Characterizing Phase Noise in a Phase-Locked Loop with a PNA1 Intensity Noise Analyzer

Application Note

Table of Contents

Introduction

Many techniques in photonics require locking two different oscillating sources to be in phase with each other by creating a signal that is sensitive to the phase difference between the two sources and feeding back on one of the sources to "lock" the two together. Usually, one of the oscillators is the reference and is assumed to be very low noise, and the other oscillator is locked to this reference. This phase-locked loop (PLL) cannot perfectly correct for the phase deviations between the two sources, so when optimizing the lock, it is crucial to be able to measure and analyze the phase error (phase noise) between the two sources. This application note will describe how to characterize phase noise by measuring a PLL error signal, as well as how to use Thorlabs' PNA1 Intensity Noise Analyzer to determine the total integrated phase error between the two oscillators.

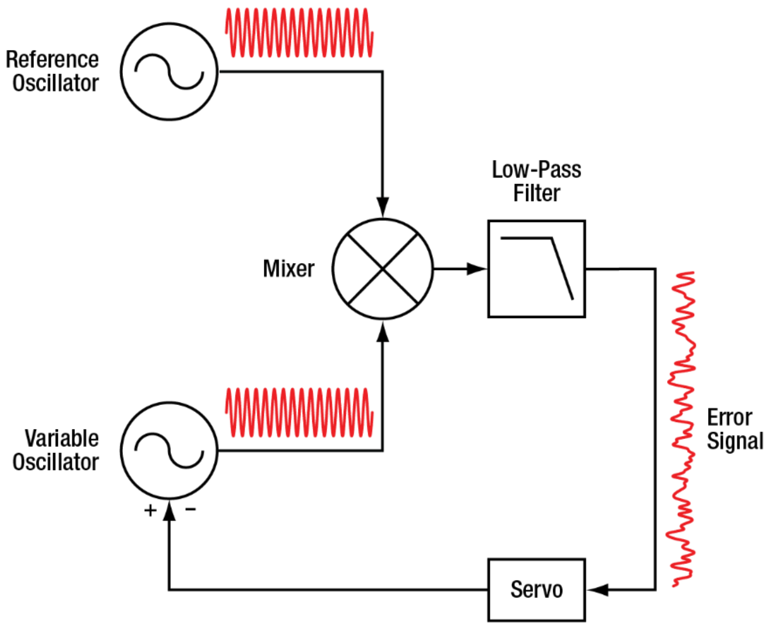

Figure 1. **General Phase-Locked Loop** A phase-locked loop consists of a reference oscillator, a variable oscillator to be phase-locked to the reference, and control electronics. The signals from the two oscillators are combined by a mixer, and the combined signal passes through a low-pass filter (LPF), producing an error signal. The error signal is fed into a servo, which provides a continuous feedback signal to the variable oscillator. For characterization of the lock quality and optimization of the servo parameters, the error signal can also be tapped off and analyzed using an apparatus such as Thorlabs' PNA1 Intensity Noise Analyzer, which is the focus of this Application Note.

Estimating Phase Error from Typical Error Signals

In a typical phase-locking application, such as locking two RF frequencies or two laser sources together, the magnitude of the error signal will be a sinusoidal function of the phase error φ between the two oscillators, perhaps with a DC offset (see the Phase Locking section for details):

$$V_{error\ signal} = A\sin(\varphi) + V_0$$

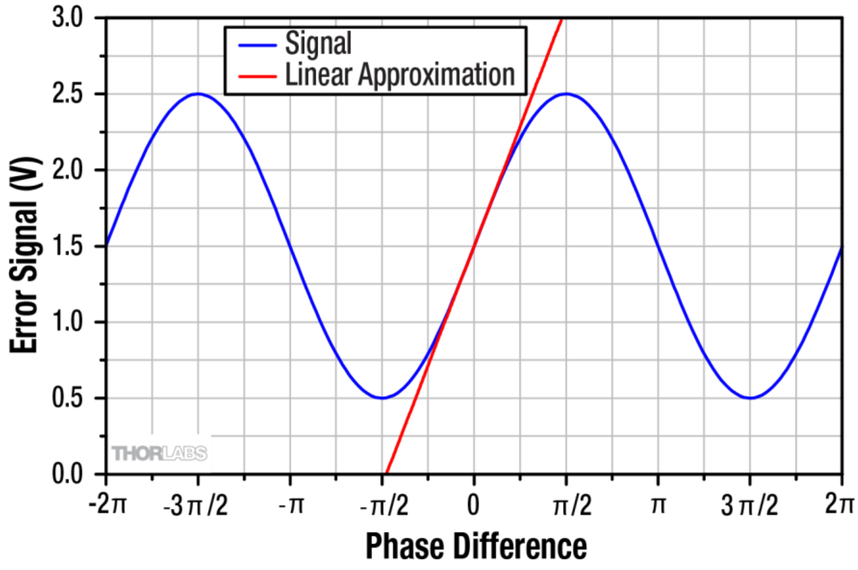

Figure 2. Typical Error Signal This graph shows the error signal for a typical phase-locked loop as a function of the phase difference between the oscillators. The sinusoidal relationship (blue) is well-approximated as linear (red) when the phase difference is small, but the approximation fails for larger phase differences. Note that as the dependence is periodic, the phase difference is only significant up to a factor of 2π.

The amplitude (A) and voltage offset (V_0) can be determined for a particular setup by varying φ and measuring the maximum and minimum voltages V_{max} and V_{min}:

$$A = \frac{(V_{max} - V_{min})}{2}$$

$$V_0 = \frac{(V_{max} + V_{min})}{2}$$

To lock the two sources together, some form of actuation is used to continually drive the error signal to the voltage offset. In the region of zero phase difference, which occurs when the two sources are nearly phase-locked, the phase error is well-approximated by:

$$\varphi \approx \frac{(V_{meas} - V_0)}{A}$$

where V_{meas} is the measured voltage. In Figure 2, $A = 1.0$ V and $V_0 = 1.5$ V, so a measured voltage of 1.6 V is a 0.1 V amplitude error, approximately corresponding to a phase error of 0.1 rad.

Phase Locking

A diagram of a phase-locked loop is shown below, where a voltage-controlled oscillator (VCO) at frequency f_{VCO} is locked to a reference radio frequency (RF) oscillator at frequency f_{ref}, using servo feedback based on the measured phase difference between the two oscillators.

The signals from the two oscillators are fed into a mixer, which produces two new frequencies: $f_{ref} - f_{VCO} = \Delta f$ and $f_{ref} + f_{VCO}$. The higher frequency $f_{ref} + f_{VCO}$, as well as any remaining signal at f_{ref} or f_{VCO}, are filtered out by a low-pass filter. If the two signals are at different frequencies, the phase difference will be a linear function of time, proportional to Δf. Owing to the sinusoidal voltage-phase difference relationship described above, the error signal will oscillate sinusoidally at Δf. If the oscillators are at the same frequency, the error signal reflects the instantaneous phase difference, which may be due to fluctuations in either source.

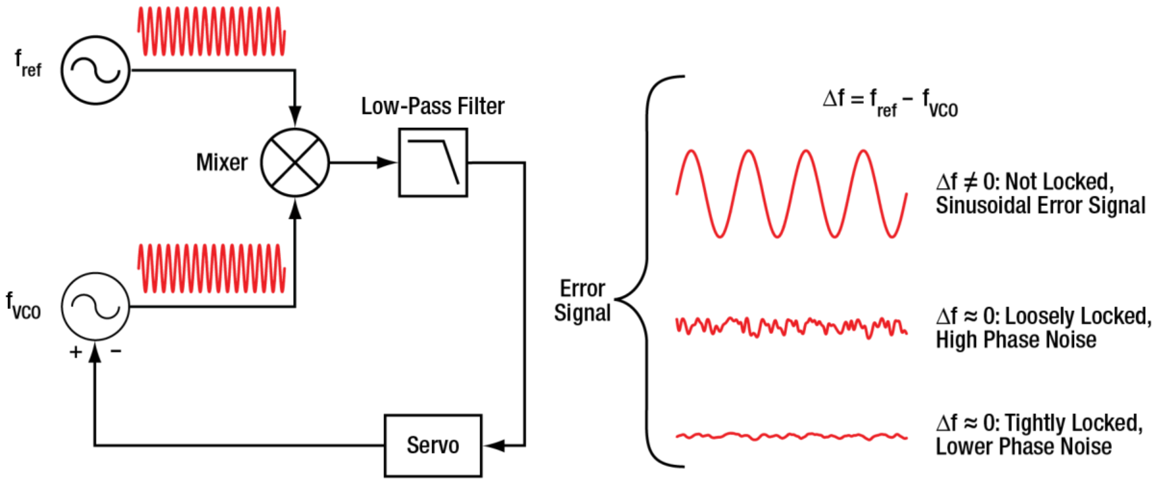

Figure 3. **RF Phase-Locked Loop** The left side of this schematic shows a typical RF phase-locked loop. The two oscillator signals are combined by the mixer and pass through the low-pass filer (LPF), producing the error signal. Three common types of error signals and the conditions that produce them are shown to the right. The error signal is fed into the servo, which adjusts the VCO to compensate for the measured phase error.

In Figure 3 three examples of error signals are shown:

1) An error signal from an unlocked loop, where there is a difference in frequency between the reference and VCO. Under this condition, the phase difference increases with time. As a result, the error signal automatically sweeps through the maximum and minimum amplitudes and enables the user to determine the error signal amplitude and offset voltage values described in the previous section.

2) An error signal from a loop where the VCO is being locked to the reference, but there is a large amount of residual fluctuation in the phase difference between the two signals. This is often seen when the two signals are first locked, but the servo PID settings have not yet been optimized.

3) An error signal from a loop where the VCO is more tightly locked to the reference source's phase, making the residual fluctuations smaller. This is typically seen after the servo settings have been optimized.

Measuring Residual Phase Noise (In-Loop)

To measure the quality of the phase lock, or to optimize the PID settings on the servo, part of the error signal can be split off and analyzed on a device such as an oscilloscope, an optical spectrum analyzer, or Thorlabs' PNA1 intensity noise analyzer.

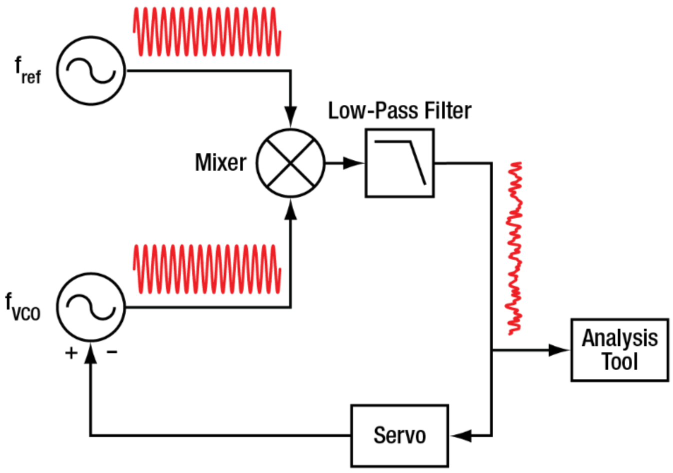

Figure 4. Measuring Residual Phase Noise In order to analyze residual phase noise after phase-locking, the error signal can be split off and sent to an analysis tool such as an oscilloscope, an optical spectrum analyzer (OSA), or the PNA1 intensity noise analyzer.

An oscilloscope easily visualizes the voltage deviations in the time domain, and one can optimize the quality of the lock by simply adjusting the controller parameters to minimize the envelope of the deviations. While this gives a good qualitative analysis, frequency-domain tools like spectrum analyzers are commonly used to get a more complete and quantitative picture of the error signal deviations from the lock point.

Spectrum analyzer devices generally measure Power Spectral Density (PSD) of an input voltage signal in units of $\frac{{V_{RMS}}^2}{Hz}$ or $\frac{dBm}{Hz}$ where dBm is the power of the signal into a 50 Ω load expressed in dB milliwatts[a]. These units can be post-processed into phase deviations via the phase-voltage relationship described above.

[a] $Power\ in\ dBm = 10log_{10}\left(\frac{Power\ in\ mW}{1\ mW}\right)$

Phase-Locking Two Continuous Wave Lasers

Phase-locked loops are commonly used to lock two continuous wave (CW) lasers, which may be operating at slightly different frequencies. Lasers can be thought of as oscillators operating at frequencies in the hundreds of THz, so the basic principles of phase-locking can be applied.

The mixer in this case is replaced by a beam splitter, which adds together the electric fields of the two lasers with frequencies of f_1 and f_2 and directs the combined beam to a photodiode for detection. Because the photodiode detects the intensity, which is proportional to the square of the electric field, the signal includes cross terms that oscillate at $(f_1 + f_2)$ and $(f_1 - f_2)$. The photodiode acts a low-pass filter, as it is not fast enough to detect the oscillations at the sum frequency $(f_1 + f_2)$ or at the original laser frequencies f_1 and f_2.

If the two signals are at different frequencies, the error signal read out from the photodiode will oscillate at the difference frequency $(f_1 - f_2)$, with voltage related to the instantaneous phase difference between the two lasers as described in the Phase Locking section. If the signals are at the same frequency, the output is related simply to the phase difference between the two lasers.

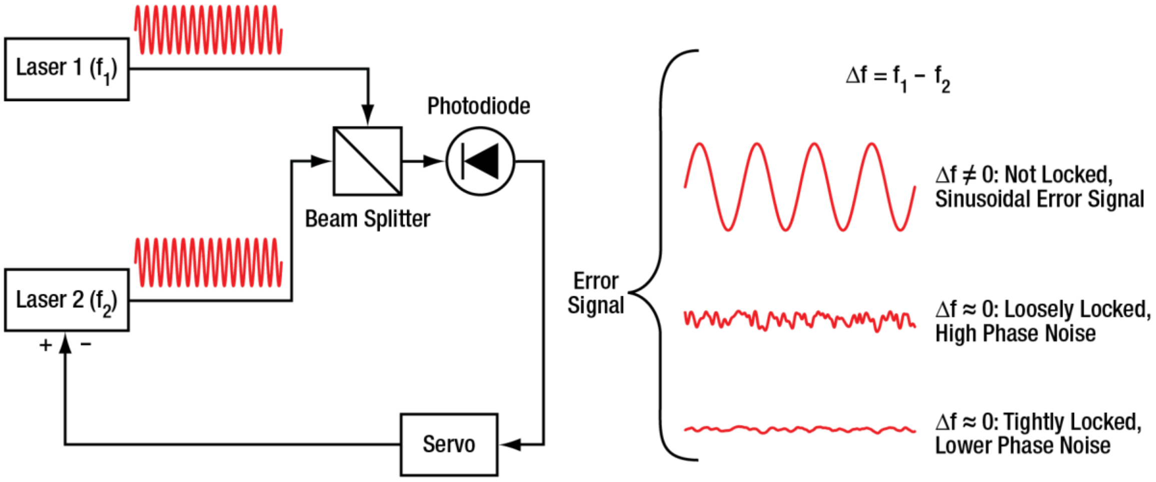

Figure 5. Two-Laser Phase-Locked Loop The left side of this schematic shows a typical phase-locked loop for two lasers. The two laser signals are combined by the beam splitter and detected by the photodiode, producing the error signal. Three common types of error signals and the conditions that produce them are shown on the right. The error signal is fed into the servo, which adjusts one of the lasers to compensate for the measured phase error.

Experimental Example: Locking Two Optical Path Lengths

In some cases, the optical path length of a system needs to be locked to a constant value, for example in fiber noise cancellation. In such cases, a Michelson interferometer configuration can be used to split a single CW laser beam into two paths and recombine onto a photodiode. The output of the photodiode in this case is a measurement of the phase difference between the two paths, and the paths can be locked to a constant optical phase difference by feeding back on one of the path lengths. This is shown in Figure 6.

This procedure can also be performed using a pulsed laser, but additional considerations related to the coherence length and the pulse repetition rate will need to be accounted for when constructing the setup.

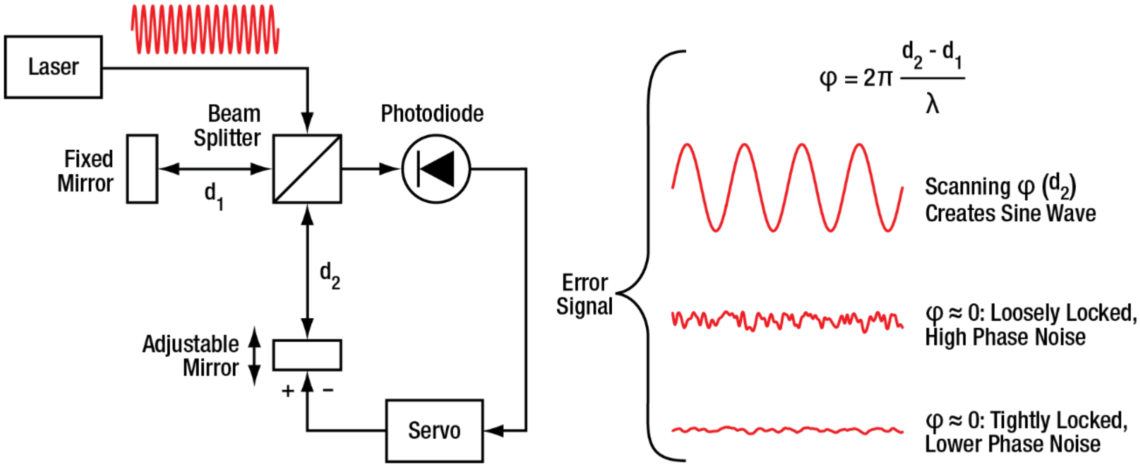

Figure 6. Optical Path Length Phase-Locked Loop The left side of this schematic shows a typical phase-locked loop for two optical paths. The signal from a single laser is split into the two paths of lengths d_1 and d_2 by the beamsplitter, reflected, and recombined in a Michelson interferometer configuration. The recombined beam is detected by the photodiode, producing the error signal. The error signal depends on the phase difference φ, which is related to d_1, d_2, and the wavelength λ by the equation in the top-right. Three common types of error signals and the conditions that produce them are shown on the right. The error signal is fed into the servo, which moves the adjustable mirror to compensate for the measured phase error.

This interferometer system will be used to demonstrate how to calculate phase noise from an analysis of the error signal.

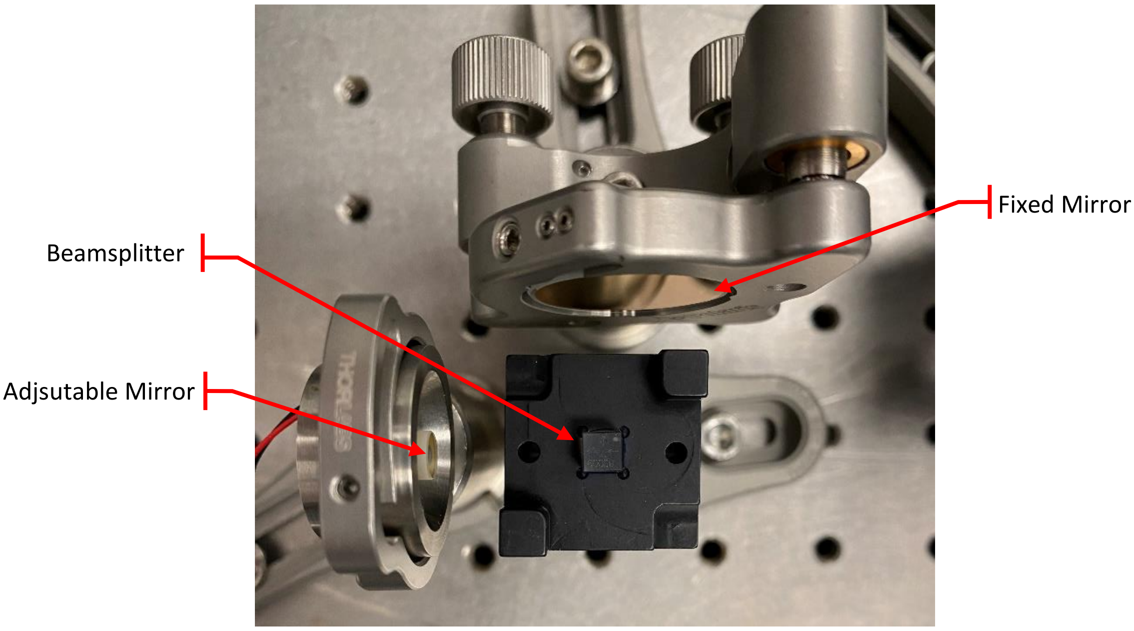

Figure 7. Experimental Interferometer Setup This photograph shows the interferometer setup used in the optical path length phase-locking experiment. The laser, photodiode, and piezo feedback and control electronics are not pictured.

Experimental Setup

A photo of the interferometer setup used for this experiment is shown in Figure 7. An SFL1550P 1550 nm External Cavity Laser in a butterfly mount was used as the laser source. This laser emits a single frequency with a typical linewidth of 50 kHz, ensuring there were no mode-hopping or mode-beating effects in the interferometer.

The laser light was split with a BS009 Non-Polarizing Cube Beamsplitter for 1100 - 1600 nm, mounted in a BS5CAM Beamsplitter Cube Adapter. The fixed arm of the interferometer consisted of a PF10-03-M01 Ø1" Protected Gold Mirror mounted in a POLARIS-K1 Kinematic Mount. The adjustable arm consisted of a PF03-03-M01 Ø7 mm Protected Gold Mirror glued to a PA44RKW Ø6.0 mm Piezo Ring Chip, which was itself glued to a custom metal adapter for mounting in a POLARIS-B1S Fixed Mount.

A DET10C2 Biased InGaAs Detector was used to detect the recombined light from the interferometer, and the signal was sent to an analog PID servo controller (SRS SIM960) as well as an oscilloscope or a PNA1 Intensity Noise Analyzer. The feedback control signal from the PID controller was sent to a BPA100 High-Voltage Piezo Amplifier connected to the piezo ring chip, to close the phase-locked loop.

Oscilloscope Measurements

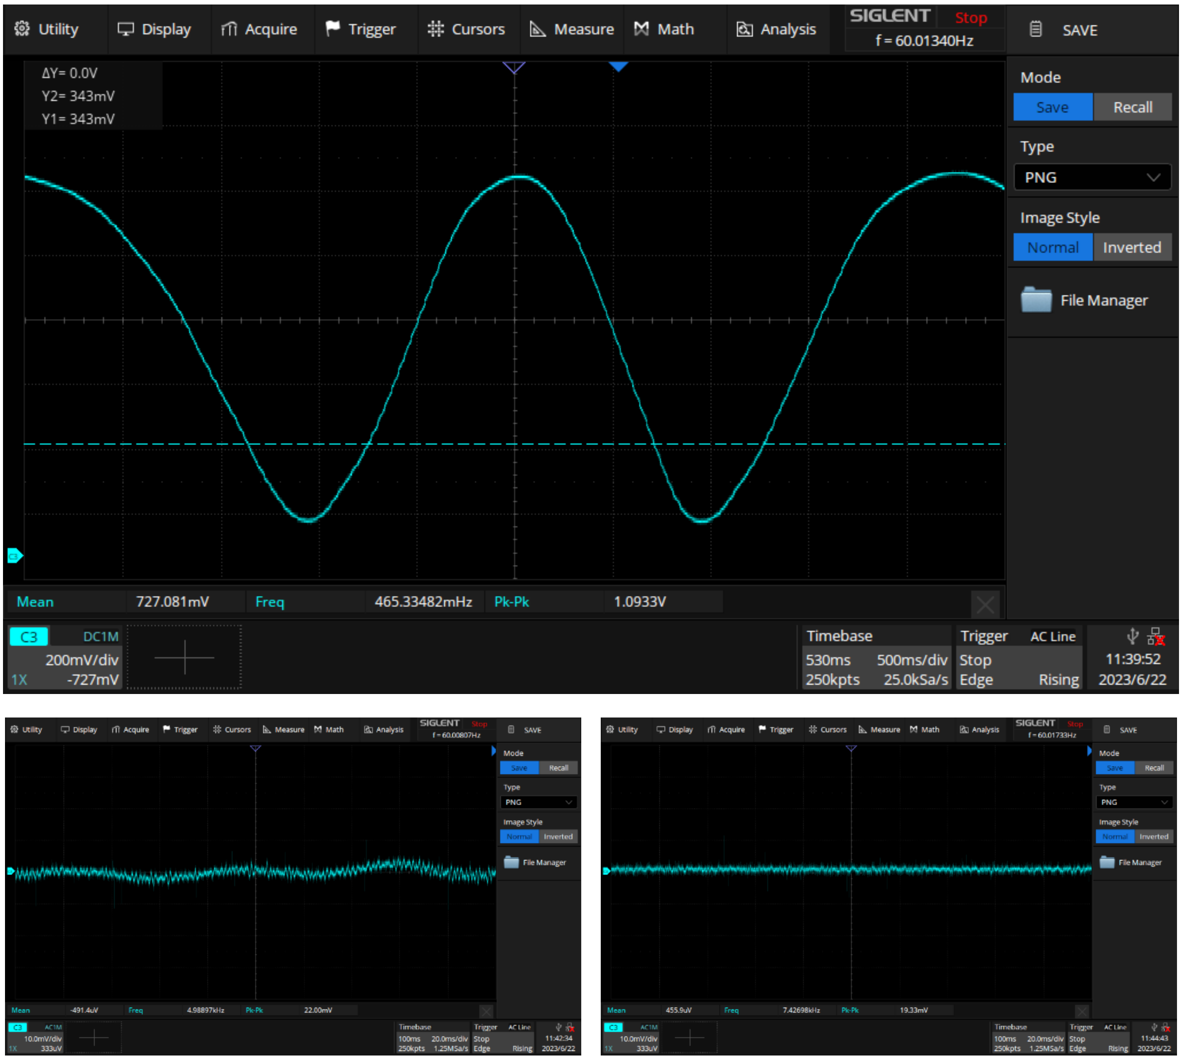

Figure 8. Oscilloscope Traces TOP (200 mV/div): This error signal was produced by scanning the piezo chip through multiple wavelengths of displacement, which is equivalent to scanning the phase difference φ through multiple periods. The trace is not exactly sinusoidal because the piezo displacement is not linear with applied voltage. BOTTOM-LEFT (10 mV/div): This is the error signal for a 1 s long window with the PID controller set to loosely lock the path lengths. Note that the voltage scale is 20 times smaller than the TOP scale. BOTTOM-RIGHT (10 mV/div): This is the error signal for a 1 s long window with the PID settings of the controller optimized to tightly lock the path lengths. The voltage deviations are visibly smaller than in the loosely-locked case.

PNA1 Analyzer Measurements

While using an oscilloscope can provide a good qualitative picture of the phase noise, the frequency spectrum of the noise is often of more interest when tuning the servo controller and identifying the sources of residual noise. The PNA1 analyzer is a convenient tool for obtaining such a spectrum.

Figure 9. PNA1 Traces This screenshot of the software included with the PNA1 analyzer shows two phase noise measurements overlaid. The Reference (blue) trace was obtained when the path lengths were weakly locked, while the Measured (green) trace was obtained after optimizing the PID controller settings for a tighter lock. This screenshot was taken in power spectral density (PSD) mode.

In this setup, the laser system was optimized to give an error signal with a sine wave amplitude (A) of 0.5 V and voltage offset (V_0) of 0.5 V, which was also the lock point of the PID servo controller.

As a result, the phase noise in dBrad^2/Hz[b] was equal to the PSD (in units of dBV^2/Hz) divided by the sine wave amplitude A (0.5 V/rad).

[b] $Phase\ in\ dBrad = 10log_{10}\left(\frac{Phase\ in\ rad}{1\ rad}\right)$

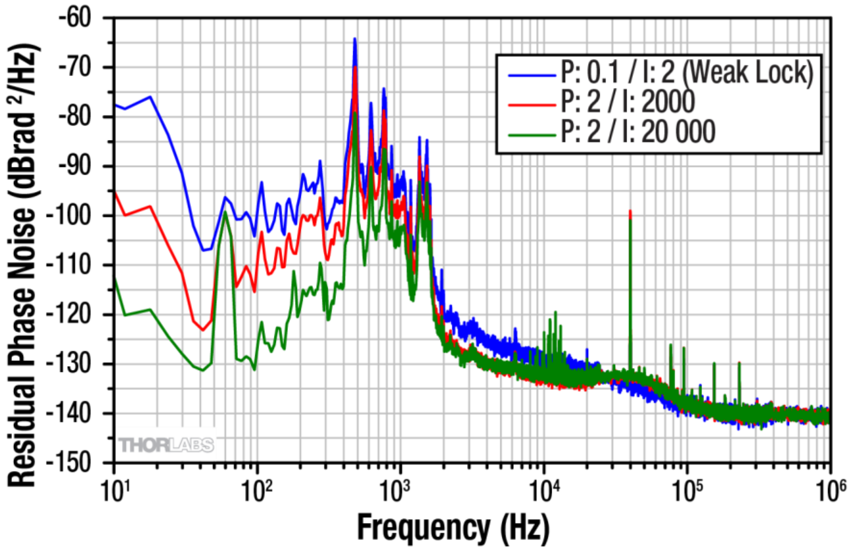

Figure 10. Residual Phase Noise This graph shows the phase noise as a function of frequency for three different PID controller settings. The phase noise decreased across most of the spectrum as the settings were adjusted from the "weak lock" settings (blue). Increasing the proportional (P) and integral (I) gain reduced the noise. The derivative (D) functionality of the controller was not used as it was found to add excess noise between 8 and 15 kHz.

Several features of note can be seen from a visual inspection of Figure 10:

1. Most of the noise in the weakly locked interferometer error signal was in the low frequency range (0 - 20 kHz). This was likely caused by air currents, acoustic noise, and seismic noise.
2. The discrete spikes from 40 kHz and up may correspond to switching power supply noise on either the laser driver or the PID electronics.
3. As the proportional (P) and integral (I) gains were increased, the noise at low frequencies reduced dramatically, with the exception of 60 Hz and harmonics. This is AC line noise and could have entered the system in a number of places, but most likely via the controller.

Figure 11 shows the effect of continuing to increase the proportional (P) and integral (I) gain. The bump that starts forming from 50 - 60 kHz as the P gain is increased is the so-called "servo bump". The system has a natural response time, which introduces lag in the actuator relative to the measured error signal. Depending on the frequency, the proportional response has a phase relative to the signal. The servo bump occurs where the response is 180° out of phase, causing noise at this frequency to be amplified rather than reduced. This effect is generally limited by the output impedance of the piezo amplifier, combined with the capacitance of the piezo chip forming a low pass filter.

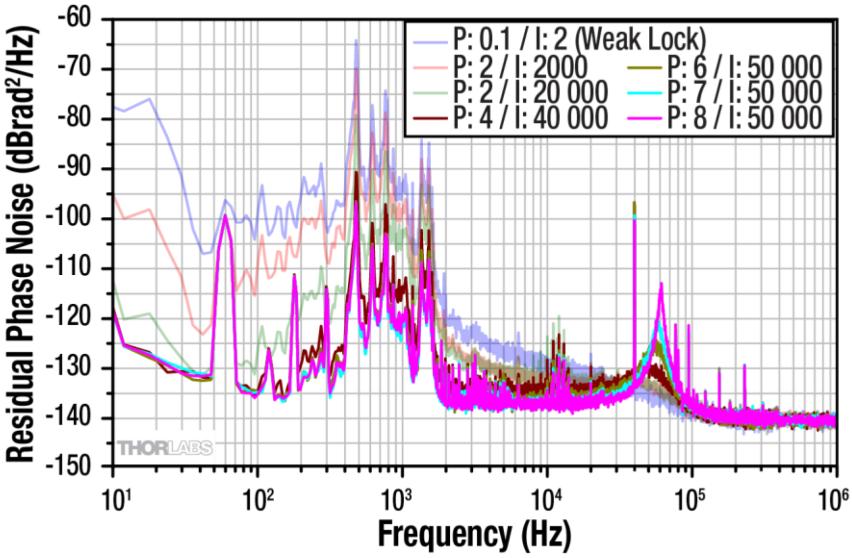

Figure 11. Residual Phase Noise for Optimized Parameters This graph shows the effect of further optimizing the servo's PID parameters. As the P and I gain settings are increased, the phase noise continues to decrease across most of the spectrum. However, increasing the P gain eventually causes the servo bump to appear around 60 kHz.

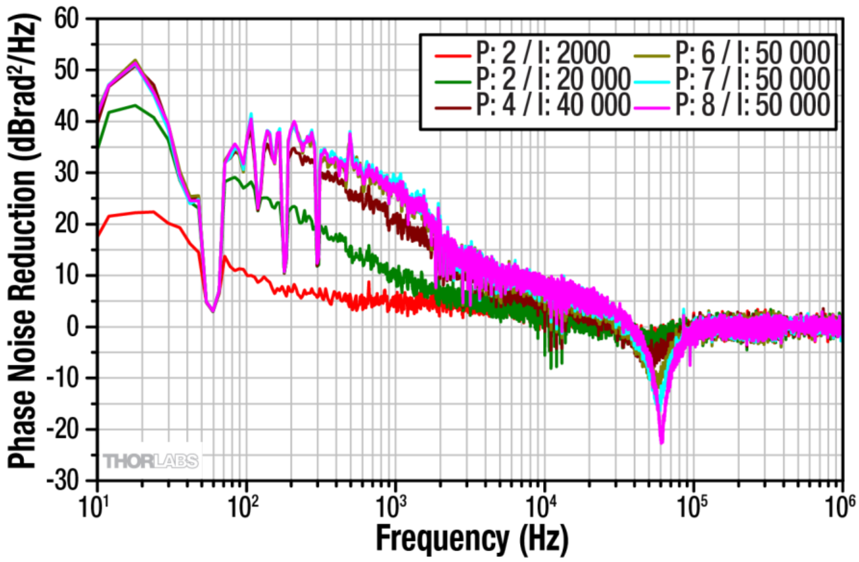

Figure 12. Phase Noise Reduction This graph shows the phase noise reduction at each setting, compared to the "weak lock" setting, where positive values indicate a reduction in phase noise. Increasing the I gain reduced noise more effectively at lower frequencies, whereas increasing the P gain reduced the noise more evenly across the bandwidth. In this example, the feedback loop reduced noise at frequencies up to approximately 30 kHz, approximately half of the frequency where the servo bump appears. At low frequencies, noise reduction of up to 50 $dBrad^2/Hz$ was achieved.

Integrated Phase Noise

To quantify the full impact of the phase noise in the interferometer, the phase noise can be integrated over the bandwidth of interest. In this case, we integrated the phase noise from 6 Hz to 1 MHz. The table below shows the total integrated phase noise for each combination of P and I settings. With weak locking, the total integrated phase noise was already very small due to the small size and passive stability of the interferometer. Even so, optimizing the lock reduced the RMS phase noise by more than an order of magnitude.

PID Controller Settings	Integrated Phase Noise (mrad)
P: 0.1 / I: 2 (Weak Lock)	2.778
P: 2 / I: 2000	1.610
P: 2 / I: 20 000	0.618
P: 4 / I: 40 000	0.277
P: 6 / I: 50 000	0.242
P: 7 / I: 50 000	0.207
P: 8 / I: 50 000	0.225

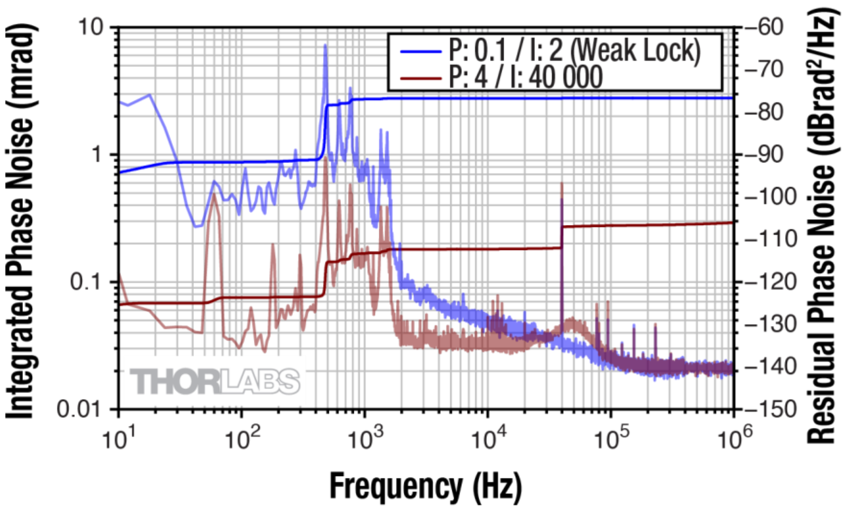

Figure 13. Integrated Phase Noise This graph shows the phase noise integrated from 6 Hz to the given frequency for two PID controller settings (solid lines), along with the phase noise (transparent lines). The sharp jump near 40 kHz for the P: 4 / I: 40 000 setting is due to the electronic noise and illustrates how the relative contributions of different frequencies can change as the controller settings are optimized. In this case, an external source of noise that cannot be reduced by changing the controller parameters becomes a dominant contributor to the integrated phase noise.

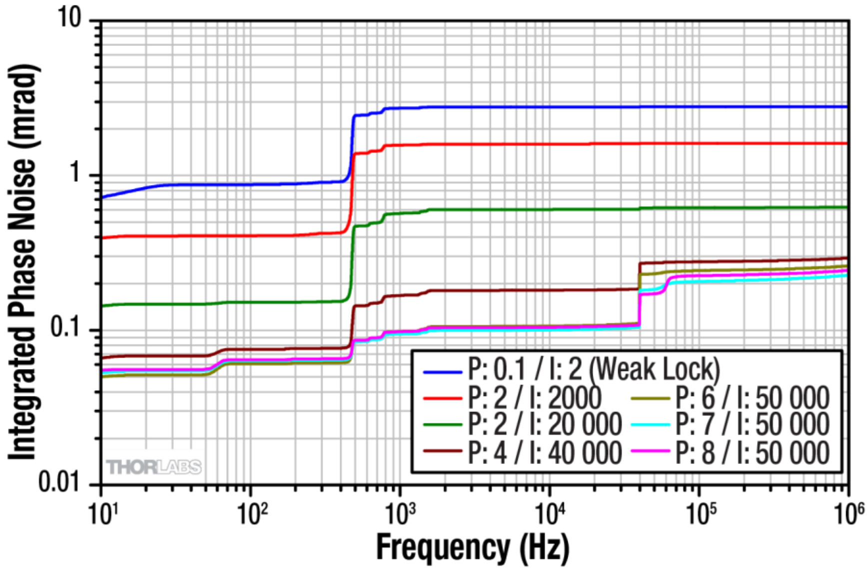

Figure 14. Integrated Phase Noise for Optimized Parameters This graph shows the integrated phase noise for all of the controller settings used. The value of each curve at 1 MHz, which is the upper end of our bandwidth of interest, can be seen to initially decrease as the P and I gains are increased. When the P gain is increased to 8, however, this value increases rather than decreasing. By a visual inspection of the graph, this effect can be attributed to the larger servo bump around 60 kHz.

Integrated phase noise can be directly correlated with noise in the relative lengths of the two arms in the setup by converting from radians to nanometers: $\Delta L \approx \frac{\varphi\lambda}{2\pi}$. This is not an equivalence, because it ignores other contributions to the apparent phase noise, including laser amplitude and frequency drift. However, a calculation ignoring those contributions can give a useful approximate figure for how closely the two lengths are locked. With the best lock settings, the two path lengths in the interferometer were locked together to within approximately $\frac{(0.2\ mrad)(1550\ nm)}{(2\pi\ rad)} = 0.05\ nm$ root-mean-square, which is approximately the diameter of a hydrogen atom.

Note that this analysis uses the in-loop error signal, and so is not a good measure of the absolute stability of the system. Systematic effects, like laser wavelength drift or responsivity changes of the detection photodiode with temperature, will make the out-of-loop absolute stability worse than the measurement performed here. However, this approach is very useful for determining the best servo lock settings and optimizing in the range from 10 Hz to a few MHz.

About Thorlabs

Thorlabs was founded in 1989 with a simple mission: make photonics lab equipment more accessible to graduate students and junior professors.

To this day, we have a love for science and an eye toward the future. We work hard to design and manufacture the tools that enable the next great steps in scientific discovery.

Contact Us

Sales and Tech Support
(973) 300-3000
www.thorlabs.com

Our Office

43 Sparta Avenue
Newton, NJ 07860

Connect

X.com/thorlabs
Facebook.com/thorlabs
Instagram.com/thorlabs
LinkedIn.com/thorlabs
Youtube.com/ThorlabsVideos