Laser Stabilization

John L. Hall, Matthew S. Taubman, and Jun Ye

JILA, University of Colorado and National Institute of Standards and Technology
Boulder, Colorado 80309-0440

1. Introduction and Overview

For laser applications in which measurement precision is a key feature, frequency stabilized lasers are preferred, if not essential. This observation was true in the gas laser days when the 10^{-6} fractional Doppler width set the uncertainty scale. Now we have diode-pumped solid state lasers with fractional tuning range approaching 10^{-2} or more, and laser diode systems with several percent tuning. Such tuning is useful to find the exact frequency for our locking resonance, but then stabilization will be essential. Locking to cavities and atomic references can provide excellent stability, even using a widely tunable laser source. Indeed, laser frequency stability between independent systems has been demonstrated at 5 x10^{-14} in 1 s averaging time, and more than a decade better at 300 s. This incredible performance enhancement is possible because of a feedback from measurement of the laser's frequency error from our setpoint, this signal being fed into a filter/amplifier system and finally to an actuator on the laser itself which changes its frequency in response. While such feedback in response to performance may be the most important principle in evolution, in machines and lasers feedback enables the design of lighter, less costly systems. The accuracy is obtained, not by great bulk and stiffness, but rather by error correction, comparing the actual output against the ideal. This continuous correction will also detect and suppress the system's nonlinearity and noise. The performance limitation ultimately is set by imprecision of the measurement, but there is a lot of care required to get into that domain: we must have a very powerful correction effort to completely hide the original sins.

This article is our attempt to lead the worker newly interested in frequency control of lasers on a guided tour of stabilized lasers, ideally providing enough insight for recruiting yet another colleague into this wonderful arena. As nonlinear optics becomes just part of our everyday tools, the buildup cavities which enhance the nonlinear couplings are taking on a more critical role: this is the reason that we focus on the taming of PZT-based systems. We then cover locking with other transducers, and present some details about their construction and use. We consider the frequency discriminator, which is a key element for these control systems. The article concludes with description of the design and performance of several full practical systems.

Quantifying Frequency Stability

In thinking about the stability of our lasers, one may first wonder whether time or frequency domain pictures will be more powerful and instructive. Experience shows that time-domain perturbations of our lasers are usually associated with unwelcome sounds – door slamming, telephone bells, loud voices. Eventually these time-localized troubles can be eliminated. But what remains is likely the sum of zillions of smaller perturbations: none too conspicuous, but too many in number to attack individually. This perspective leads to a frequency-domain discussion where we can add the Fourier amplitudes caused by the many little sources. Eventually we are led to idealize our case to a continuum of spectrally-described perturbations. This physical outlook is one reason we will mainly be specifying our performance measures in the frequency domain.

Another important issue concerns the nifty properties of Mr. Fourier's description: in the frequency domain cascaded elements are represented by the multiplication of their individual transfer functions. If we had chosen instead the time domain, we would need to work with convolutions, nonlocal in time. Today's result in time is the sum of all previous temporal events that have the proper delay to impact us now. So it seems clear that frequency domain is good for analysis. What about describing the results?

Frequency vs. time: drift — the Allan Variance method

At the other end of our laser stabilization project, describing the results, it is convenient to measure and record the frequency as a function of time. We can measure the frequency averaged over an one-second gating time, for example, and stream 100 points to a file. This would be a good way to see the variations around a mean for the 1 s time intervals. This measurement could be repeated using a succession of gate times, 3 s, 10 s, 30 s 100 s. ... Surely it will be attractive to make this measurement just once and numerically combine the data to simulate the longer gate times. Thinking this way brings us a new freedom: we can process this data to recover more than just the mean and the standard deviation. Of course, we can expect to eventually see some drift, particularly over long times. When we look at the drift and slowly varying laser frequency, one wishes for a method to allow us to focus on the random noise effects which are still visible, even with the extended gate times. This is where the resonance physics is, while the drift is mainly due to technical problems. Dave Allan introduced the use of first differences, which has come to be called the Allan Variance method.[1] If we take the difference between adjacent samples of the measured frequency, we focus on the random processes which are averaged down to small, but not insignificant values within each gate time τ. These first differences (normalized by $1/\sqrt{2}$ to account for random noise in each entry) form a new data set which is first-order insensitive to long-term processes such as drift which dominate the directly recorded data.

Essentially the Allan Variance calculation presents us with a display of the laser's fractional frequency variation, σ_y , as a function of the time over which we are interested. At medium times, say τ of a few seconds, most laser stabilization systems will still be affected by the random measurement noise arising from shot noise and perhaps laser technical noise. At longer times the increased signal averaging implies a smaller residual fluctuation due to random processes. It is easy to show that the dependence of σ_y vs. τ can be expected to be $1/\tau^{1/2}$, in the domain controlled by random (white) noise. The Allan deviation also has a great utility in compressing our statement of laser stability: we might say, for example, "the (in-)stability is 2×10^{-12} at 1 s, with the $1/\tau^{1/2}$ dependence which shows that only random noise is important out to a time of 300 s."

Allan Deviation definition

With a counter linked to a computer, it is easy to gather a file of frequency values f_i measured in successive equal gate time intervals, t_g . Usually there is also some dead time, say t_d , while the counter-to-computer data transfers occur via the GPIB connection. This leads to a sample-to-sample time interval of $t_s = t_g + t_d$. Allan's variance is one half of the average squared difference between adjacent samples, and the usually quoted quantity, the Allan Deviation, is the square root of this averaged variance,

$$\sigma_y(\tau) = \left[\frac{1}{2(N-1)} \sum_{n=1}^{N-1} (f_{n+1} - f_n)^2 \right]^{1/2} . \qquad (1)$$

The dependence of σ_y upon the measuring time τ contains information essential for diagnosis of the system performance. These values for several times can be efficiently calculated from the (large) data set of frequencies observed for a fixed minimum gate time by adding together adjacent measurements to represent what would have been measured over a longer gate time. (This procedure neglects the effects of the small dead-time t_d, which are negligible for the white frequency noise $1/\sqrt{\tau}$ of usual interest.) Fewer samples will be available when the synthetic gate time becomes very long, so the uncertainty of this noise measurement increases strongly. Usually one insists on 3 or 4 examples to reduce wild variations, and so the largest synthetic gate time τ_{max} will be the total measurement time/3. For a serious publication we might prefer 5 or 10 such synthetic measurements for the last point on the graph.

The Allan deviation has one curiosity in the presence of a distinct sinusoidal modulation of the laser's frequency: when the gate time is 1/2 the sinusoid's period, adjacent samples will show the maximum deviation between adjacent measurements, leading to a localized peak in σ_y vs. τ . Interestingly, there will be "ghosts" or aliases of this when the gate time/modulation period ratio is 1/4, 1/8, etc. For longer gate times compared with the modulation, some fractional cycle memories can be expected also. So a clean slope of -1/2 for a log-log plot of σ_y vs. τ makes it clear that there is no big coherent FM process present.

Historically Allan's Variance has been valuable in locating time scales at which new physical processes must be taken into account. For example, at long times it is usual for a laser or other stable oscillator to reach a level of unchanging σ_y vs. τ . We speak of this as a "flicker" floor. It arises from the interplay of two opposite trends: the first is the decreasing random noise with increasing τ (decreasing σ_y vs. τ). At longer times one sees an increasing σ_y vs. τ, due to drifts in the many system parameters (electronic offsets, temperature…), which make our lasers lock at points increasingly offset from the ideal one. If we wait long enough, ever larger changes become likely. So for several octaves of time, the combination of one decreasing and one increasing contribution leads to a flat curve. Eventually significant drift can occur even within one measurement time, and this will be mapped as a domain of rising σ_y increasing as the +1 power of τ.

It is useful to note that the frequency/time connection of the Allan Variance transformation involves very strong data compression and consequently cannot at all be inverted to recover the original data stream in the way we know from the Fourier transform pair. However in the other direction, we can obtain the Allan Deviation from the Phase Spectral Density.[2]

Spectral Noise Density

As noted earlier, when the number of individual contributions to the noise becomes too large to enumerate, it is convenient to move to a spectral density form of representation. Two natural quantities to use would be the frequency deviations occurring at some rate and the narrow bandwidth within which they occur. To work with a quantity which is positive definite and has additive properties, it is convenient to discuss the squared frequency deviations $<(f^2_N)>$ which occur in a noise bandwidth B around the Fourier frequency f. This Frequency Noise Power Spectral Density, $S_f \equiv <(f^2_N)>/B$, will have dimensions of Hz^2 (deviation^2)/Hz (bandwidth). The summation of these deviations over some finite frequency interval can be done simply by integrating S_f between the limits of interest.

Connecting Allan Deviation and Spectral Density

Sometimes one can estimate that the system has a certain spectrum of frequency variations described by S_f (f), and the question arises of what Allan Deviation this would represent. We prefer to use the Allan presentation only for experimental data. However Ref. 2 indicates the weighted transform from S_f to Allan Variance.

Connecting linewidth and Spectral Density

A small surprise is that an oscillator's linewidth generally will not be given by the summation of these frequency deviations! Why? The answer turns on the interesting properties of Frequency Modulated (FM) signals. What counts in distributing power is the Phase Modulation Index, β , which is the ratio of the peak frequency excursion compared with the modulation rate. Speaking of pure tone modulation for a moment, we can write the phase-modulated field as

$$
\begin{aligned}
E(t) &= \sin(\Omega t + \beta \sin(\omega t)) \\
&= J_0(\beta)\exp(i\Omega t) + \sum_{n=1}^{\infty} J_n(\beta)\exp(i(\Omega + n\omega)t) + \sum_{n=1}^{\infty} J_n(\beta)(-1)^n \exp(i(\Omega - n\omega)t)
\end{aligned}
\qquad (2)
$$

where Ω is the "carrier" frequency, and $\omega = 2\pi$ f and its harmonics are the modulation frequencies. The frequency offset of one of these "sidebands", say the n-th one, is n times the actual frequency of the process' frequency f. The strength of the variation at such an n-th harmonic decreases rapidly for n > β according to the Bessel function $J_n(\beta)$. We can distinguish two limiting cases.

Large excursions, slow frequency rate. This is the usual laboratory regime with solid state or HeNe and other gas lasers. The perturbing process is driven by laboratory vibrations that are mainly at low frequencies. The extent of the frequency modulation they produce depends on our mechanical design, basically how efficient or inefficient an "antenna" have we constructed to pick up unwanted vibrations. Clearly a very stiff, lightweight structure will have its mechanical resonances at quite high frequencies. In such case both laser mirrors will track with nearly the same excursion, leading to small differential motion, i.e. low pickup of the vibrations in the laser's frequency. Heavy articulated structures, particularly mirror mounts with soft springs, have resonances in the low audio band and lead to big FM noise problems. A typical laser construction might use a stiff plate, say 2 inches thick of Al or honeycomb-connected steel plates. The mirror mounts would be clamped to the plate, and provide a laser beam height of 2 inches above the plate. Neglecting air pressure variations, such a laser will have vibration-induced excursions $(\langle f^2_N \rangle)^{1/2}$ of << 100 kHz. An older concept used low expansion rods of say 15 mm diameter Invar, with heavy Invar plates on the ends, and kinematic but heavy mirror mounts. This system may have a vibration-induced linewidth $(\langle f^2_N \rangle)^{1/2}$ in the MHz range. Only when the "rods" become several inches in diameter is the axial and transverse stiffness adequate to suppress the acceleration induced forces. With such massive laser designs we have frequency excursions of 10's to thousands of kHz, driven by vibrations in a bandwidth B < 1 kiloHertz. In this case $(\langle f^2_N \rangle)^{1/2} \gg B$, and the resulting line shape is Gaussian. The linewidth is given by Ref. 3, $\Delta f_{FWHM} = [8 \ln(2) (\langle f^2_N \rangle)]^{1/2} \cong 2.355 (\langle f^2_N \rangle)^{1/2}$.

The broadband fast, small excursion limit. This is the domain in which we can usually end up if we can achieve adequate servo gain to reduce the vibration-induced FM. Since the drive frequency is low, it is often feasible to obtain a gain above 100, particularly if we use a speedy transducer such as an AOM or an EOM. In general we will find a noise floor fixed, if by nothing else than the broadband shot noise which forms a minimum noise level in the measurement process. Here we can expect small frequency excursions at a rapid rate, $(\langle f^2_N \rangle)^{1/2} \ll B$, leading to a small phase modulation index. If we approximate that the Spectral Noise Frequency Density $S_f = (\langle f^2_N \rangle)/B$ is flat, with the value S_f Hz^2 (deviation^2)/Hz (bandwidth), then the linewidth in this domain is Lorentzian,[3] with the $\Delta f_{FWHM} = \pi S_f = \pi (\langle f^2_N \rangle)/B$.

This summary of frequency-domain measures is necessarily brief and the interested reader may find additional discussion useful.[3,4,5] A number of powerful consequences and insights flow from reworking the above discussions in terms of a Phase Noise Power Spectral Density, $S_\phi = S_f / f^2$. The NIST Frequency and Time Division publishes collections of useful tutorial and overview articles from time to time. The currently available volume[2] covers these topics in more detail. Vendors of rf-domain spectrum analyzers also have useful application notes.[6]

2. Servo Principles and Issues[7,8]

Bode representation of a Servo System

We will describe our systems by transfer functions, output/input, as a function of Fourier frequency ω. We begin purely in the domain of electronics. The amplifier gain is G(ω). The electrical feedback is represented as H(ω). Both will have voltage as their physical domain, but are actually dimensionless in that they are output/input ratios. Considering that we will have to represent phase of these AC signals, both G(ω) and H(ω) will generally be complex. It will be fundamental to view these functions with their dependence on frequency, for both the amplitude and phase response.

Imagine a closed loop system with this amplifier as the forward gain G(ω) between input V_i and output V_o. Some fraction of the output is tapped off and sent back to be compared with the actual input. For more generality we will let H(ω) represent this feedback transfer ratio. The actual input, minus this sampled output will be our input to our servo amplifier G(ω). After a line of algebra we find the new gain of the closed loop – in the presence of feedback – is

$$A_{cl} = \frac{V_0}{V_i} = \frac{G(\omega)}{1+G(\omega)H(\omega)} \quad . \qquad (3)$$

A particularly instructive plot can be made for the product G(ω)H(ω), called the "open loop gain," which appears in the denominator. In this so-called Bode plot, the gain and phase are separately plotted. Also, from inspection of Eq. 3 we can learn one of the key advantages which feedback brings us: if the feedback factor GH were >>1, the active gain G would basically cancel out and we would be left with $A_{cl} \sim 1/H$. We imagine this feedback channel will be passive, formed from nearly ideal non-distorting components. The noise, exact value of the gain, and distortion introduced by it are seen to be nearly unimportant, according to the large magnitude of 1 + GH. Gentle amplifier overload will lead to overtone production, but could alternatively be represented by a decrease of G with signal. Since the output doesn't depend sensitively upon G anyway, we are sure these distortion products and internally-generated noise will be suppressed by the feedback. We can identify the denominator 1 + GH as the noise and distortion reduction factor.

What is the cost of this reduced dependence on the active components G(ω) and their defects? Basically it is that the gain is reduced and we must supply a larger input signal to obtain our desired output. For a music system one can then worry about the distortion in the pre-amplifier system. However, we want to make quiescent lasers, without the slightest hint of noise. So it is nice that the amplification of internal noise is reduced.

To be concrete, the circuit of Fig. 1 represents a common building block in our servo design, and represents a simple case of feedback. We show it as a current summing input node: the subtraction at the input happens here because the sign of the gain is negative. With the nearly ideal high gain operational amplifiers now available, G>>1 and we can closely approximate the closed loop gain by 1/ H(ω), yielding a flat gain above and a rising gain below some corner frequency $\omega_0 = 1/\tau_0$, with $\tau_0 = R_f C$. Remember 1/H(ω) is the closed loop gain between V_o and V_i. To find the exact relationship between the signal V_s and V_o, we notice the related voltage divider effect gives $V_i = (1 - H(\omega))V_s$, which leads to

$$\frac{V_o}{V_s} = -\frac{R_f}{R_i}\frac{\left(1 + j\,\omega/\omega_0\right)}{j\,\omega/\omega_0} = -\frac{R_f}{R_i}\frac{\left(1 + j\omega\tau_0\right)}{j\omega\tau_0} \quad . \tag{4}$$

The negative sign arises from the fact that the forward gain is negative. When the corner frequency ω_0 is chosen to be sufficiently high, we may have to consider the bandwidth issue of the OpAmp: G(ω) could start to roll off and no longer satisfy the approximation of G>>1. A more complex network is needed to compensate for the gain roll off and that is exactly the topic of feedback we wish to cover below.

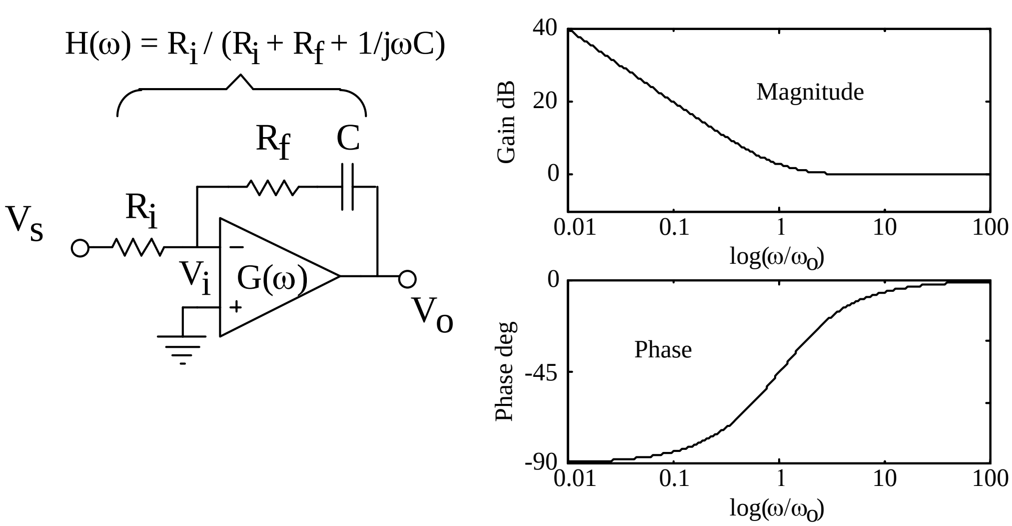

Fig. 1. Phase and amplitude response of a Proportional-Integral (PI) Amplifier Circuit. The PI function is implemented using an inverting OpAmp.

Phase and amplitude responses vs. frequency

We can plot[9] the gain magnitude and phase of this elementary feedback example in Fig. 1, where we can see the flat gain at high frequencies and the rising response below ω_0 . Our laser servo designs will need to echo this shape, since the drift of the laser will be greater and greater at low frequencies, or as we wait longer. This will require larger and larger gains at low frequencies (long times) to keep the laser frequency nearby our planned lock point. The phase in Fig. 1 shows the lag approaching 90° at the lowest frequencies. (An overall minus sign is put into the subtractor unit, as our circuit shows an adder.) The time domain behavior of this feedback system is a prompt inverted output, augmented later by the integration's contribution.

As a first step toward modeling our realistic system, Fig. 2 shows the laser included in our control loop. The servo system's job is to keep the laser output at the value defined by the reference or setpoint input. Some new issues will arise at the high frequency end with the physical laser, as its PiezoElectric Transducer (PZT) will have time delay, finite bandwidth and probably some resonances.

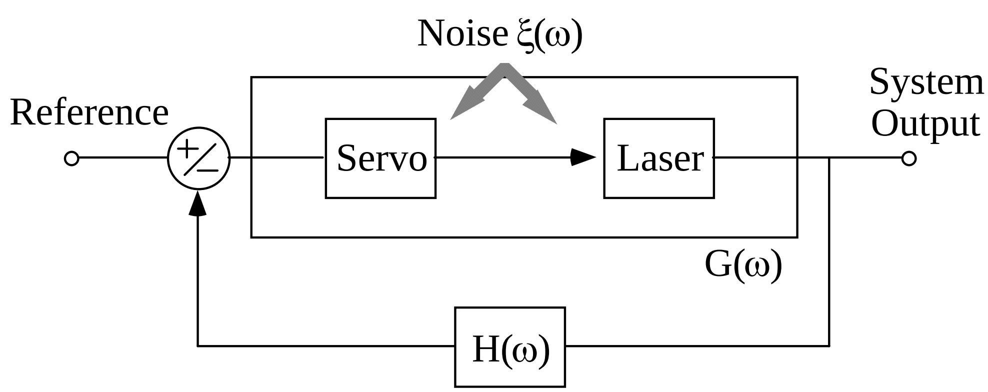

Fig. 2. Schematic model of laser system, with frequency noise included, as part of a servo control loop.

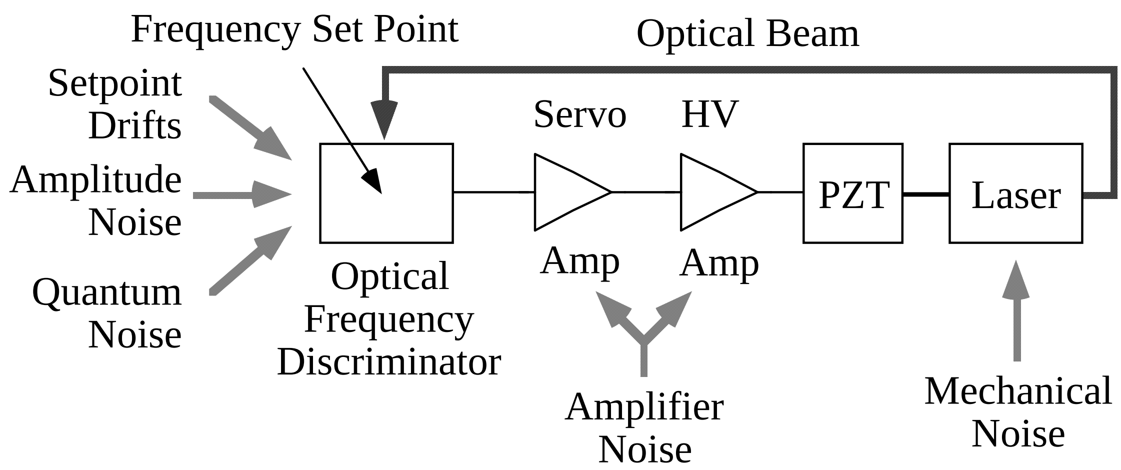

Fig. 3. Detailed model of frequency-controlled laser.

One way we should expand the model is to include the laser's operation as a frequency transducer, converting our control voltage to a frequency change. Probably the laser will have some unwanted frequency noises, and in Fig. 3 we can indicate their unwanted contributions arriving in the optical frequency discriminator, which functions like a summing junction. The emitted laser field encounters an optical frequency discriminator and the laser frequency is compared with the objective standard, which we will discuss below. In our diagram we show this laser frequency discriminator's role as an optical frequency–to–voltage converter element. More exactly, laser frequency *differences* from the discriminator's reference setpoint are converted to a voltage output. Laser amplitude noises (due to the intrinsic property of the laser itself or external beam propagation) and vibration effects on the discriminator will appear as undesired additive noises also.

The first simple idea is that the feedback loop can be closed when the servo information, carried as a voltage signal in our amplifier chain, is converted to a displacement (in m) by the PZT, then into laser frequency changes by the laser's standing-wave boundary condition. As the length changes, the "accordian" in which the waves are captive is expanded or compressed, and along with it the wavelength and frequency of the laser's light.

A second truth becomes clear as well: there is freedom in designating the division into the forward gain part and the feedback path part. Actually, we probably would like the laser to be tightly locked onto the control cavity/discriminator, and then we will tune the whole system by changing the set-point which is the discriminator's center frequency. This leads us to view the optical frequency discriminator as the summing junction, with the amplifier and PZT transducer as the forward gain part. The output is taken as an optical frequency, which would be directly compared to the set-point frequency of the discriminator. So the feedback path H = 1.

We should consider some magnitudes. Let K_{pzt} represent the tuning action of the PZT transducer, expressed as displacement (m)/Volt. A typical value for this would be K_{pzt} = 0.5 nm/Volt. The laser tunes a frequency interval c/2L for a length change by $\lambda/2$, so the PZT tuning will be ~ 600 Volts/order at 633 nm.

$$K_V = K_{PZT}\frac{2}{\lambda}\frac{c}{2L} \quad (5)$$

So we obtain a tuning sensitivity K_V ~ 800 kHz/Volt tuning for a foot-long laser, assuming a disk-type PZT geometry. See the section below on PZT design.

Measurement Noise as a performance limit -- it isn't

Usually our desire for laser stability exceeds the range of the possible by many orders, and we soon wonder about the ultimate limitations. Surely the ultimate limit would be due to measurement noise. However, we rarely encounter the shot-noise limited case, since the shot noise limited S/N of a 100 μW locking signal is ~6 $\times 10^6$ in a 1 Hz bandwidth. (See section on Cavity Frequency Discriminators, below.) Rather we are dealing with the laser noise remaining because our servo gain is inadequate to reduce the laser's intrinsic noise below the shot noise limit, the clear criterion of gain sufficiency. So our design task is to push up the gain as much as possible to reduce the noise, limited by the issue of *stability* of the thus-formed servo system.

Servo stability: Larger gain at lower frequencies, decreasing to unity gain and below…

Our need for high gain is most apparent in the low-frequency domain ~ 1 kHz and below. Vibrations abound in the dozens to hundreds of Hz domain. Drifts can increase almost without limit as we wait longer or consider lower Fourier frequencies. Luckily, we are allowed to have more gain at low frequencies without any costs in stability. At high frequencies, it is clear we will not help reduce our noise if our correction is applied too late and so no longer has an appropriate phase. One important way we can characterize the closed-loop behavior of our servo is by a time delay t_{delay} . Here we need to know the delay time before any servo response appears; a different (longer) time characterizes the full scale or 1/e response. The latter depends on the system gain, while the ultimate high-speed response possible is controlled by the delay until the first action appears. A good criterion is that the useful unity gain frequency can be as high as $f_\tau = 1/2\pi\, t_{delay}$, corresponding to 1 rad extra phase-shift due to the delay. Below this ultimate limit we need to increase the gain - increase it a lot - to effectively suppress the laser's increased noise at low frequencies. This brings us to address the closed-loop stability issue.

Closed-loop Stability Issues

One can usefully trace the damping of a transient input as it repetitively passes the amplifier and transducer, and is re-introduced into the loop by the feedback. Evidently stability demands that the transient is weaker on each pass. The settling dynamics will be more beautiful if the second-pass version of the perturbation is reduced in magnitude and is within say ±90° of the original phase. Ringing and long decay times result when the return phasor approaches -1 times the input signal vector, as then we are describing a sampled sinewave oscillation. These time domain pictures are clear and intuitive, but require treatment in terms of convolutions, so here we will continue our discussion from the frequency-domain perspective that leads to more transparent algebraic forms. We can build up an arbitrary input and response from a summation of sinusoidal inputs. This leads to an output as the sum of corresponding sinusoidal outputs, each including a phase shift.

In our earlier simple laser servo example, no obvious limitation of the available closed-loop gain was visible. The trouble is we left out two fundamental laboratory parasites: time delay, as just noted, and mechanical resonances. We will usually encounter the mechanical resonance problem in any servo based on a PZT transducer. For design details, see the "Practical Issues" section below. A reasonable unit could have its first longitudinal resonance at about 25 kHz, with a Q ~10. In servo terms, the actual mechanical PZT unit gives an added 2-pole roll-off above the resonance frequency and a corresponding asymptotic phase lag of 180°. Including this reality in our model adds another transfer function $R_{PZT} = \omega_0^{\,2}/(\omega_0^{\,2} + 2\omega\,\eta\,\omega_0 + \omega^2)$, where ω_0 is 2π times the resonance frequency, and $\eta = 1/2Q$ is the damping factor of the resonance. This response is shown below in Fig. 4.

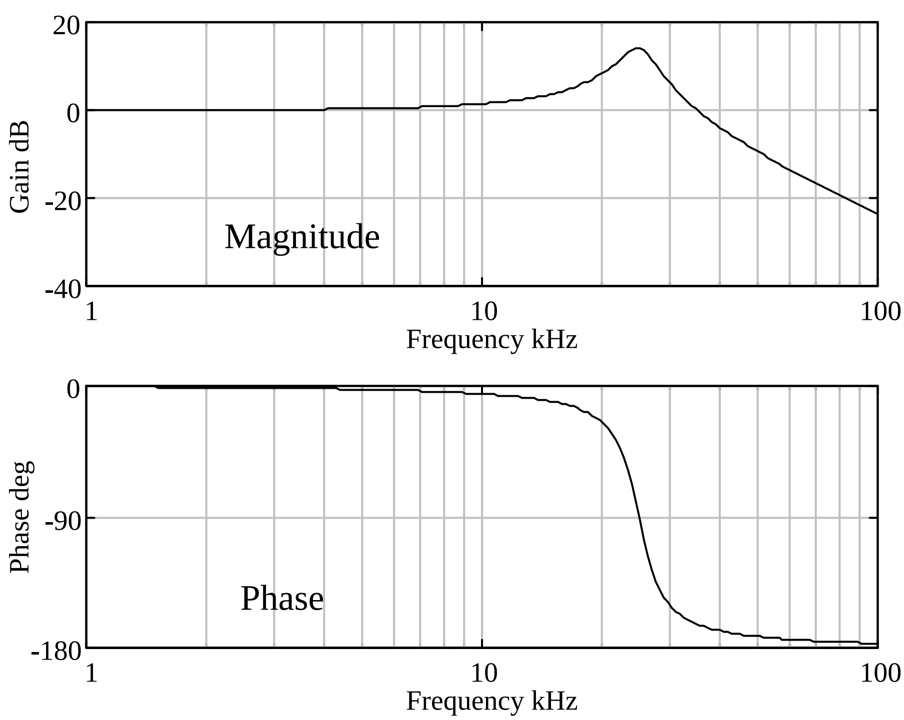

Fig. 4. The amplitude and phase response of a tubular PZT transducer and an 8mm-diameter by 5mm-thick mirror. The resonance is at 25 kHz with a Q of 10. (This mechanical design is discussed below in Section 3.)

We now talk of stabilizing this system. The first appealing option is to try a pure integrator. The problem then is that we are limited in gain by the peak height of the resonance which must remain entirely below unity gain to avoid instability. In Fig. 5 case (a) we see that the unity gain frequency is limited to a value of 1.5 kHz. Some margin is left to avoid excessive ringing near the resonant frequency, but it is still visible in the time domain. Techniques that help this case include a roll-off filter between the unity gain and PZT resonance frequencies.

Fig. 5 shows the "open loop" gain function GH of the feedback equation, and the corresponding phase response. We already noted the dangerous response function of -1 where the denominator of Eq. (3) vanishes. In the time domain iterative picture, the signal changes sign on successive passes and leads to instability/oscillation. We need to deal with care as we approach near this point in order to obtain maximum servo gain: it is useful to consider two stability margins. The **phase stability margin** is the phase of the open loop function when the gain is unity. It needs to be at least 30°. The **gain margin** is the closed loop gain when the phase is 180 degrees. In Fig. 5 case (a) we see that the phase is not shifted very much until we really "sense" the amplitude increase from the resonance. So this resonance may tend to fix an apparently solid barrier to further servo improvement. But as shown in Fig 5 case (b), just a low-pass to push down the PZT resonance is very helpful.

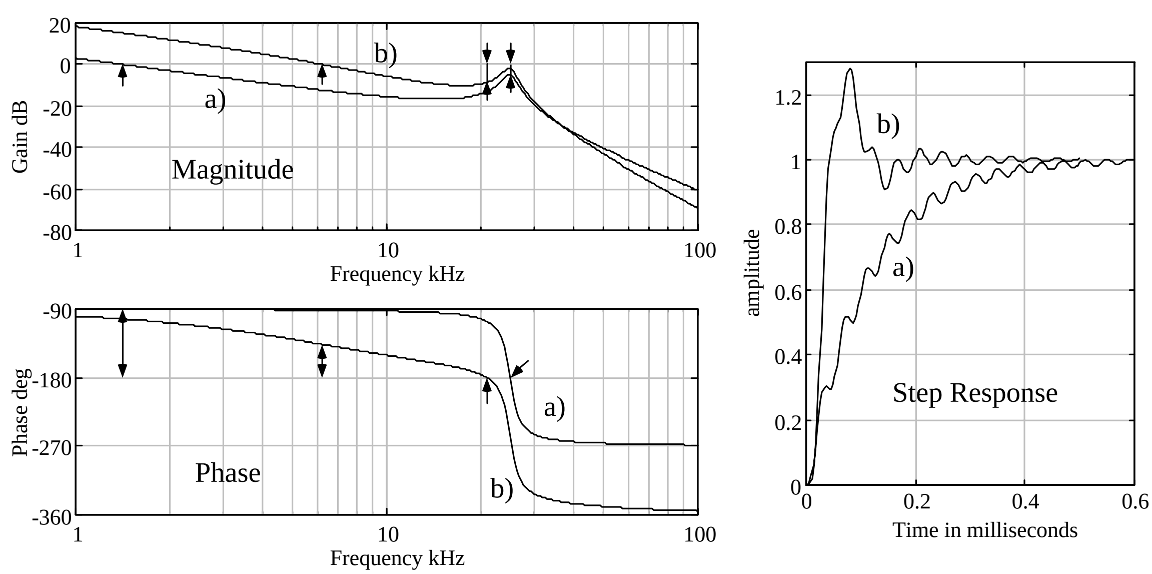

Fig. 5. a) Integrator gain function alone. Gain must be limited so that gain is <1 even at the resonance. b) Single-pole low pass at 6 kHz inserted. Now unity gain can increase to 6 kHz and time response is ~3-fold faster. Small arrows in the graph indicate the phase

margin at the unity gain frequency (gain = 0 dB) and gain margin at a phase shift of -180^0.

In fact, there are many ways of improving the low frequency gain of this system. They include imposing yet another high frequency roll-off (or multi-pole low pass filter) just before the resonance thus pushing its height down and allowing the open loop transfer function to come up; adding lag compensators before the resonance to push the low frequency gain up while keeping the high frequency response relatively unchanged; adding lead compensator just above the resonance to advance the phase and increase the unity gain point; or placing a notch at the resonant frequency to "cut it out" of the open loop transfer function. The last two options in this list are quite promising and are discussed in more detail below.

Proportional Integral Derivative (PID) Controller versus Notch Filters

Like many "absolute" barriers, it is readily possible to shoot ahead and operate with a larger closed loop bandwidth than that represented by the first PZT resonance. The issue is that we must control the lagging phase that the resonance introduces. A good solution is a Differentiator stage, or a phase lead compensator, which could also be called a high frequency boost gain step circuit. In Fig. 6 case (a) we show the Bode plot of our PZT-implemented laser frequency servo, based on a PID (Proportional Integral Differentiator) controller design. Just a few moments of design pays a huge benefit, as the unity gain frequency has now been pushed to 40 kHz, almost a factor of 2 *above* the PZT's mechanical resonance. For this PID controller example, unity gain occurs at a 7-fold increased frequency compared with Fig. 5 case (b). Thus at the lower frequencies we would hope to have increased the servo gain by a useful factor of 7x or 17 dB. However comparison of Fig 5 (b) and 6 (a) shows that the low frequency gain is hardly changed, even though we greatly increased the servo bandwidth.

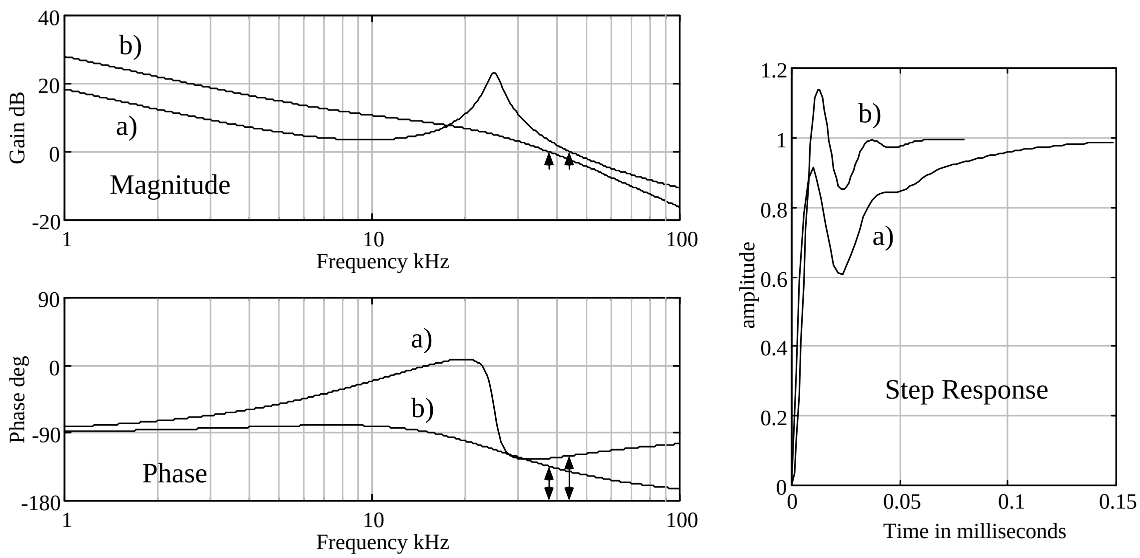

Fig 6. Two methods of working through and beyond a resonance. a) PID controller where the Derivative term advances the phase near the resonance. b) Adding a notch is a better approach, where the notch function approximates the inverse of the resonance peak. Transient response settles much more quickly. Again, we use the small arrows in the graph to indicate the phase margin at the unity gain frequency.

How *do* we go forward? We could in principle continue to increase the gain and unity gain frequency, but this is not really practical however, since we will again be limited by additional structure resonances that exist beyond the first resonance. Also, the Derivatives needed to tame these resonances cost low frequency gain, and it is hard to win. To make progress, we use a notch as an alternative technique to suppress the resonance. Now a D term is not needed, and we can conserve the gain at low frequencies. The notch filter, combined with a PI stage, gives unity gain at higher frequencies, and increases gain for low ones. See Fig. 6 case (b). Then Fig. 7 compares adding another PI stage to the two cases of Fig. 6, PID in (a) and Notch plus PI in (b). The time domain approach, shown in Fig. 7, shows case (b) settles rather nicely. AND the gain has increased more than 20 dB at frequencies of 1 kHz and below. So this is very encouraging.

While we have come to the cascaded-integrators approach cautiously in this discussion, in fact at least 2 integrators would always be used in practice. Workers with serious gain requirements, for example the LIGO and VIRGO gravitational wave detector groups, may use the equivalent of 4 cascaded integrators! Such a design is “conditionally stable” only, meaning that the gain cannot be smoothly reduced or increased. Such aggressive stabilizer designs have their place, but not for a first design!

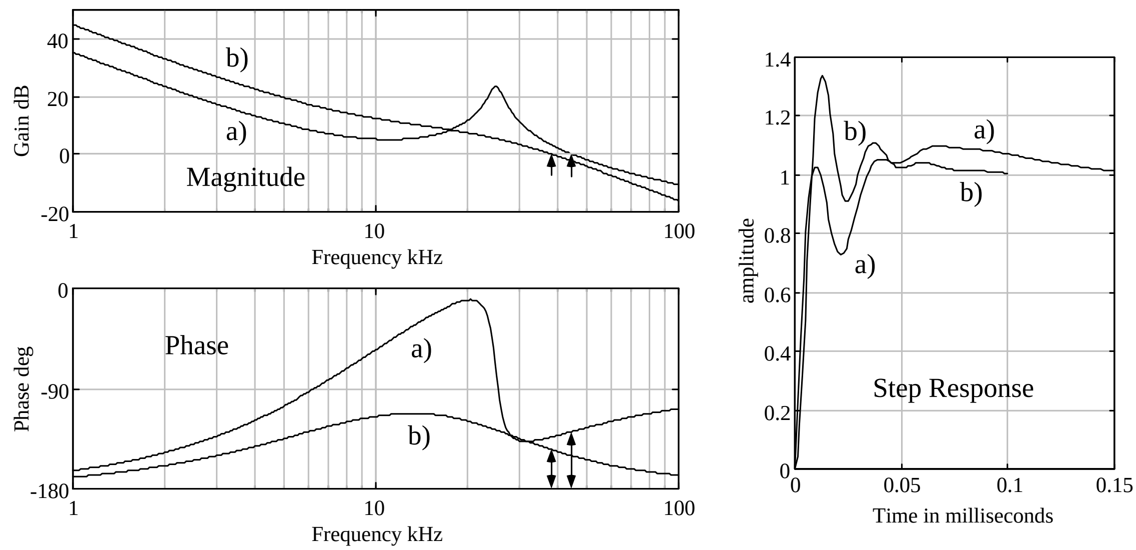

Fig 7. Adding an additional PI stage to (a) the PID and (b) the PI-plus-notch stabilizers of Fig. 6. Note that low frequency gain is strongly increased.

"rule of thumb" PID design for system with a transducer resonance

Optimizing servo performance is an elegant art, turned into science by specification of our "cost function" for the system performance shortcomings. In the case that we wish to minimize the time-integrated magnitude of the residuals following a disturbance, one comes to the case studied by Ziegler and Nichols for the PID controller used in a system with a combined roll-off and time delay.[8] Such a case occurs also in thermal controllers. With only the P term, one first looks for the frequency f_{osc} where the system first oscillates when the gain is increased. The PD corner is then set 1.27x higher than this f_{osc}, the P gain is set at 0.6 of the oscillation gain, and the PI corner is set at 0.318 times the oscillation frequency. This "rule of thumb" design of of the phase compensation produces a transient response which settles reasonably well, so as to minimize the Time Integrated Error. For phase-locking lasers, a cost function with more emphasis on long-lasting errors leads to another kind of "optimum" tuning, but with qualitatively similar results.

When a notch is used to suppress the resonance, there is no longer an anomalous gain at the resonant frequency and one is returned to the same case as in its absence. A reasonable servo approach to using two PI stages is to design with only one, achieving the desired unity gain frequency. The second PI is than added to have its corner frequency at this same point or up to 10-fold lower in frequency, depending on whether we wish the most smooth settling or need the highest feasible low frequency gain. The Figures 6b and 7 shows the Bode plot of such designs, along with the system's closed-loop transient response. An elegant strategy is to use adaptive clamping to softly turn on the extra stage when the error is small enough, thus dynamically increasing the order of the controller when it will not compromise the dynamics of recovery.

3. Practical Issues

Here we offer a number of important tidbits that are useful background material for a successful application of the grand schemes discussed above.

Frequency Discriminators for Laser Locking

So far we have devoted our main effort addressing the issues of the feedback scheme. Of equal importance is the subject of frequency reference system. After all, a good servo eliminates intrinsic noises of the plant (laser), and replaces them with the measurement noise associated with the reference system. Indeed, development of prudent strategies in high precision spectroscopy and the progress of laser stabilization have been intimately connected to each other through the years,[10] with the vigorous pursuit of resolution and sensitivity resulting in amazing achievements in both fields.

To stabilize a laser, one often employs some kind of resonance information to derive a frequency/phase-dependent discrimination signal. The resonance can be of material origin, such as modes of an optical interferometer; or of natural origin, such as atomic or molecular transitions. If the desired quantity of a stabilized laser, an optical frequency standard, is its long term stability or reproducibility, the use of a natural resonance is preferred. Reproducibility is a measure of the degree to which a standard repeats itself from unit to unit and upon different occasions of operation. The ultimate reproducibility is limited to the accuracy of our knowledge of the involved transitions of free atoms or molecules. The term "free" means the resonance under study has a minimum dependence on the laboratory conditions, such as the particle moving frame (velocity), electromagnetic fields, collisions, and other perturbations. To realize these goals, modern spectroscopy has entered the realm of quantum limited measurement sensitivities and exquisite control of internal and external degrees of freedom of atomic motions.

A careful selection of a high quality resonance can lead to superior system performance and high working efficiency. For example, the combined product of the transition quality factor Q and the potential signal-to-noise ratio (S/N) is a major deciding factor, since this quantity controls the time scale within which a certain measurement precision (fractional frequency) can be obtained. This importance is even more obvious when one considers the waiting time for a systematic study is proportional to the inverse square of (Q x S/N). A narrower transition linewidth of course also helps to reduce the susceptibility to systematic errors. The resonance line shape is another important aspect to explore. By studying the line shape we will find out whether we have come to a complete understanding of the involved transition and whether there are other unresolved small lines nearby ready to spoil our stabilization system.

Sometimes it may be not sufficient to use the natural resonance alone for stabilization work, or may not be necessary. The saturation aspect of the atomic transition limits the attainable S/N. To stabilize a noisy laser we need to use, for example, an optical resonator, which can provide a high contrast and basically unlimited S/N of the resonance information. Careful study of the design and control of the material properties can bring the stability of material reference to a satisfactory level. See below for a more detailed discussion on this topic.

Ideally, a resonance line shape is even symmetric with respect to the center frequency of the resonance, and deviations from this ideal case will lead to frequency offsets. However, for the purpose of feedback, the resonance information needs to be converted to an odd symmetric discriminator shape: we need to know in which direction the laser is running away from the resonance. A straightforward realization of an error signal using direct absorption technique is to have the laser tuned to the side of resonance.[11] The slope of the line is used to convert the laser frequency noise to amplitude information for the servo loop. This technique is essentially a DC approach and can suffer a huge loss in S/N due to the low frequency amplitude noise of the laser. A differential measurement technique using dual beams is a requirement if one wishes to establish a somewhat stable operation. With a dual beam approach, the information about the laser noise can be measured twice and therefore it is possible to completely eliminate the technical noise and approach the fundamental limit of shot noise using clever designs of optoelectronic receivers. Conventional dual beam detection systems use delicate optical balancing schemes,[12] which are often limited by the noise and drift of beam intensities, residual interference fringes, drift in amplifiers, and spatial inhomogenity in the detectors. Electronic auto cancellation of the photo detector currents has provided near shot noise limited performance.[13] Although this process of input normalization helps to increase S/N of the resonance, the limitation on the locking dynamic range remains as a problem. The servo loop simply gets lost when the laser is tuned to the tail or over the top of the resonance. Further, it is found that transient response errors basically limit the servo bandwidth to be within the cavity linewidth.[14] Another effective remedy to the DC measurement of resonances is the use of zero-background detection techniques, for example, polarization spectroscopy.[15,16] In polarization spectroscopy the resonance information is encoded in the differential phase shifts between two orthogonally polarized light beams. Heterodyne detection between the two beams can reveal an extremely small level of absorption-induced polarization changes of light, significantly improving the detection sensitivity. However, any practical polarizer has a finite extinction ratio (ε) which limits the attainable sensitivity. Polarization spectroscopy reduces the technical noise level by a factor of $\sqrt{\varepsilon}$, with $\varepsilon \sim 10^{-7}$ for a good polarizer. Polarization techniques do suffer the problem of long term drifts associated with polarizing optics.

Modulation techniques are of course often used to extract weak signals from a noisy background. Usually noises of technical origins tend to be more prominent in the low frequency range. Small resonance information can then be encoded into a high frequency region where both the source and the detector possess relatively small noise amplitudes. Various modulation schemes allow one to compare on-resonant and off-resonant cases in quick succession. Subsequent demodulations (lock-in detection) then simultaneously obtain and subtract these two cases, hence generating a signal channel with no output unless there is a resonance. Lorentzian signal recovery with the frequency modulation method has been well documented.[17] The associated lock-in detection can provide the first, second, and third derivative type of output signals. The accuracy of the modulation waveform can be tested and various electronic filters can be employed to minimize nonlinear mixing among different harmonic channels and excellent accuracy is possible. In fact, the well-established 633 nm HeNe laser system[18] is stabilized on molecular iodine transitions using this frequency dither technique and third harmonic (derivative) signal recovery. Demodulation at the third or higher order harmonics helps to reduce the influence of other broad background features.[19] The shortcoming of the existence of dither on the output beam can be readily cured with an externally implemented "un-dithering" device based on an AOM.[20] However, in this type of modulation spectroscopy the modulation frequency is often chosen to be relatively low to avoid distortions on the spectral profile by the auxiliary resonances associated with modulation-induced spectral sidebands. An equivalent statement is that the line is distorted because it can not reach an equilibrium steady state in the face of the rapidly-tuning excitation. This low-frequency operation (either intensity chopping or derivative line shape recovery) usually is still partly contaminated by the technical noise and the achievable signal-to-noise ratio (S/N) is thereby limited. To recover the optimum signal size, large modulation amplitudes (comparable to the resonance width) are also employed, leading to a broadened spectral linewidth. Therefore the intrinsic line shape is modified by this signal recovery process and the direct experimental resolution is compromised.

A different modulation technique was later proposed and developed in the microwave magnetic resonance spectroscopy and similarly in the optical domain.[21,22,23] The probing field is phase-modulated at a frequency much larger than the resonance linewidth under study. When received by a square-law photodiode, the pure FM signal will generate no photo-current at the modulation frequency unless a resonance feature is present to upset the FM balance. Subsequent heterodyne and rf phase-sensitive detection yield the desired signal. The high sensitivity associated with the FM spectroscopy is mainly due to its high modulation frequency, usually chosen to lie in a spectral region where the amplitude noise level of the laser source approaches the quantum (shot noise) limit. The redistribution of some of the carrier power to its FM sidebands causes only a slight penalty in the recovered signal size. Another advantage of FM spectroscopy is the absence of linewidth broadening associated with low-frequency modulation processes. The wide-spread FM spectra allows each individual component to interact with the spectral features of interest and thereby preserves the ultrahigh resolution capability of contemporary narrow-linewidth lasers.

Since its invention, FM spectroscopy has established itself as one of the most powerful spectroscopic techniques available for high sensitivity, high resolution and high-speed detection. The high bandwidth associated with the radio frequency (rf) modulation enables rapid signal recovery, leading to a high Nyquist sampling rate necessary for a high bandwidth servo loop. The technique has become very popular in nonlinear laser spectroscopy,[24] including optical heterodyne saturation spectroscopy,[23] two-photon spectroscopy,[25] Raman spectroscopy,[26] and heterodyne four-wave mixing.[27] Recent developments with tunable diode lasers have made the FM technique simpler and more accessible. The field of FM-based laser diode detection of trace gas and remote sensing is rapidly growing. In terms of laser frequency stabilization, the rf sideband based Pound-Drever-Hall locking technique[28] has become a uniformly adopted fast stabilization scheme in the laser community. The error signal in a high-speed operating regime is shown to correspond to the instantaneous phase fluctuations of the laser, with the atom or optical cavity serving the purpose of holding the phase reference. Therefore a properly designed servo loop avoids the response time of the optical phase/frequency storage apparatus and is limited only by the response of frequency correcting transducers.

In practice some systematic effects exist to limit the ultimate FM sensitivity and the resulting accuracy and stability. Spurious noise sources include residual amplitude modulation (RAM), excess laser noise, and étalon fringes in the optical system.[29] A number of techniques have been developed to overcome these problems. In many cases FM sidebands are generated with electro-optic modulators (EOM). A careful design of EOM should minimize the stress on the crystal and the interference between the two end surfaces (using angled incidence or anti-reflection coatings). Temperature control of the EOM crystal is also important and has been shown to suppress the long-term variation of RAM.[30] The RAM can also be reduced in a faster loop using an amplitude stabilizer[31] or a tuning filter cavity.[32] The étalon fringe effect can be minimized by various optical or electronic means.[33] An additional low-frequency modulation (two-tone FM[34]) can be used to reduce drifts and interference of the demodulated baseline.

In closing this section we note that a laser is not always stabilized to a resonance but is sometimes referenced to another optical oscillator.[35] Of course the working principle does not change: one still compares the frequency/phase of the laser with that of reference. The technique for acquiring the error information is however more straightforward, often with a direct heterodyne detection of the two superposed waveforms on a fast photo detector. The meaning of the fast photo detector can be quite extensive, sometime referring to a whole table-top system that provides THz-wide frequency gap measurement capabilities.[36,37,38] Since it is the phase information that is detected and corrected, an optical phase locked loop usually provides a tight phase coherence between two laser sources. This is attractive in many measurement applications where the relative change of optical phase is monitored to achieve a high degree of precision. Other applications include phase-tracked master-slave laser system where independent efforts can be made to optimize laser power, tunability and intrinsic noise.

the Optical Cavity-based Frequency Discriminator

It is difficult to have both sensitive frequency discrimination and short time delay, unless one uses the reflection mode of operation: these issues have been discussed carefully elsewhere.[28] With ordinary commercial mirrors, we can have a cavity linewidth of 1 MHz, with a contrast C above 50%. We can suppose using 200 μW optical power for the rf sideband optical frequency discriminator, leading to a dc photo current i_0 of ~100 μA and a signal current of ~25 μA. The shot noise of the dc current is $i_n = \sqrt{2ei_0}$ in a 1 Hz bandwidth, leading to a S/N of ~ 4 x10^6 . The frequency noise-equivalent would then be 250 milliHertz/√Hz. If we manage to design enough useful gain in the controller to suppress the laser's intrinsic noise below this level, the laser output frequency spectrum would be characterized by this power spectral density. Under these circumstances the output spectrum would be Lorentzian, of width $\Delta\nu_{FWHM} = \pi S_f = \pi\,(0.25\text{ Hz})^2/\text{Hz} \sim 0.8$ Hz . One comes to impressive predictions in this business! But usually the results are less impressive.

What goes wrong? From measurements of the servo error, we can see that the electronic lock is very tight indeed. However, the main problem is that vibrations affect the optical reference cavity's length and hence its frequency. For example measurements show the JILA Quiet Room floor has a horizontal seismic noise spectrum which can be approximated by 4 x10^{-9} m rms/√Hz from below 1 Hz to about 20 Hz, breaking there to an f^{-2} roll-off. Below 1 Hz the displacement noise climbs as f^{-3}. Accelerations associated with these motions lead to forces on the reference cavity that will lead to mechanical distortion and hence frequency shifts. In the axial direction, holding the cavity in the mid-plane seems wise as the net length change would tend be cancelled: one half is under compression, the other half is under tension at a particular moment in the ac vibration cycle. We denote this cancellation by symmetry as ε, with $0 \le \varepsilon \le 1$. The asymmetry value observed for our pendulum mountings is ε =0.05. Simple approximate analysis leads to a dynamic modulation of the cavity length l by the acceleration a, as

$$\frac{\Delta l}{l}\Big|_{axial} = \frac{al\varepsilon}{2Y}, \tag{6}$$

where Y ~70 GPa is the Young's modulus for the ULE or Zerodur spacer. For ε = 1, l = 10 cm and a = 1 g, we expect Δl/l = - Δf/f ~ 7.5 MHz/g, supposing λ =633 nm and not yet counting the nearly symmetric mounting. Inclusion of this factor makes our horizontal sensitivity 375 kHz/g. Vertically accelerating the interferometer produces length changes through the distortion coupling between the lateral and lengthwise dimension, the effect of "extrusion of the toothpaste", with a displacement reduction by the Poisson ratio σ = 0.17. Also the vertical height is really the spacer's diameter ϕ, which is about 5-fold less than the length.

$$\frac{\Delta l}{l}\Big|_{vertical} = \frac{a\phi\sigma}{2Y}. \tag{7}$$

We come to a vertical sensitivity of 250 kHz/g. Integrating the acceleration produced by the mid-band vibration spectrum quoted above leads to a broadband noise of a few Hz in both H and V planes. Left out however is the 1 milli-"g" vibration near 30 Hz due to ac motors in JILA (Pepsi refrigerators!). So we should have a vibration-induced linewidth of something like 1/2 kHz, which correlates well with experience. Active anti-vibration measures suppress this linewidth below 10 Hz, while improved passive mountings at NIST have recently led to sub-Hertz cavity-locked laser linewidths.[39]

Quantum resonance absorption[40]

Establishing a long-term stable optical frequency standard requires a natural reference of atomic or molecular origin. Historically the use of atomic/molecular transitions was limited to those that had accidental overlap with some fixed laser wavelengths. With the advent of tunable lasers, research on quantum absorbers has flourished. A stabilized laser achieves fractional frequency stability $\frac{\delta\nu}{\nu} = \frac{1}{Q}\frac{1}{S/N}\frac{1}{\sqrt{\tau}}$, where Q is the quality factor of the transition involved, S/N is the recovered signal-to-noise ratio of the resonance information, and τ is the averaging time. Clearly one wishes to explore the limits on both resolution and sensitivity of the detected signal. The nonlinear nature of a quantum absorber, while on one hand limiting the attainable S/N, permits sub-Doppler resolutions. With sensitive techniques such as FM-based signal modulation and recovery, one is able to split a MHz scale linewidth by a factor of 10^4 – 10^5, at an averaging time of 1 s or so. Sub-Hertz long-term stability can be achieved with carefully designed optical systems where residual effects on baseline stability are minimized. However, a pressing question is: How accurate is our knowledge of the center of the resonance? Collisions, electromagnetic fringe fields, probe field wave-front curvature, and probe power can all bring undesired linewidth broadening and center shifts. Distortion in the modulation process and other physical interactions can produce asymmetry in the recovered signal line shape. These issues will have to be addressed carefully before one can be comfortable talking about accuracy. A more fundamental issue related to time dilation of the reference system (second order Doppler effect) can be solved in a controlled fashion: one simply knows the sample velocity accurately (for example, by velocity selective Raman process), or the velocity is brought down to a negligible level using cooling and trapping techniques.

The simultaneous use of quantum absorbers and an optical cavity offers an attractive laser stabilization system. On one hand, a laser pre-stabilized by a cavity offers a long phase coherence time, reducing the need of frequent interrogations of the quantum absorber. In other words, information of the atomic transition can be recovered with an enhanced S/N and the long averaging time translates into a finer examination of the true line center. On the other hand, the quantum absorber's resonance basically eliminates inevitable drifts associated with material standards. The frequency offset between the cavity mode and atomic resonance can be bridged by an AOM. In this case the cavity can be made of totally passive elements: mirrors are optically contacted to a spacer made of ultra low expansion material such as ULE or Zerodur. In case that the cavity needs to be made somewhat tunable, an intracavity Brewster plate driven by Galvo or a mirror mounted on PZT are often employed. Of course these mechanical parts bring additional thermal and vibrational sensitivities to the cavity, along with non-linearity and hysteresis. Temperature tuning of a resonator is potentially less noisy but slow. Other tuning techniques also exist, for example, through the use of magnetic force or pressure (change of intracavity refractive index or change of cavity dimension by external pressure). An often-used powerful technique called frequency-offset-locking brings the precision rf tuning capability to the optical world.[14]

Transducers

PZT Transducer design: disk vs tube designs

We will usually encounter the mechanical resonance problem in any servo based on a PZT transducer: Small mirrors clearly are nice as they can have higher resonance frequencies. A mirror, say 7.75 mm Φ x 4 mm high, might be waxed onto a PZT disk 10 mm diameter x 0.5 mm thick. The PZT, in turn, is epoxied onto a serious backing plate. This needs to be massive and stiff, since the PZT element will produce a differential force between the mirror and the backing plate. At short times there will be a "reduced mass" kind of splitting of the motion between the mirror and the support plate. At lower frequencies, one hates to get a lot of energy coupled into the mirror mount since it will have a wealth of resonances in the sub-kHz range. For this size mirror, the backing plate might be stainless steel, 1 inch diameter by ~3/4 inch wedged thickness, and with the PZT deliberately decentered to break down high Q modes. The piston mode will be at ~75 kHz.

tubular PZT design

Often it is convenient to use a tubular form of PZT, with the electric field radially-applied across a thin wall of thickness t. This gives length expansion also, transverse to the field using a weaker d_{31} coefficient, but wins a big geometric factor in that the transverse field is generating a length response along the entire tube height h. The PZT tube could be 1/2" Φ x 1/2 high, with a wall thickness t =1.25 mm. This geometry leads to a ~ 7-fold sensitivity win, when $d_{31} \approx 0.7\, d_{33}$ is included. Typical dimensions for the mirror might be 12.5 mm diameter x 7 mm high. The PZT tube also is epoxied onto a serious backing plate. For the high voltage isolation of the PZT electrodes at the tube ends, a thin sheet (say < 0.5 mm) of stiff ceramic, alumina for example, will suffice. An alternative way to provide the electrical isolation of the ends involves removing the silver electrodes for several mm at the end. A new technique uses a diamond-charged tubular core drill mounted into a collet in a lathe. The active tool face projects out only 2 mm so that hand-held PZT grinding leads to clean electrode removal, inside and out. This end of the PZT tube is attached to the backing mass with strong epoxy. The mirror is attached to the open PZT tube end with melted wax. This is vastly better than epoxy in that it does not warp the optic, and the small energy dissipation occurs at the best place to damp the Q of the PZT assembly. If done well, this unit will have its first longitudinal resonance at about 25 kHz, with a Q ~10. As noted above, in servo terms, the actual mechanical PZT unit gives an added 2-pole roll-off above the resonance frequency and a corresponding asymptotic phase lag of 180°. So it is useful to design for high resonant frequency and low Q.

Comparing disk and tubular designs, the disk approach can have a three-fold higher resonance frequency, while the tubular design is ~7-fold more sensitive. Perhaps more important is the tube's reduced stiffness, moving the PZT/mirror resonance down into the 20 kHz domain. This brings us to the subject of spectral shaping of the amplifier gain and limitations of servo performance due to electronic issues.

Amplifier strategies for PZT driver

We enjoy the tubular PZT for its large response per volt and its relatively high resonance frequencies. But it gives a problem in having a large capacitance, for example of 10 nf in the above design. Even with the high sensitivity of 70 V/order, achieving a tight lock requires high frequency corrections and can lead to a problem in supplying the necessary ac current, supposing that we ask the HV amplifier alone to do the job. An apparent answer is to use a pair of amplifiers, one fast and the other HV, separately driving the two sides of the PZT. This alone doesn't solve the problem, as the big high-frequency ac current is only returned via the HV amplifier. The answer is to use a crossover network on the HV amplifier side. A capacitor to ground, of perhaps 3 or 5-fold larger value than the PZT will adequately dump the fast currents coming through the PZT's capacitance. A resistor to this PZT/shunt capacitor junction can go to the HV amplifier. Now this HV amp has indeed more capacitance to drive, but is only needed to be active below a few hundred Hz where the current demand becomes reasonable. An alternative topology sums the two inputs on one side of the PZT.

Other useful transducers - slow but powerful

Commercial multiple wafer designs utilize 100 or more thin PZT sheets mechanically in series and electrically in parallel to produce huge excursions such as 10 μm for 100 V. Of course the capacitance is ~0.1 μF and the stability leaves something to be desired. These are useful for applications that can tolerate some hysteresis and drift, such as grating angle tuning in a diode laser. When a large dynamic range is needed to accommodate wide tuning range or to correct for extensive laser frequency drifts at low frequencies, a galvo-driven Brewster plate can be used inside the optical cavity. Typically a Brewster plate inflicts an insertion loss less than 0.1% if its angular tuning range is limited within ±4 degrees. Walk-off of the optical beam by the tuning plate can be compensated with a double-passing arrangement or using dual plates. In the JILA-designed Ti:Sapphire laser, we use the combination of PZT and Brewster plate for the long term frequency stabilization. The correction signal applied to the laser PZT is integrated and then fed to the Brewster plate to prevent saturation of the PZT channel. At higher frequencies we use much faster transducers, such as AOM and EOM, which are discussed below.

Temperature control of course offers the most universal means to control long-term drifts. Unfortunately the time constant associated with thermal diffusion is usually slow and therefore the loop bandwidth of thermal control is mostly limited to Hz scale. However, thoughtful designs can sometimes push this limit to a much higher value. For example, a Kapton thin-film heater tape wrapped around the HeNe plasma tube has produced a thermal control unity gain bandwidth in excess of 100 Hz.[41] The transducer response is reasonably modeled as an integrator above 0.3 Hz and excessive phase shifts associated with the thermal diffusion does not become a serious issue until ~ 200 Hz. This transfer function of the transducer can be easily compensated with an electronic PI filter to produce the desired servo loop response. Radiant heating of a glass tube by incandescent lamps has achieved a time delay < 30 ms and has also been used successfully for frequency control of HeNe lasers.[42] If a bipolar thermal control is needed, Peltier-based solid state heat pumps (thermoelectric coolers) are available and can achieve temperature differences up to 70 ^{0}C, or can transfer heat at a rate of 125 W, given a proper configuration of heat sinking. Parallel use of these Peltier devices result in a greater amount of heat transfer while cascaded configuration achieves a larger temperature difference.

Combining various servo transducers in a single feedback loop requires though understanding of each actuator, their gains and phase shifts, and the overall loop filter function one intends to construct. Clearly, to have an attractive servo response in the time-domain, the frequency transfer functions of various gain elements need to crossover each other smoothly. A slow actuator may have some resonance features in some low frequency domain, hence the servo action needs to be relegated to a faster transducer at frequency ranges beyond those resonances. The roll-off of the slow transducer gain at high frequencies needs to be steep enough, so that the overall loop gain can be raised without exciting the associated resonance. On the other hand, the high frequency channel typically does not have as large a dynamic range as the slow ones. So one has to pay attention not to overload the fast channel. Again, a steep filter slope is needed to rapidly relinquish the gain of the fast channel towards the low frequency range. However, we stress here that the phase difference between the two channels at the crossover point needs to be maintained at less than 90^0. In the end, pre-determined gains and phase shifts will be assigned to each transducer so that the combined filter function resembles a smooth single channel design. Some of these issues will be addressed briefly in the section below on example designs.

Servo Design in the face of time delay: additional transducers are useful

As one wishes for higher servo gain, with stability, it means a higher closed-loop bandwidth must be employed. Eventually the gain is sufficiently large that the intrinsic laser noise, divided by this gain, has become less than the measurement noise involved in obtaining the servo error signal. This should be sufficient gain. However it may not be usable in a closed-loop scenario, due to excessive time delay. If we have a time delay of t_{delay} around the loop from an injection to the first receipt of correction information, a consideration of the input and response as vectors will make it clear that no real servo noise suppression can occur unless the phase of the response at least approximates that required to subtract from the injected error input to reduce its magnitude on the next cycle through the system. A radian of phase error would correspond to a unity-gain frequency of $1/2\pi t_{delay}$, and we find this to be basically the upper useful limit of servo bandwidth. One finds that to correct a diode laser or dye laser to leave residual phase errors of 0.1 rad, it takes about 2 MHz servo bandwidth. This means a loop delay time, at the absolute maximum, of $t_{delay} = 1/2\pi$ 2MHz = 80 ns. Since several amplifier stages will be in this rf and servo-domain control amplifier chain, the individual bandwidths need to be substantially beyond the 12 MHz naively implied by the delay spec. In particular rf modulation frequencies need to be unexpectedly large, 20 MHz at least, and octave rf bandwidths need to be utilized, considering that the modulation content can only be 1/2 the bandwidth. Suppression of even-order signals before detection is done with narrow resonant rf notches.

Of course a PZT transducer will not be rated in the ns regime of time delay. Rather, one can employ an **AOM** driven by a fast-acting Voltage-Controlled Oscillator to provide a frequency shift. Unfortunately the acoustic time delay from the ultrasonic transducer to the optical interaction seems always to be 400 ns, and more if we are dealing with a very intense laser beam and wish to avoid damage to the delicate AO transducer. The AOM approach works well with diode-pumped solid state lasers, where the bandwidth of major perturbations might be only 20 kHz. By double-passing the AOM the intrinsic angular deflection is suppressed. Usually the AOM prefers linear polarization. To aid separation of the return beam on the input side, a spatial offset can be provided with a collimating lens and roof prism, or with a cat's-eye retro-reflector. Amplitude modulation or leveling can also be provided with the AOM's dependence on rf drive, but it is difficult to produce a beam still at the shot noise level after the AOM.

The final solution is an **EOM phase modulator**. In the external beam, this device will produce a phase shift per volt, rather than a frequency shift. So we will need to integrate the control input to generate a rate of change signal to provide to the EOM, in order to have a frequency relationship with the control input.[43] Evidently this will bring the dual problems of voltage saturation when the output becomes too large, and a related problem, the difficulty of combining fast low-delay response with high voltage capability. The standard answer to this dilemma was indicated in our PZT section, namely, one applies fast signals and high-voltage signals independently, taking advantage of the fact that the needed control effort at high frequencies tends to cover only a small range. So fast low voltage amplifier devices are completely adequate, particularly if one multi-passes the crystal several times. A full discussion of the crossover issues and driver circuits will be prepared for another publication.

Representative/Example Designs

Diode-Pumped Solid State Laser

Diode-pumped solid state lasers are viewed as the most promising coherent light sources in diverse applications, such as communications, remote detection and high precision spectroscopy. Nd:YAG laser is among the first developed diode-pumped solid state lasers and has enjoyed continuous improvements in its energy efficiency, size, lifetime and intrinsic noise levels. The laser's free-running linewidth of ~10 kHz makes it a straightforward task to stabilize the laser via an optical cavity or an optical phase locked loop. In our initial attempt to stabilize the laser on a high finesse (F ~ 100,000, linewidth ~ 3 kHz) cavity, we employ an external acousto-optic modulator (AOM) along with the laser internal PZT which is bonded directly on the laser crystal. The frequency discrimination signal between the laser and cavity is obtained with 4 MHz FM sidebands detected in cavity reflection. The PZT corrects any slow but potentially large laser frequency noise. Using the PZT alone allows the laser to be locked on cavity. However, the loop tends to oscillate around 15 kHz and the residual noise level is more than 100 times higher than that obtained with the help of an external AOM. The AOM is able to extend the servo bandwidth to ~ 150 kHz, limited by the propagation time delay of the acoustic wave inside the AOM crystal. The crossover frequency between the PZT and AOM is about 10 kHz. Such a system has allowed us to achieve a residual frequency noise spectral density of 20 mHz/√Hz. The laser's linewidth relative to cavity is thus a mere 1.3 mHz,[44] even though the noise spectral density is still 100 times higher than the shot noise. This same strategy of servo loop design has also been used to achieve a microradian level phase locking between two Nd:YAG lasers.[45]

It is also attractive to stabilize the laser directly on atomic/molecular transitions, given the low magnitude of the laser's intrinsic frequency noise. Of course the limited S/N of the recovered resonance information will not allow us to build speedy loops to clean off the laser's fast frequency/phase noise. Rather we will use the laser PZT alone to guide the laser for a long term stability. An example here is the 1.064 μm radiation from the Nd:YAG, which is easily frequency doubled to 532 nm where strong absorption features of iodine molecules exist.[46,47] The doubling is furnished with a noncritical phase-matched $KNbO_3$ crystal located inside a buildup cavity. 160 mW of green light output is obtained from an input power of 250 mW of IR. Only mW levels of the green light are needed to probe the iodine saturated absorption signal. Low vapor pressure (~ 0.5 Pa) of the iodine cell is used to minimize the collision-induced pressure shift and to reduce the influence on baseline by the linear Doppler absorption background. The signal size decreases as the pressure is reduced. However, this effect is partly offset due to the reduced resonance linewidth (less pressure broadening) which helps to increase the slope of the frequency locking error signal. A lower pressure also helps to reduce power-related center frequency shifts since a lower power is needed for saturation. With our 1.2-m long cells, we have achieved a S/N of 120 in a 10 kHz bandwidth, using the modulation transfer spectroscopy.[48] (Modulation transfer is similar to FM except that we impose the frequency sideband on the saturating beam and rely on the nonlinear medium to transfer the modulation information to the probe beam which is then detected.) Normalized to 1-s averaging time, this S/N translates to the possibility of a residual frequency noise level of 10 Hz when the laser is locked on the molecular resonance, given the transition linewidth of 300 kHz. We have built two such iodine-stabilized systems and the heterodyne beat between the two lasers permits systematic studies on each system and checks the reproducibility of the locking scheme.[49] With a 1-s counter gate time, we have recorded the beat frequency between the two lasers. The standard deviation of the beat frequency noise is ~ 20 Hz, corresponding to ~ 14 Hz rms noise per IR laser, basically a S/N limited performance. The beat record can be used to calculate the Allan standard deviation: starting at 5×10^{-14} at 1-s, decreasing with a slope of $1/\sqrt{\tau}$ up to 100-s. (τ is the averaging time.) After 100-s the deviations reach the flicker noise floor of ~ 5×10^{-15}. At present, the accuracy of the system is limited by inadequate optical isolation in the spectrometer and the imperfect frequency modulation process (residual amplitude noise, RAM) used to recover the signal. This subject is under intense active study in our group.

External Cavity Diode Lasers

Diode lasers are compact, reliable and coherent light sources for many different applications.[50] The linewidth of a free-running diode laser is limited by the fundamental spontaneous emission events, enhanced by the amplitude-phase coupling inside the gain medium. With a low noise current driver, a typical mW scale AlGaAs diode laser has a linewidth of several MHz. To reduce this fast frequency noise, one typically employs an external cavity formed between one of the diode laser facets and a grating (or an external mirror that retro-reflects the first order grating diffraction).[51,52,53] This optical feedback mechanism suppresses the spontaneous emission noise, replaced by much slower fluctuations of mechanical origin. The linewidth of the grating-stabilized external cavity diode laser (ECDL) is usually between 100 kHz and 1 MHz, determined by the quality factor of the optical feedback. The ECDL also offers much better tuning characteristics compared against a solitary diode. To do such tuning, the external grating (or the mirror that feeds the grating-dispersed light back to the laser) is controlled by a PZT for scanning. Synchronous tuning of the grating dispersion and the external cavity mode can be achieved with a careful selection of the grating rotation axis position. Similarly, this PZT controlled grating can be used to stabilize the frequency of an ECDL. However, owing to the low bandwidth limited by the mechanical resonance of PZT, a tight frequency servo is possible only through fast transducers such as the laser current or intracavity phase modulators.

This hybrid electro-optic feedback system is attractive, and ECDLs have been demonstrated to show Hz level stability under a servo bandwidth of the order of 1 MHz. For a solitary diode, feedback bandwidth of tens of MHz would have been needed in order to bring the frequency noise down to the same level. However, considering that the optical feedback has a strong impact on the laser frequency noise spectrum, one finds the frequency response of the compound laser system is clearly dependent upon the optical alignment. Therefore for each particular ECDL system, we need to measure the frequency response function of the laser under the optimally aligned condition. We are dealing with a multi-channel feedback system (for example, PZT plus current), so that designing smooth crossovers between different transducers requires knowledge of the transfer functions of each transducer. Normally the current-induced FM of a solitary diode has a flat response up to 100 kHz, and then starts to roll off in the region between 100 kHz and 1 MHz, initially with a single-pole character. This is due to the time response of the current-induced thermal change of the refractive index inside the diode. (At a faster time scale, the carrier density variation will dominate the laser frequency response.) Design of a fast feedback loop needs to take into account of this intrinsic diode response. Fortunately the time delay associated with the current response is low, typically below 10 ns.

In our example system, the frequency discrimination signal of the ECDL is obtained from a 100 kHz linewidth cavity with a sampling frequency of 25 MHz. The error signal is divided into three paths: PZT, current modulation through the driver, and direct current feedback to the diode head. The composite loop filter function is shown in Fig. 8. The crossover between the slow current channel and the PZT usually occurs around 1 kHz, in order to avoid the mechanical resonance of PZT at a few kHz. In our system, the frequency response of the PZT/grating is 10 GHz/V. To furnish this in-loop gain of ~ 1000 at 1 kHz, we need to supply an electronic gain of 0.1, given that the error signal has a slope of 1 V/1 MHz. Towards the lower frequency range the PZT gain increases by 40 dB/decade (double integrators) to suppress the catastrophically rising laser frequency noise. It is obvious from Fig. 8 that the intermediate current channel tends to become unstable at a few hundred kHz, due to the excessive phase shift there. The fast current loop, bypassing the current driver to minimize additional time delay and phase shift, has a phase lead compensator to push the unity gain bandwidth to 2 MHz. With this system we can lock the ECDL robustly on the optical cavity, with a residual noise spectral density of 2 Hz/√Hz, leading to a relative linewidth of 12 Hz. The achieved noise level is about 100 times higher than the fundamental measurement limit set by shot noise. We note in passing that when an ECDL gradually goes out of alignment, the previously adjusted gain of the current loop will tend to make the servo oscillate so a new alignment is needed. The laser FM sideband used to generate the locking signal is produced directly by current modulation. An electronic filter network is employed to superimpose the slow servo, fast servo and modulation inputs to the diode. Exercise caution when accessing the diode head, as a few extra mA current increase can lead to drastic output power increase and melted laser facets, all in 1 μs!

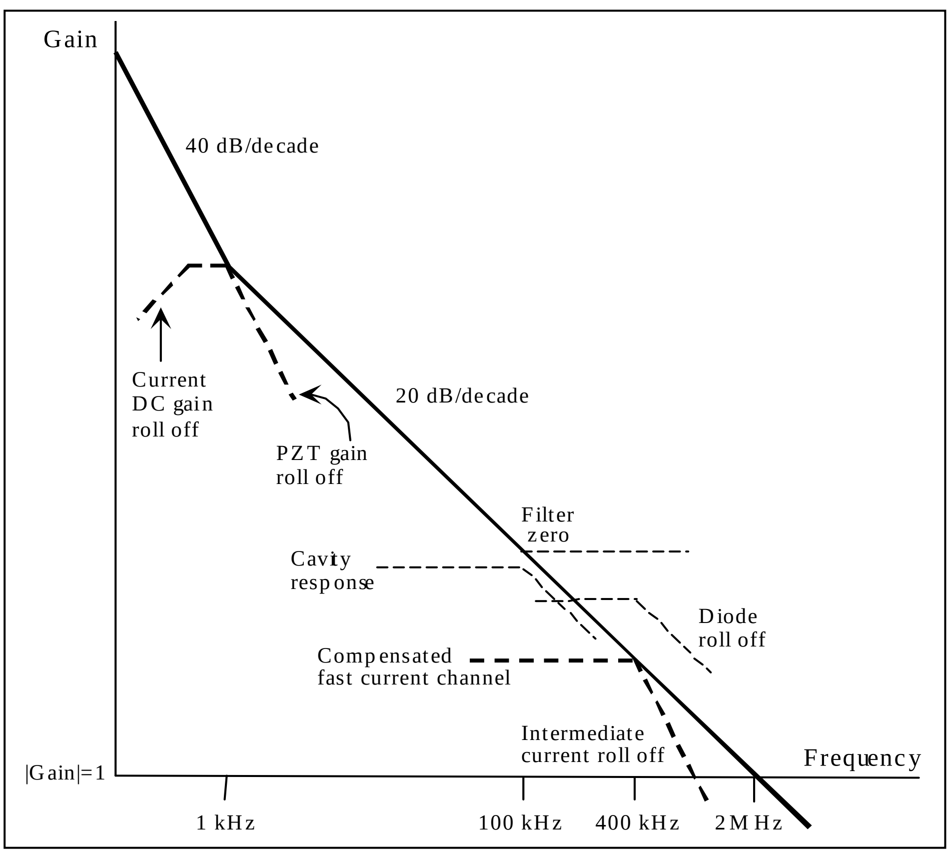

Fig.8 The combined loop filter function for ECDL frequency stabilization.

Summary and Outlook

The technology of laser frequency stabilization has been refined and simplified over the years and has become an indispensable research tool in any modern laboratories involving optics. Research on laser stabilization has been and still is pushing the limits of measurement science. Indeed, a number of currently active research projects on fundamental physical principles benefit a great deal from stable optical sources and will need a continued progress of the laser stabilization front.[54] Using extremely stable phase coherent optical sources, we will be entering an exciting era when picometer resolution can be achieved over a million kilometer distance in the space[55] or a few Hertz linewidth of an ultraviolet resonance can be probed with a high S/N.[56] One has to be optimistic looking at the stabilization results of all different kinds of lasers. To list just a few examples of cw tunable lasers, we notice milliHertz linewidth stabilization (relative to a cavity) for diode-pumped solid state lasers; dozen milliHertz linewidth for Ti:Sapphire lasers; and sub-Hertz linewidths for diode and dye lasers. Long-term stability of lasers referenced to atoms and molecules have reached mid 10^{-15} level in a short averaging time of ~ 300 s. Phase locking between different laser systems can be achieved, even for diode lasers that have fast frequency noise.

Quantum noise is the usual limit of the measurement process and therefore will be the limit of the stabilization process as well. To circumvent the quantum noise is altogether an active research field itself.[57] We, however, have not reached this quantum limit just yet. For instance, we have already stated that the Nd:YAG laser should be able to reach microHertz stability if the shot noise is the true limit. What have we done wrong? A main part of the deficiency is due to the inadequacy of the measurement process, namely the lack of accuracy. This is because the signal recovery effort — modulation and demodulation process is contaminated by spurious optical interference effects and RAM associated with the modulation frequency. Every optical surface along the beam path can be a potential time bomb to damage the modulation performance. In cases that some low contrast interference effects are not totally avoidable, we would need to have the whole system controlled in terms of the surrounding pressure and temperature. The degree to which we can exert control of course dictates the ultimate performance.

Conclusions and Recommendations

It becomes clear that there are many interlinking considerations involved in the design of laser stabilization systems, and it is difficult to present a full description in an article such as this. Still it is hoped that the reader will see some avenues to employ feedback control methods to the laser systems of her current interest. We are optimistic that some of this technology may become commercially available in the future, thus simplifying the user's task.

Acknowledgements

The work discussed here has profited from interactions with many colleagues, postdoctoral researchers and graduate students over many years. In particular we must thank Leo Hollberg and Miao Zhu for their earlier contributions. One of us (JLH) is particularly grateful to his wife Lindy for patience 'beyond the call of duty' during these many years of laser research. The work at JILA has been supported over the years by the Office of Naval Research, the National Science Foundation, the Air Force Office of Scientific Research, and the National Institute of Standards and Technology, as part of its frontier research into basic standards and their applications.

References

1. D. W. Allan, Proc. IEEE **54**, 221 (1966).
2. “Characterization of Clocks and Oscillators”,
   D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, Editors,
   NIST Technical Note **1337**, U. S. Government Printing Office, Washington DC, 1990.
3. D. S. Elliot, R. Roy and S. J. Smith, Phys. Rev. A **26**, 12 (1982).
4. J. L. Hall and M. Zhu, in Proc. of *the International School of Phys. ‘Enrico Fermi’*, Course CXVIII, *Laser manipulation of atoms and ions*, 671, E. Arimondo, W. D. Phillips and F. Strumia, Eds., North Holland, 671 (1992).
5. M. Zhu and J. L. Hall, J. Opt. Soc. Am. B **10**, 802 (1993).
6. For similar issues in the microwave/rf field, see the application note
   “Time keeping and Frequency Calibration”,
   Hewlett Packard Company, Palo Alto, CA.
7. For general references on feedback systems, please refer to
   “Modern Control Systems”, 3-rd edition
   Richard C. Dorf, Addison-Wesley Publishing Co., Reading MA (1980).
8. “Feedback Control of Dynamic Systems”, 3rd edition
   Gene F. Franklin, J. David Powell and A. Emami-Naeini
   Addison-Wesley Publishing Co., Reading MA (1994).
9. The simulations presented here are performed with the following software,
   Matlab Control Systems Toolbox: User's Guide
   Matlab Simulink: User's Guide
   The Math Works Inc., Nattick MA (1996).
10. See, e.g., J. L. Hall, in *Frequency-Stabilized Lasers and Their Applications*, Y. C. Chung, ed., SPIE Proc. **1837**, 2 (1993).
11. R. L. Barger, M. S. Sorem, and J. L. Hall, Appl. Phys. Lett. **22**, 573 (1973).
12. G. D. Houser and E. Garmire, Appl. Opt. **33**, 1059 (1994).
13. K. L. Haller and P. C. D. Hobbs, in *Optical Methods for Unltrasensitive Detection and Analysis: Techniques and Applications*, B. L. Fearey, ed., Proc. SPIE **1435**, 298 (1991).
14. J. Helmcke, S. A. Lee, and J. L. Hall, Appl. Opt. **21**, 1686 (1982).
15. C. E. Wieman and T. W. Hänsch, Phys. Rev. Lett. **36**, 1170 (1976).
16. T. W. Hänsch and B. Couillaud, Opt. Comm. **35**, 441 (1980).

17. H. Wahlquist, J. Chem. Phys. **35**, 1708 (1961).
18. T. J. Quinn, Metrologia **30**, 524 (1994).
19. J. Hu, E. Ikonen, and K. Riski, Opt. Comm. **120**, 65 (1995).
20. M. S. Taubman and J. L. Hall, Opt. Lett., in press (2000).
21. B. Smaller, Phys. Rev. **83**, 812-820 (1951); and R. V. Pound, Rev. Sci. Instrum. **17**, 490 (1946).
22. G. C. Bjorklund, Opt. Lett. **5**, 15 (1980).
23. J. L. Hall, L. Hollberg, T. Baer, and H. G. Robinson, Appl. Phys. Lett. **39**, 680-682 (1981).
24. G. C. Bjorklund and M. D. Levenson, Phys. Rev. A **24**, 166-169 (1981).
25. W. Zapka, M. D. Levenson, F. M. Schellenberg, A. C. Tam, and G. C. Bjorklund, Opt. Lett. **8**, 27 (1983).
26. G. J. Rosasco and W. S. Hurst, J. Opt. Soc. Am. **B2**, 1485-1496 (1985).
27. J. H. Shirley, Opt. Lett. **7**, 537 (1982).
28. R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, Appl. Phys. B **31**, 97 (1983).
29. P. Werle, Appl. Phys. B **60**, 499 (1995).
30. J. L. Hall, J. Ye, L.-S. Ma, K. Vogel, and T. Dinneen, in *Laser Spectroscopy XIII*, Z.-J. Wang, Z.-M. Zhang, and Y.-Z. Wang, eds. (World Scientific, Singapore, 1998) pp. 75-80.
31. N. C. Wong and J. L. Hall, J. Opt. Soc. Am. B **2**, 1527 (1985).
32. M. S. Taubman and J. L. Hall, work in progress (1999).
33. D. R. Hjelme, S. Neegård, and E. Vartdal, Opt. Lett. **20**, 1731 (1995).
34. G. R. Janik, C. B. Carlisle, and T. F. Gallagher, J. Opt. Soc. Am. **B 3**, 1070 (1986).
35. J. L. Hall, L.-S. Ma, and G. Kramer, IEEE J. Quan. Electro. **QE-23**, 427 (1987).
36. M. Kourogi, K. Nakagawa, and M. Ohtsu, IEEE J. Quan. Electron. **QE-29**, 2693 (1993).
37. Th. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch, Opt. Lett. **24**, 881 (1999).
38. S. A. Diddams, L.-S. Ma, J. Ye, and J. L. Hall, Opt. Lett. **24**, 1747 (1999).
39. B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist, Phys. Rev. Lett. **82**, 3799 (1999).
40. Interested readers are referred to the following book for more detailed discussions. "The Quantum Physics of Atomic Frequency Standards", volumes I and II, Jacques Vanier and Claude Audoin, Adam Hilger, Institute of Physics Publishing Ltd, Bristol, UK (1989).
41. T. M. Niebauer, J. E. Faller, H. M. Godwin, J. L. Hall and R. L. Barger, Appl. Opt. **27**, 1285 (1988).
42. Jun Ishikawa, private communications (1996).
43. J. L. Hall and T. W. Hänsch, Opt. Lett. **9**, 502 (1984).
44. J. Ye, L.-S. Ma, and J. L. Hall, J. Opt. Soc. Am. B **15**, 6 (1998).
45. J. Ye and J. L. Hall, Opt. Lett. **24**, 1838 (1999).

46. A. Arie and R. L. Byer, J. Opt. Soc. Am. B **10**, 1990 (1993).
47. M. L. Eickhoff and J. L. Hall, IEEE Trans. Instrum. & Meas. **44**, 155 (1995); and P. Jungner, M. Eickhoff, S. Swartz, J. Ye, J. L. Hall and S. Waltman, IEEE Trans. Instrum. & Meas. **44**, 151 (1995).
48. J. Ye, L. Robertsson, S. Picard, L.-S. Ma, and John L. Hall, IEEE Trans. Instrum. & Meas. **48**, 544 (1999).
49. J. L. Hall, L.-S. Ma, M. Taubman, B. Tiemann, F.-L. Hong, O. Pfister, and J. Ye, IEEE Trans. Instrum. & Meas. **48**, 583, (1999).
50. C. E. Wieman and L. Hollberg, Rev. Sci. Instrum. **62**, 1 (1991).
51. M. G. Littman and H. J. Metcalf, Appl. Opt. **17**, 2224 (1978).
52. B. Dahmani, L. Hollberg, and R. Drullinger, Opt. Lett. **12**, 876 (1987).
53. K. MacAdam, A. Steinbach, and C. E. Wieman, Am. J. Phys. **60**, 1098 (1992).
54. P. Fritschel, G. González, B. Lantz, P. Saha, and M. Zucker, Phys. Rev. Lett. **80**, 3181 (1998).
55. K. Danzmann, Class. Quantum Grav. **14**, 1399 (1997).
56. J. C. Bergquist, W. M. Itano, and D. J. Wineland, in *Frontiers in Laser Spectroscopy*, 359 (1994); J. C. Bergquist, private communications (1999).
57. "Quantum Measurement", V. B. Braginsky and F. Ya. Khalili, Cambridge University Press, Cambridge, UK (1992).