation frequency that often occurs is attributable to improper alignment of the modulator crystal with respect to the laser's polarization axis. However, the residual level attributable to these sources is indicated by the lower trace. The coherent spike increases by 20 dB when locked, and is most likely caused by one or more transverse cavity modes near the modulation frequency. The affected sideband would be slightly phase shifted if any power was coupled into a transverse mode. Consequently, the RF modulation detected on the cavity reflection would not cancel at the exact center of the cavity. For applications in which the small offset from the center of the cavity resonance is important, attention to the proper choice of modulation frequency, cavity geometry, and spatial coupling is required. Note that the coherent peak is still a very small fraction of the power at the mixer input when the laser is unlocked but near the cavity mode. The power when the system is unlocked may be measured with an oscilloscope as the laser is swept over a cavity fringe, and for this system it is on the order of 0 dBm, or -50 dBm/Hz when referenced to the 100 kHz spectrum analyzer bandwidth of Figure 1.12a.

At this point the performance of the laser lock should be adequate for most applications that require narrow linewidths and high spectral resolution. It is now appropriate to expand the capabilities of the lock and start to take advantage of the switches to make the lock more robust and more flexible.

1.5.3 Evaluating the Locked System

With the servo working reasonably well, one might be curious as to the linewidth of the locked laser and how to measure it. This problem is actually trickier than one might guess. The best way is to beat it against a more stable laser on a photodiode and evaluate the resulting fluctuations. The strength of this technique is that it reveals both of the major contributors to the laser's frequency-noise spectrum: the residual noise of the laser lock, and the fluctuations of the cavity length. However, this requires access to a more stable laser of the same color. One can of course build a replica of the first locked laser system and make a comparison. An independent Fabry-Pérot

cavity can also be used as a frequency discriminator. However, both of these approaches will suppress any noise that is common-mode to the two systems under comparison, so the resulting beat-note noise spectrum may be artificially clean. This is a serious problem because a frequency fluctuation of kilohertz magnitude corresponds to a vibration of the cavity length below the picometer level (25 cm cavity length, 800 nm wavelength). The self-heterodyne techniques involving long optical fibers are impractical, given the long coherence length of a laser locked to a high-finesse cavity, and the additional phase noise caused by the fiber.

For many applications including evaluation of the lock itself, it is useful to measure the noise within the servo loop relative to the cavity. It is very important to emphasize that measurements of the noise within servo loops are not good indications of the actual laser linewidth or lineshape. One reason is that the servo gain acts to minimize the noise at the integrator input, *even if the noise is not caused by frequency fluctuations*. For instance if some of the signal fluctuations at the integrator input are actually amplitude fluctuations, the servo will suppress this signal by increasing the frequency fluctuations. Nonetheless, the noise measurements are useful in determining the performance of the electronics, transducers, and servo components.

In the absence of excess noise of the electronics or cavity (usually *not* the case), one can make an estimate of the performance of the servo system. In particular, consider the case in which the noise spectrum is flat out to some bandwidth, and whose noise level is small compared to that bandwidth (typically the case for a well-stabilized laser). Then the contribution of this noise to the laser line shape is a Lorentzian with a linewidth equal to $\Delta\nu = \pi(\Delta\nu_{\text{rms}})^2/B$, where $\Delta\nu_{\text{rms}}$ is the root-mean-squared fluctuation derived from the spectrum analyzer noise level divided by the discriminator slope, and B is the bandwidth for the spectrum analyzer measurement. Many spectrum analyzers give the noise level as spectral power, in dBm/Hz; one needs only to convert dBm to volts² and then divide by the square of the discriminator slope to convert the measurement into hertz. While our noise spectrum is not usually completely flat, one can nonetheless use this formula to make estimates that are reasonable enough, especially in light of the fact that this approach neglects fluctuations of the cavity itself, which for most good locks, probably are the dominant noise source. There are also mathematical techniques for mapping spectra that are more complicated into laser line shapes [19, 20].

1.5.4. Making the Lock More Robust

While the method described thus far should produce a good laser lock, there are several minor modifications we can add to make the locking

system more robust and reliable. In particular, the system will be easier to work with if it can keep the laser locked in the face of moderate disturbances and quickly recover from more severe events. Increased servo gain is the answer to acoustic perturbations, as the loop will reduce the error-signal fluctuations by nearly the inverse of the gain. For example, at acoustic frequencies the system gain may be 100 dB, or 10^5 . The frequency noise of the laser caused by voices will be reduced to 10^{-5} of what it would have been in the unlocked case. More challenging problems come with sharp raps on the mechanical system caused, for instance by dropping a tool on the optical table, as the system has much less gain at a few hundred kilohertz. As a first step toward reducing the sensitivity to such effects, it is advisable to use solid mechanical structures for the optical components in the system. Good vibration isolation of the reference cavity and the laser box will reduce the instability of the system and hence the servo requirements. Other effective passive approaches include the use of robust mounts (i.e. no laboratory posts) for the mirrors and beam splitter.

Even with good mechanical mounting, the system will still be susceptible to out of lock events caused by sharp mechanical perturbations. Here the inclusion of two electronic switches in the control circuit, one manual, and one electronic, serves to improve the lock and make it more robust. The first of these, S5 in the circuit diagram, simply turns the op-amp "c" from a proportional gain stage to a pure integrator. This increases the low-frequency gain and drives the signal after the first stage to zero as desired. During acquisition it is preferable to have this switch closed (i.e. integrator off); otherwise, this op-amp will sit on the rail (because of even tiny offsets, an integrator will drift to rail quickly when the laser is out of lock) and make acquisition unidirectional at best. When lock is achieved, one can simply flip the switch to increase the low-frequency gain.

A second switch, S4 in the circuit diagram, can be inserted to make the lock considerably more robust. This switch may be configured so that it simply turns off the input to the PZT integrator when the laser momentarily goes out of lock for any reason. The TTL signal that controls the switch is derived from a comparator that monitors whether the transmission through the cavity (i.e. the voltage from a detector placed after the cavity) is greater than a reference voltage (set to be roughly half of the detector voltage when the laser is locked to the cavity). When the laser is pulled into resonance by the current, the switch is closed and the PZT channel is turned on. If a perturbation momentarily knocks the laser out of lock, the switch automatically opens and the value on the capacitor is held, so that when the perturbation is over the laser frequency should be close to the previous value and relocking to the cavity can occur easily. We have found that the addition of this switch will keep the laser locked in the presence of fairly

significant noise. This makes it possible to work on the table without knocking the laser out of lock.

Note that if the laser stays out of lock for more than a short time, the integrator will start to drift away from the “lock” value toward a rail and relocking is less likely. To reacquire the lock, usually one needs to close S5 to turn off the integrator. In principle, this can pull the laser frequency quite far from that of the desired cavity mode, especially if the previous “lock” value was large. We have found that with occasional monitoring, this correction voltage from the amplifier can be kept near zero simply by tuning the DC knob on the PZT driver. Then relocking to the same cavity mode is usually straightforward.

During momentary unlock events the input switches (S1 and S2) to op-amp “a” could also be opened in a manner analogous to using S4 as discussed above. This would allow the amplifier feedback switch S3 to remain open, in integrator mode. However, the RC time constant of this amplifier is much shorter than amplifier “c,” and our experience has been that relocking is much less likely if switch S3 remains open.

There are other approaches to building “smart” servos that automatically relock the laser. For a processor-controlled system, it is possible to control the locking electronics and the laser in such a manner as to mimic manually relocking of the laser. Alternatively, this can be achieved using analog electronics. For instance, one approach upon an accidental unlock is to sweep the laser back and forth to once again “find” the cavity mode. Our implementation uses a single-pole, double-throw switch on the input to the first op-amp, in place of S1. The normally closed switch setting would connect the mixer output to the op-amp input. The other switch position connects the op-amp to a small oscillating current, generated by a low-pass filtered 555-timer circuit. When the cavity transmission falls below threshold, S1 is set to sweep the laser. Once the laser relocks, S1 is switched and the sweep is no longer sent to the circuit. This works quite well, as pounding on the optical table will cause the laser to unlock, reacquire, and lock within a fraction of a second. The circuit is shown in Figure 1.13.

1.6 Repetitive Locking for Cavity Ring-Down Spectroscopy

We conclude this chapter with some modifications to the locking circuits that make it well suited for cavity ring-down spectroscopy (CRDS) experiments. Many CRDS measurements are realized with pulsed lasers

or continuous-wave (CW) lasers that are swept through the cavity fringe. CW locking of the laser to the cavity has advantages that have been previously recognized and discussed in the literature [21–24]. These advantages come from the ability to excite a single-cavity mode, and with more power than is possible by simply sweeping past the mode. Such experiments use repeated locking and intentional unlocking to generate ring-down decays from the cavity, improving the signal-to-noise through averaging. While we briefly discussed relocking issues in the previous section, the process of rapid unlocking and relocking of a laser to a cavity requires a more detailed understanding of the reacquisition process. In this final section, we will discuss repetitive locking and show how the addition of a few more switches to the loop filter can accommodate its requirements.

A limitation during reacquisition of the lock is the cavity ring-up time, which of course is the same as the ring-down time. In other words, one expects that the system cannot lock instantly, and that the time period required for the cavity transmission to reach a stable intensity will be no shorter than the cavity-decay time. For a cavity with a high Q , we expect that the contribution of the servo to the time required to reach an equilibrium state is relatively minor, because the servo bandwidth is approximately 1 MHz, while the timescale to equilibrium for a typical high- Q cavity for ring-down spectroscopy is on the order of hundreds of microseconds. As expected, we find that during the relocking process the error signal returns to zero well before the cavity transmission stabilizes. There are two caveats to this; the first being that switching the servo “on” does not mean that the laser will immediately be driven with high gain to the resonance center. In fact if the laser frequency is off resonance by more than a few cavity linewidths (but less than the modulation frequency), the laser could still be driven in the proper direction, but relatively slowly, as the error signal is near zero. Thus, one expects that acquisition will take longer than the microsecond timescale that corresponds to the servo bandwidth.

The second caveat is that for a system with a frequency control “knob” other than the laser injection current (such as the piezoelectric length of an extended cavity), the laser frequency and the transmitted power through the cavity may not be perfectly correlated. For instance, in response to a transient that shifts the cavity resonance towards the blue, the servo may quickly return the laser frequency to the center of the cavity resonance by decreasing the injection current and hence also the laser power. With the laser frequency fixed once again on the cavity resonance, the circuit in Figure 1.10 will slowly bring the injection current back up to the original value by acting on the PZT, decreasing the laser cavity’s mechanical length. Consequently, we expect the “ring-up” of power in the cavity to the previous stable value may take slightly longer than the ring-down, because

the laser power was lowered to relock the system. This process occurs in the reverse sense also, causing a slightly faster return to equilibrium if the laser power is increased to relock.

In order to switch the laser in and out of lock repetitively, we need some fast way to turn the servo off and back on. Simply switching the error signal off (with S1) at the first amplifier's input serves to unlock the laser, because the servo loses access to the error (i.e. correction) signal. Switching the output switch S6 has the same effect. We keep the PZT integrator switch (S5 in Fig. 1.10) open during the unlocking and locking cycle so that the laser's PZT drive voltage remains appreciably unchanged. Furthermore, for the same reason the integrator input switch S4 remains open during "unlock" and closed during "lock." If the input switch S4 is not utilized, the output of the integrator may be dependent on the duty-cycle of the unlocking and locking. This is a function of the current through the integrator's input resistor during the "unlock" portion of the cycle. If it is nonzero, the integrator output will attempt to steer the laser away from the resonance, which will result in the system breaking lock after only a few cycles of the relocking sequence.

To allow the cavity ring-down time constant to be measured, CRDS experiments have most often used acousto-optic modulators to switch the input beam away from the cavity mirror. Switching the laser to a sub-threshold value has also been employed [25]. Rapid frequency shifting of the laser to a stable off-resonance frequency has been used to allow heterodyne ring-down measurements in reflection [26]. Observing the ring-down decay in transmission by frequency shifting is also possible for some cavity configurations. The shift must be accomplished in a manner such that the cavity is a high-attenuation filter. This requires attention to the cavity's transverse-mode structure in the vicinity of the fundamental mode, and the cavity's vibrational stability. Here we have used a half-symmetric resonator constructed from a plane mirror and a 30 cm radius mirror, separated by about 25 cm. There is a 15 MHz wide-frequency region to the blue side of each fundamental mode of this cavity that appears free from transverse modes. A modal analysis indicates that in this region there are no modes with combined transverse-mode orders ($n+m$) of less than 100. Thus by shifting the laser frequency so that the laser sits in this "dark" frequency region during cavity ring-down, we should have minimal contamination of the ring-down signal attributable to leakage of the carrier into the cavity. This approach may require neglecting some data points at the beginning of the ring-down decay if the laser does not shift frequency quickly enough. To ensure the locking sidebands at ± 29 MHz do not leak into the cavity, they are turned off with another CMOS analog switch at the start of each ring-down event. The switch was arranged in the center of a "pi"-shaped 50 Ω

configuration, which provided 30 dB of “off” attenuation and about 1 dB of “on” attenuation [14]. We note that more attenuation can be accomplished by using three switches, one in series with each resistor of the “pi” attenuator.

The laser-frequency shift is accomplished with another analog switch, which is not shown in Figure 1.10. The switch is positioned on the diode laser side of the 2 k Ω coupling resistor, and, when closed, connects the diode anode to ground through a variable resistor. The switch is closed by the same TTL signal that unlocks the laser. Several hundred microseconds later, after the ring-down measurement is completed, the switch is opened and the laser frequency returns to the vicinity of the cavity fringe (i.e. within the PDH capture range). A single cycle of this unlock, relock sequence is illustrated by Figure 1.14. During this cycle the laser frequency is monitored by using an auxiliary reference cavity, positioning the resonance such that

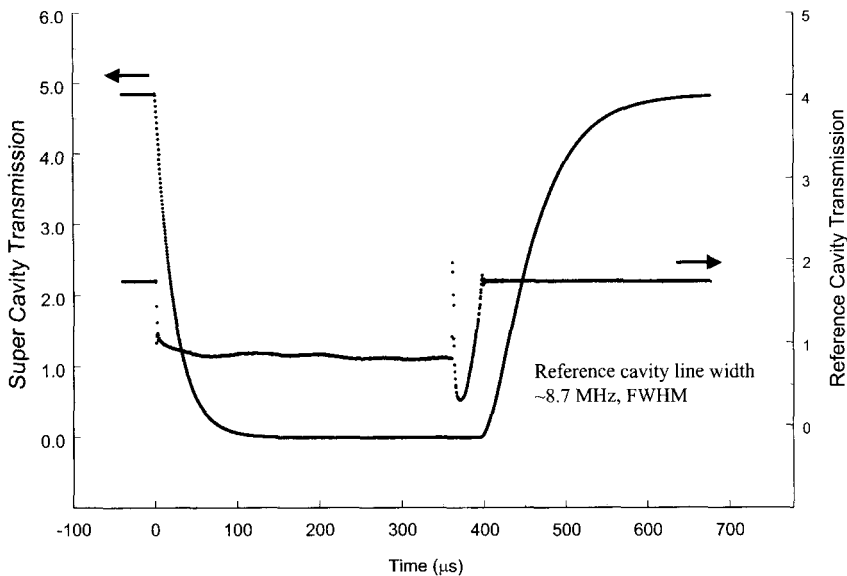


FIG. 1.14. Data from a single ring-down and relocking cycle. The exponential decay from the high-finesse measurement cavity corresponds to the left scale, while the laser’s optical frequency is indicated by using the transmission from an auxiliary reference cavity. Prior to the ring-down event (triggered at $t = 0$), the laser is locked to the high-finesse cavity and the reference cavity transmission is approximately 75%. The laser is shifted to the blue approximately 3–4 MHz by shunting 0.3 mA of current away from the diode when it is unlocked (see text). The relock signal happens at 355 μ s, and reacquisition and locking occurs in the next 100 μ s.

the laser frequency is on the side of a fringe. The sample rates of both the ring-down data and the reference cavity data are $0.4\text{ }\mu\text{s}$ per point. Within three data points ($\approx 1.2\text{ }\mu\text{s}$), the laser frequency has shifted about 3–4 MHz away from the high-finesse resonance.

We find that the repetitive locking works better with only one full integrator in the system, so we keep switch S3 closed for both locking and unlocking. A system undergoing this repetitive locking is shown in Figure 1.15. Although only six cycles are shown, the system is capable of indefinitely unlocking and relocking. Using software to set the TTL “unlock” and “lock” delays, we have cycled through 10^6 ring-down sequences. However, some caution must be exercised before adding a large number of decays together to measure the exponential decay time. The limitation on the measurement time and subsequent number of averages will be given by any nonrandom process that may change the cavity loss, for instance a subtle change in gas pressure or temperature. In such a situation, although the random noise will appear to decrease with averaging, the

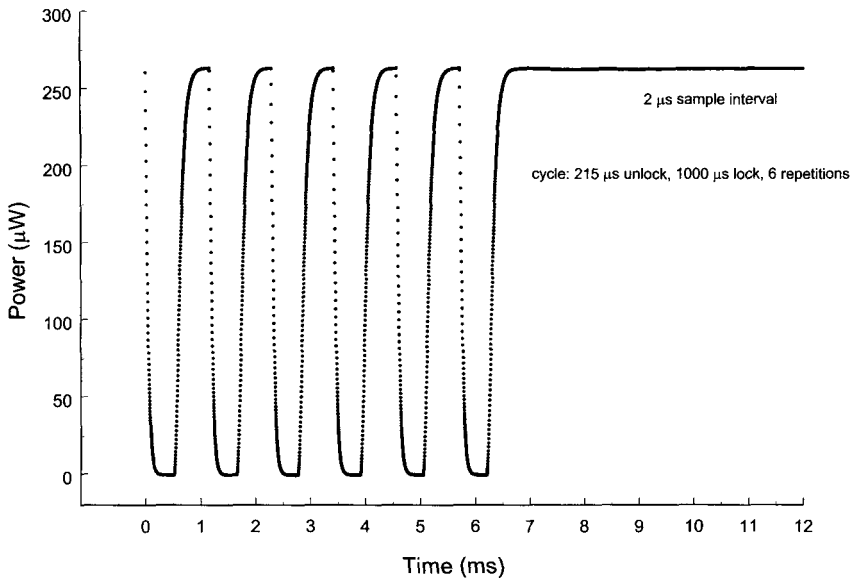


FIG. 1.15. Repetitive locking of the laser is possible by opening and closing the input switches or the output switch. At $t = 0$, the first ring-down cycle is triggered, and the laser is simultaneously unlocked and frequency shifted to the blue by decreasing the current to the diode. The laser is shifted back near the cavity resonance and the input switch closed $215\text{ }\mu\text{s}$ later. The transmitted power through the cavity returns to an equilibrium value, and the process is triggered again at a 820 Hz repetition rate.

measurement repeatability will not improve. Shown in Figure 1.15 is an arbitrary $1000\mu\text{s}$ delay between the end of each “unlock” period and the start of the next “unlock” period. Higher throughputs (faster averaging times) are possible by reducing this dead time. As may be observed in Figure 1.14, this system returns to the steady-state transmission in about $350\mu\text{s}$ from the TTL “lock” signal. Even though the system will relock to the cavity center in a shorter amount of time, this does not necessarily mean that the cavity transmittance has stabilized at its previous value.

A set of twenty-five sequential decays is averaged together and is shown in Figure 1.16. The exponential decay appears linear on a log scale, until the trace reaches the noise floor set by the detector noise level. Fitting an exponential curve to this data results in a $1/e$ time constant of $24.2\mu\text{s}$. The exponential fit was weighted for shot noise (square root of the signal level), and includes a constant term for the offsets that account for the analog-to-digital converter and detector noise.

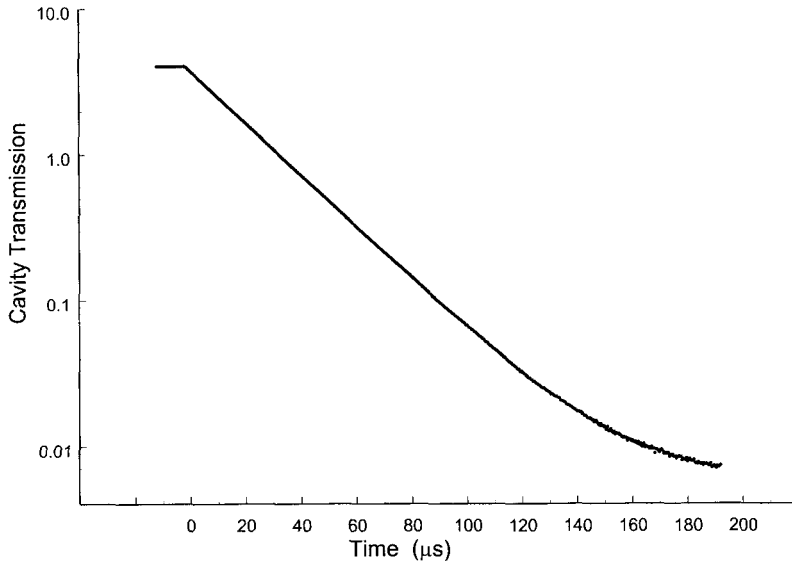


FIG. 1.16. The average of twenty-five sequential cavity ring-downs, shown on a log scale. The cavity was in a closed chamber, with the loss caused by a coincidence between the cavity mode and an absorption line in the water band near 830 nm . The data was taken with a 16 bit analog-to-digital converter with a sample period of $0.4\mu\text{s}$. The trigger point is at $t=0$, and the subsequent exponential decay time constant is $24.2\mu\text{s}$. When an exponential fit is performed on each of the ring-down decays individually, resulting in twenty-five time constants, the standard deviation is $0.01\mu\text{s}$.

1.7 Conclusions

In this chapter we have described in some detail a fairly simple approach for locking an ECDL to a high-finesse Fabry-Pérot cavity. In our lab, we have used this approach to lock diode lasers for many different applications including precision metrology of atom-trapping experiments, high-resolution atomic spectroscopy, cavity ring-down spectroscopy, and length metrology. While the different lasers required different values for the electronic components, we have found the overall approach to be quite flexible and yield good results on a fairly short time scale. Our experience has shown that laser linewidths of well under 1 kHz can be readily achieved with this technique. We note that many other laboratories have used a similar approach to achieve excellent results as well. We have attempted to make this chapter “user-friendly” in the sense that it includes details not normally included in the standard research papers. We even added descriptions of a few “bells and whistles” in the form of switches that might make the locking system more reliable and easy to work with. Of particular interest and importance in the present context is the demonstration of rapid locking and unlocking of an extended-cavity diode laser to a high-finesse optical cavity. This technology can greatly enhance the signal-to-noise ratio and data rate in cavity ring-down spectroscopy.

References

1. 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, 97–105 (1983).
2. M. Zhu and J. L. Hall, Stabilization of optical-phase frequency of a laser system—application to a commercial dye-laser with an external stabilizer, *J. Opt. Soc. Am. B* 10, 802–816 (1993).
3. H. R. Telle, Narrow linewidth laser diodes with broad, continuous tuning range, *Appl. Phys. B* 49, 217–226 (1989).
4. H. R. Telle, Stabilization and modulation schemes of laser diodes for applied spectroscopy, *Spectrochim. Acta Rev.* 15, 301–327 (1993).
5. A. S. Arnold, J. S. Wilson, and M. G. Boshier, A simple extended-cavity diode laser, *Rev. Sci. Instrum.* 69, 1236–1239 (1998).
6. P. Zorabedian and W. R. Trutna, Jr., Alignment-stabilized grating-tuned external-cavity semiconductor laser, *Opt. Lett.* 15, 483–485 (1990).

7. A. Schoof, J. Grunert, S. Ritter, and A. Hemmerich, Reducing the linewidth of a diode laser below 30 Hz by stabilization to a reference cavity with a finesse above 10^5 , *Opt. Lett.* 26, 1562–1564 (2001).
8. B. Dahmani, L. Hollberg, and R. Drullinger, Frequency stabilization of semiconductor lasers by resonant optical feedback, *Opt. Lett.* 12, 876–878 (1987).
9. H. Li and H. R. Telle, Efficient frequency noise reduction of GaAlAs semiconductor lasers by optical feedback from an external high-finesse resonator, *IEEE J. Quantum. Electron.* 25, 257–263 (1989).
10. H. Patrick and C. E. Wieman, Frequency stabilization of a diode laser using simultaneous optical feedback from a diffraction grating and a narrow-band Fabry–Pérot cavity, *Rev. Sci. Instr.* 62, 2593–2595 (1991).
11. S. Ohshima and H. Schnatz, Optimization of injection current and feedback phase of an optically self-locked laser diode, *J. Appl. Phys.* 71, 3114–3117 (1992).
12. A. E. Siegman, “Lasers.” University Science Books, Mill Valley, California, 1986.
13. C. Fabre, R. G. Devoe, and Brewer, R. G. Ultrahigh-finesse optical cavities, *Opt. Lett.* 11, 365–367 (1986).
14. C. Hutchinson, Ed., “The ARRL Handbook for Radio Amateurs.” The American Radio Relay League, Newington, Connecticut, 2002.
15. R. C. Dorf, “Modern Control Systems.” Addison-Wesley Publishing, Reading, Massachusetts, 1989.
16. P. Tremblay and R. Ouellet, Frequency response of a Fabry–Pérot interferometer used as a frequency discriminator, *IEEE Trans. Instrum. Meas.* 40, 204–207 (1991).
17. P. C. D. Hobbs, Ultrasensitive laser measurements without tears, *Appl. Opt.* 36, 903–920 (1997).
18. P. Dubé, L.-S. Ma, J. Ye, P. Jungner, and J. L. Hall, Thermally induced self-locking of an optical cavity by overtone absorption in acetylene gas, *J. Opt. Soc. Am. B* 13, 2041–2053 (1996).
19. D. S. Elliott, R. Roy, and S. J. Smith, Extracavity laser band-shape and bandwidth modification, *Phys. Rev. A* 26, 12 (1982).
20. R. W. Fox, Trace Detection with Diode Lasers. Ph.D. Thesis, University of Colorado, Boulder, 1995.
21. D. Romanini, A. A. Kachanov, N. Sadeghi, and F. Stoeckel, CW cavity ring-down spectroscopy, *Chem. Phys. Lett.* 264, 316–322 (1997).
22. B. A. Paldus, C. C. Harb, T. G. Spence, B. Wilke, J. Xie, J. S. Harris, and R. N. Zare, Cavity-locked ring-down spectroscopy, *J. Appl. Phys.* 83, 3991–3997 (1998).

23. C. R. Bucher, K. K. Lehmann, D. F. Plusquellic, and G. T. Fraser, Doppler-free nonlinear absorption in ethylene by use of continuous-wave cavity ring-down spectroscopy, *Appl. Opt.* 39, 3154–3164 (2000).
24. T. G. Spence, C. C. Harb, B. A. Paldus, R. N. Zare, B. Willke, and R. L. Byer, A laser-locked cavity ring-down spectrometer employing an analog detection scheme, *Rev. Sci. Instrum.* 71, 347–353 (2000).
25. A. A. Kosterev, A. L. Malinovsky, F. K. Tittel, C. Gmachl, F. Capasso, D. L. Sivco, J. N. Baillargeon, A. L. Hutchinson, and A. Y. Cho, Cavity ring-down spectroscopic detection of nitric oxide with a continuous-wave quantum-cascade laser, *Appl. Opt.* 40, 5522–5529 (2001).
26. M. D. Levenson, B. A. Paldus, T. G. Spence, C. C. Harb, R. N. Zare, M. J. Lawrence, and R. L. Byer, Frequency-switched heterodyne cavity ring-down spectroscopy, *Opt. Lett.* 25, 920–922 (2000).