Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

- Optics Track -

Optical Resonators & FM Spectroscopy

Overarching objectives

The *Optics Track* in Ph77 is an eight-week set of laboratory and prelab exercises focusing on optics and optical instrumentation. Although we cannot include every hardware specialty in this brief time, we try to cover some especially common types of optical equipment together with a foundational description of measurement techniques and mathematical principles. One of our primary goals is for you to gain familiarity with the kinds of optical hardware you might soon encounter in a university research lab in graduate school. However, given how ubiquitous optics has become in modern science and technology, you may well find this knowledge quite useful if your career path involves using or designing a large variety of laboratory experiments, lasers, spacecraft, telescope instrumentation, optical telecom technology, commercial products, or any other equipment that incorporates optics.

The supplementary documents *General Intro.pdf* and *eNotebook Example.pdf* (available on Canvas) contain important information that is common to all the Ph77 Tracks. Please read both these short documents before proceeding. If you are really serious about learning optics at a professional level, we recommend that you purchase the textbook *Optics*, by Eugene Hecht, and look through it independently while working on this Track.

This handout is divided into two Parts. You should complete **Part 1** for the first two-week Assignment in Canvas and **Part 2** for the second two-week Assignment.

Figure 1. The optical table used for the Cavity Lab in Room 210 East Bridge.

Part 1. Optical resonators

Optical resonators (also called *optical cavities* or *Fabry-Perot etalons*) are ubiquitous elements in optical physics and technology, being useful for many diverse applications such as sensitive wavelength discriminators, as stable frequency references, and for building up large field intensities with low input powers. Also, all lasers begin with an optical resonator. Optical resonators are often made from two curved high-reflectivity surfaces separated by a distance L as shown in Figure 2. We can typically ignore the outer flat surfaces because these have high-quality anti-reflection (AR) coatings.

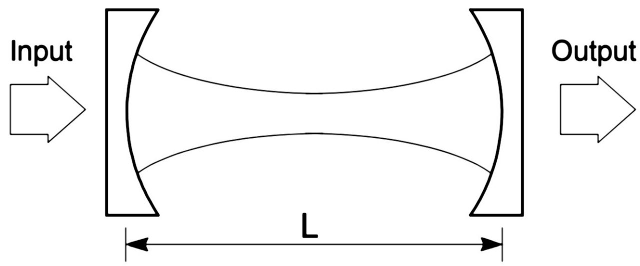

Figure 2. The basic Fabry-Perot cavity. The curved surfaces of the mirrors are coated for high reflectivity, while the flat surfaces are anti-reflection coated and have negligible reflectivity. The curved lines inside the cavity represent the shape of the resonance optical mode.

The plane-wave approximation

Although curved mirrors are the norm in optical resonators, much of the underlying physics can be understood by considering the flat-mirror case shown in Figure 3, which reduces the problem to one dimension. Physically, this case can be realized if the flat mirrors have effectively infinite extent and the input light can be approximated by a perfect plane wave, and this is a limiting case of the curved-mirror cavity. For two identical flat mirrors, each with reflectivity R and transmission T (and we assume lossless mirrors with $R+T=1$), the amplitude of the transmitted and reflected electric field amplitudes (which you will calculate as an exercise) are given by

$$E_t = \frac{Te^{i\delta}}{1-Re^{2i\delta}}E_0 \tag{1}$$

$$E_r = \frac{(1-e^{i2\delta})\sqrt{R}}{1-R\ e^{i2\delta}}E_0 \tag{2}$$

where E_0 is the amplitude of the incident light and $\delta = 2\pi L/\lambda$ is the phase shift of the light after propagating through the cavity (we assume the index of refraction is unity between the mirrors). The transmitted light intensity is then

$$\frac{I}{I_0} = \left|\frac{E_t}{E_0}\right|^2 = \left|\frac{Te^{i\delta}}{1-Re^{2i\delta}}\right|^2 \tag{3}$$

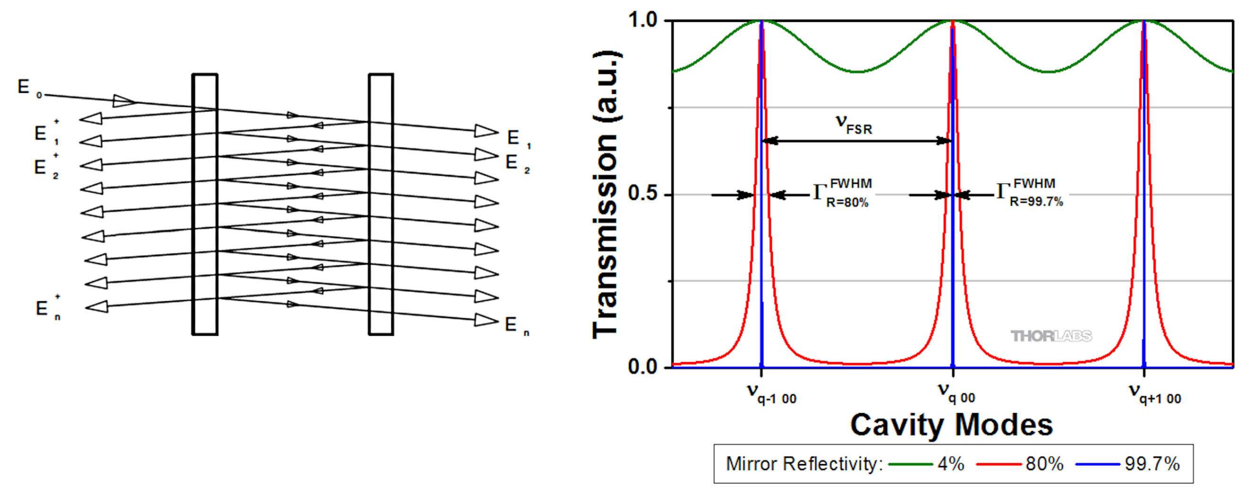

Figure 3. The optical properties of a Fabry-Perot cavity result from the interference of reflected and transmitted waves. (Left) Here the outer surfaces of the glass plates are assumed to have perfect AR coatings, so reflections only occur on the inner surfaces. A nonzero incident angle is shown for clarity. (Right) This graph shows the cavity transmission as a function of light frequency, shown for different mirror reflectivities.

The cavity transmission peaks when $e^{2i\delta} = 1$, or equivalently at frequencies $f_m = mc/2L$, where m is an integer, c is the speed of light, and L is the mirror spacing. At these frequencies the cavity length is an integer number of half-wavelengths of light, giving a standing wave inside the cavity. Note that the peak transmission for this lossless case is $I/I_0 = 1$, regardless of R.

The separation between adjacent peaks is called the *free-spectral range*, given by

$$\Delta f_{FSR} = \frac{c}{2L} \tag{4}$$

If the mirror reflectivity is high (for the cavity mirrors in this lab it is approximately 99.5 percent) then the transmission peaks will be narrow compared with Δf_{FSR}. The full-width-at-half-maximum, Δf_{FWHM}, (*i.e.* the separation between two frequencies where the transmission is half the peak value) is written as

$$\Delta f_{FWHM\ (of\ transmission\ peak)} \approx \frac{\Delta f_{FSR}}{F} \tag{5}$$

where F is called the *cavity finesse*. If $T \ll 1$ the finesse can be shown from the above to be approximately

$$F \approx \frac{\pi\sqrt{R}}{1-R} \approx \frac{\pi}{T} \tag{6}$$

If we add some absorption or other loss of light inside the cavity, the transmitted electric-field intensity becomes

$$E_t = \frac{T\alpha e^{i\delta}}{1 - R\alpha^2 e^{2i\delta}} E_0 \tag{7}$$

where α is a loss parameter and the fractional intensity loss from a single pass through the cavity is equal to $\varepsilon = 1 - \alpha^2$. This gives a peak cavity transmission

$$\frac{I_{peak}}{I_0} \approx \frac{T^2(1-\varepsilon)}{(T+\varepsilon)^2} \approx \frac{T^2}{(T+\varepsilon)^2} \tag{8}$$

and a finesse

$$F \approx \frac{\pi\alpha\sqrt{R}}{1-\alpha^2 R} \approx \frac{\pi}{T+\varepsilon} \tag{9}$$

Exercise 1. Derive the above expressions for E_r and E_t as a function of δ when $\varepsilon = 0$. [Hint: make a sketch showing the first several reflected and transmitted beams and include the electric field amplitudes on both sides of each of the two reflecting surfaces. Use reflected and transmitted amplitudes r and t, where $r^2 = R$ and $t^2 = T$. After a few terms you can see how to sum the series (the transmission sum is simpler, so start with that). Use the fact that if the reflected amplitude is r for light entering the cavity from outside, is $-r$ for light reflecting from inside the cavity (this direction-dependent sign flip appears for all reflecting surfaces).] Then square the electric field amplitudes to plot the transmitted and reflected intensities, I_{tran}/I_0 and $I_{reflect}/I_0$ of a cavity as a function of δ for $R = 0.8$, 0.95, 0.99 and $\varepsilon = 0$.

In the time domain

It is instructive to think about what happens when you first turn on a laser beam incident on a cavity. At its first encounter with the front mirror, the beam mostly reflects off, because R is close to unity. But some light passes through the mirror and ends up reflecting back-and-forth inside the cavity (i.e., between the mirrors). If the spacing L between the mirrors is nothing special, then the light that continues leaking into the cavity from the incident plane wave will randomly interfere (both constructively and destructively) with the existing light inside the cavity. The various phase relationships rapidly change with time, with the result that little light builds up inside the cavity. In steady-state, there is always a small but fluctuating amount of light rattling around inside the cavity, and some of that light leaks out each of the two mirrors. But nearly all the light just reflects off the front mirror.

The situation becomes more interesting when the incoming light resonates with a standing wave inside the cavity, which happens when $L = m\lambda/2$, where m is any integer. For these special cases, the light that continues leaking into the cavity from the incident beam always constructively interferes with the existing light inside the cavity. The moment you turn on the incident beam, therefore, the intra-cavity intensity begins to increase with time. And it continues building up until a substantial amount of light leaks out both the front and back mirrors. When you do the math, the light leaking out the front mirror (so in the backward direction) exactly cancels the "prompt" reflected light from the incident beam (before it enters the cavity), so now no light reflects off the front mirror. Meanwhile the light leaking out the back mirror has exactly the same intensity as the incoming plane wave. To the outside observer, it appears that the light has simply passed through the cavity. Moreover, the intra-cavity light intensity is now $\sim 1/T$ times that of the input beam.

Gaussian Optics

The main problem with flat-mirror optical resonators is that the plane-wave approximation does not fare well in real life. If the incident light is not precisely normal to the mirror surface, the light will "walk off" the mirror after multiple reflections. Moreover, diffraction effects mean that precisely normal incidence is not even possible in principle with finite-size beams. Early lasers were built using flat mirrors, but they worked poorly and were difficult to align. Researchers soon realized that curved mirrors provided a better method for constructing optical resonators. The resulting technology is known as *Gaussian optics*, and this topic is much studied in laser physics and anything relating to optical resonators.

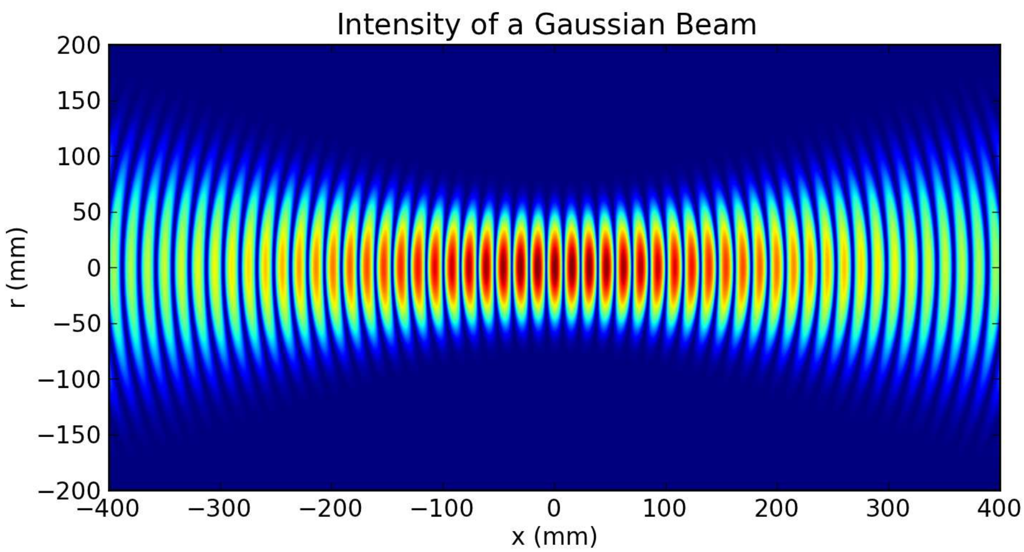

Figure 4. A typical standing-wave intensity pattern inside an optical cavity. Along the optical axis, the standing-wave nodes are separated by λ/2. In the radial direction, the light intensity exhibits a Gaussian profile, which is a nearly ubiquitous property of laser beams. In this example the beam diameter goes through a local minimum called the beam "waist". In an optical cavity, the surfaces of the two curved mirrors match the wave fronts.

What sets Gaussian optics apart from plane-wave analysis and ray optics is that diffraction is built into the mix from the outset. We will not delve into the theory in any meaningful way in this laboratory course, but Figure 4 shows you the fundamental result in a nutshell. You can think of this as a snapshot of a propagating laser beam or, if you place curved mirrors so they line up with the wavefronts at both ends, this could be a picture of a standing wave inside an optical resonator. Lining up the wavefronts requires mirror surfaces with a spherical shape, and the radius-of-curvature of the surfaces is typically much larger than the size of the mirrors.

Looking in the radial r direction in Figure 4, the overall beam intensity has a profile $I(r) \sim \exp(-2r^2/w^2)$, where $w(z)$ defines the radial size of the beam at any point in z along the beam propagation direction. The minimum value of w is called the beam *waist*, often denoted w_0, which is located at the center of Figure 4. When you focus a laser beam to a small spot, this is what it looks like. The size of w_0 depends on the size of the initial laser beam and the focal length of the focusing lens. A tight focus means w_0 is small, while a longer focal length gives a larger w_0. In all cases, however, the radial profile $I(r)$ maintains its Gaussian form. A "collimated" laser beam (like what comes out of a typical laser pointer) just means that w is nearly constant over a considerable distance. But w cannot be truly constant at all z because diffraction dictates that finite-size laser beam must expand as it propagates. For a typical laser pointer, for example, the beam exits the laser with some Gaussian diameter ~$2w$. And from there it must expand with some minimum diffraction angle $\theta \approx \lambda/w$.

When two spherical mirrors with radius-of-curvature ρ are placed some distance L apart (and $\rho > L$), then there will always be stable cavity "modes" that look something like that shown in Figure 4. The wave fronts have to match up with the mirror surfaces, and this fact determines $w(z)$ inside the cavity. Moreover, the cavity free-spectral-range and finesse will be nearly identical to the plane-wave values given above. The plane-wave case is just a limiting case of Gaussian-optics, although plane-wave mirrors have *unstable* cavity modes. Gaussian cavity modes made using spherical mirrors are not just stable, they are also somewhat insensitive to mirror alignment errors.

Because real cavity mirrors have a finite size, there is always some optical loss as the beam spills over the edge of the mirrors. But the Gaussian function drops off very quickly with r/w, so usually these losses are quite small. Absorption in the mirror coatings and dirty mirror surfaces is usually more important than losses from the Gaussian tails.

Laguerre-Gaussian modes

The image in Figure 4 shows the simplest standing-wave pattern inside an optical cavity with spherical mirrors, and these are call "longitudinal" or TEM $_{00}$ cavity modes, or sometimes just "00" modes. Optical cavities, however, much like acoustic cavities, are three-dimensional beasts that

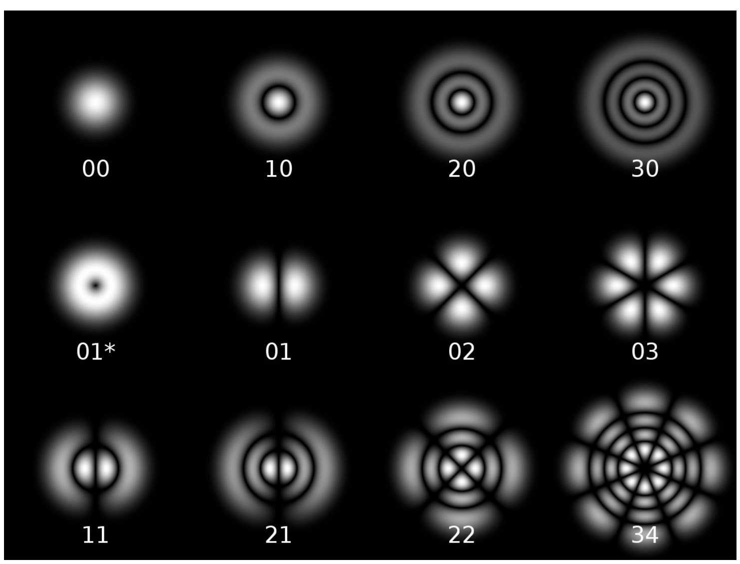

Figure 5. Several Laguerre-Gaussian modes, which are the electromagnetic normal modes inside a cylindrical Fabry-Perot cavity. The TEM01 mode, called a doughnut mode, is a superposition of two (degenerate) TEM01 modes rotated 90 degrees with respect to one another.*

support a whole family of standing-wave patterns with additional complexity. For cavities with cylindrical symmetry, these normal modes of the electromagnetic field are called *Laguerre-Gaussian modes*; the transverse mode patterns are described by a combination of a Gaussian beam profile with a Laguerre polynomial. The modes are labeled by TEM $_{p\ell}$, where p and ℓ are integers labeling the radial and angular mode orders. The intensity at a point r, ϕ (in polar coordinates) is given by

$$I_{p\ell}(r,\phi) = I_0 \rho^{\ell} [L_p^{\ell}(\rho)]^2 \cos^2(\ell\phi) e^{-\rho} \tag{10}$$

where $\rho = 2r^2/w^2$ and L_p^{ℓ} is the associate Laguerre polynomial of order p and index ℓ. The radial scale of the mode is given by w, and modes preserve their general shape during propagation.

A sample of some low-order Laguerre-Gaussian modes is shown in Figure 5, which displays the transverse mode profiles; the longitudinal profile of all the modes is essentially that of a standing wave inside the cavity with some number n of nodes, as shown in Figure 4. The various modes with different n, p, and ℓ in general all have different resonant frequencies.

The TEM $_{00}$ mode has a simple Gaussian beam profile, and this is the mode one usually wants to excite inside the cavity. Lasers typically use a 00 mode to generate Gaussian output beams with no unwanted transverse nodes, and the goal in many research/technology applications is to excite a fundamental 00 mode. As you will see in the lab, however, it is not always trivial to excite just the 00 mode inside a cavity.

An optical spectrum analyzer

One popular use of optical cavities is in *optical spectrum analyzers*, which can be used to measure the optical power in a laser beam as a function of frequency. The basic idea is shown Figure 6, where the mirror spacing L is scanned using a piezoelectric transducer (PZT). This essentially scans the

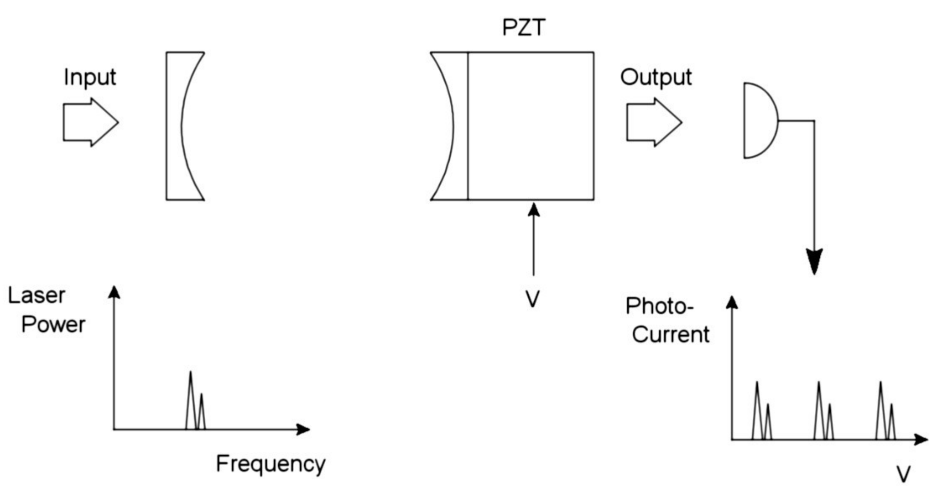

Figure 6. Using a confocal cavity as an optical spectrum analyzer. Here the input laser power as a function of frequency P(f) is shown with a multi-mode structure. By scanning the cavity length with a piezoelectric (PZT) tube, the laser's mode structure can be seen in the photodiode output as a function of PZT voltage I(V). Note the signal repeats with the period of the cavity free-spectral range.

phase δ you used in your calculations, and thus scans the cavity resonant frequencies f_m. If the laser beam contains power at frequencies in a range around some f_0, then by scanning the PZT one can record the laser spectrum, as is shown in Figure 6.

Note that there is some ambiguity in the signal from an optical spectrum analyzer; a single laser frequency f_0 produces peaks in the spectrum analyzer output at $f_0 + j\Delta f_{FSR}$, where j is any integer. If a laser contains two closely spaced modes, as in the example shown in the figure, then the output signal is obvious. But if the laser modes are separated by more than Δf_{FSR}, then the output may be difficult to interpret.

Using a tunable cavity

In the lab, you will be setting up a series of different optical resonators of different lengths and characteristics, using different input lasers as well. During the process, you will learn about the physics of lasers and optical cavities, measure some of their properties, and generally gain some hands-on experience setting up laser optics. And the theory will make more sense once you start playing with the hardware. We will start with a fixed-wavelength laser and a tunable optical cavity.

Optical setup

Your first task is to set up the optical layout shown in Figure 7, and everything you need should be found on the Cavity-Lab optical table. As a general rule, it is best to leave the larger components fixed on the table (the laser, electronics boxes, etc.) while moving the smaller optical components (mirrors, lenses, photodetectors, etc.) to make your setup.

Start by turning on the Helium-Neon (He-Ne) laser using the white timer box behind the laser. If you press the 4-hour button, the laser will shut itself off after four hours. We use this timer because lasers are expensive and they wear out over time, so it's best to leave them off when not in use. You should see a bright red laser beam that has a fixed wavelength of 633 nm. You will not

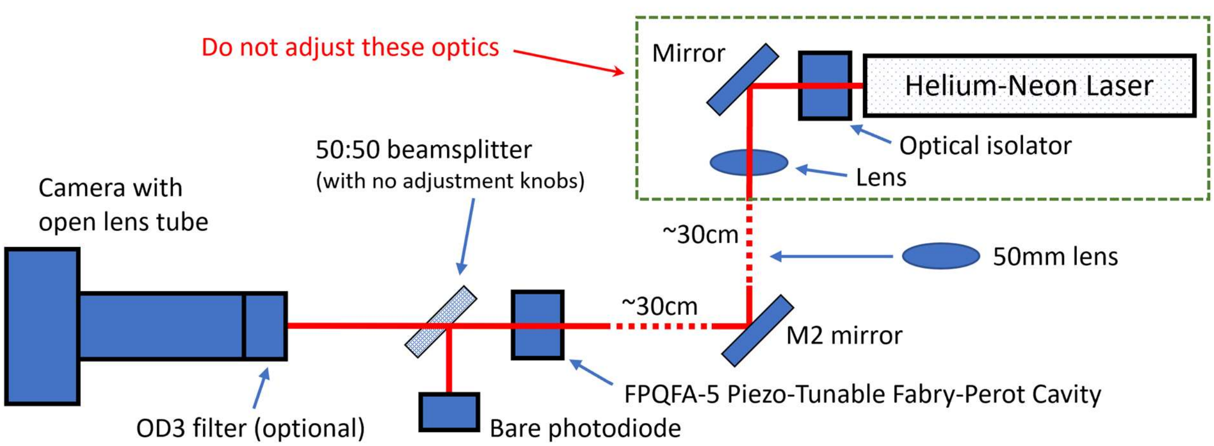

Figure 7. The optical layout for sending a Helium-Neon laser beam through a 5-mm-long tunable cavity and observing the transmitted light. There is no lens on the camera, just an open tube (here shown with an OD3 attenuation filter).

need to wear laser goggles when using this laser. The beam is visible and the intensity is fairly low, so it is comparable to the Sun in terms of eye safety. You can hurt your eyes by staring directly at the Sun, but billions of people don't do that every day. Low-power visible lasers like laser pointers and He-Ne lasers are safe for the same reason; just don't stare into the beam.

The He-Ne is a gas laser that sends a high-voltage electrical discharge through a glass tube inside the aluminum tube you see on the table. The energetic electrons in the discharge excite Helium atoms, which subsequently excite Neon atoms via collisions in the gas. The goal of this is to excite a specific metastable energy level in Neon, which then decays back down by emitting 633nm photons. Why the roundabout process involving electrons and Helium atoms? Because it works, and simpler things do not. You can read all about how this laser works online, if you wish.

Before setting up all the optics, insert a 50mm lens into the beam as shown in Figure 7, with M2 not yet placed in the beam. Hold a white card in the beam about a meter downstream, and you will see a fairly clean Gaussian beam profile, for which the light intensity drops off following

$$I(r) = I_0 \exp\left[-\frac{2r^2}{w^2}\right] \tag{11}$$

and you can observe that w changes with distance along the beam. If you removed the lens, w changes more slowly with distance, but it always changes because of diffraction.

Gas lasers tend to have nearly perfect Gaussian beams, and our He-Ne laser is no exception. You may see some small non-Gaussian intensity variations, but these mostly come from dirt and other optical imperfections. These imperfections are not as bad as they look; you can see them mainly because your eye is a logarithmic detector that is especially good at noticing small variations in intensity.

Next remove the 50mm lens and start setting up the other optical components shown in Figure 7. Be sure to ask if you cannot find something or just have questions. Remember that humans are social animals, and you will learn more efficiently if you interact with other people. Before you get started placing the various parts on the table, consider these practical "rules" for setting up laser optics:

Aligning laser optics

Laser optical setups in research labs can easily involve hundreds of optical components, all of which need to be carefully aligned. It is useful, therefore, to learn some of the accepted standard practices for setting such systems up, which researchers have developed through much trial-and-error.

1) As you are handling and aligning the optics, do not touch any optical surfaces. This is important because optical coatings are fragile and easily damaged. The acids deposited by fingerprints tend to etch the coatings away, thereby ruining the expensive optical surfaces. Some people insist that gloves be worn when working with optics, but we will not go that far. If you accidentally touch any optical surfaces, or notice any fingerprints, please let us know so we can clean them (which involves standard practices as well).
2) Keep all laser beams at a constant height above the optical table, four inches in this lab. This

eliminates one degree of freedom for all optical elements (i.e. height), and that simplifies the alignment process.

3) Use right-angle beam steering when possible. This is not a hard-and-fast rule, but following it keeps all the beams on a rectangular grid, again reducing degrees of freedom and thus simplifying the alignment process. In Figure 7, for example, you see all the beams running on a rectangular grid. Use the hole pattern on the optical table to define your coordinate system. This rule has become so commonplace that many ultra-high-reflectivity mirrors must be used with right-angle reflections.
4) Keep laser beams away from the edges of optical components. Laser beams should always reflect from, or pass through, the central regions of every mirror, beamsplitter, filter, polarizer, lens, etc. Light hitting the edge of an optical mount is called *vignetting*, and clearly this should be avoided.
5) Keep mirror mounts near the middle of their pointing range. The pointing knobs are for fine adjustments only, so avoid running these fine-pitch screws near the ends of their travel.
6) Make an alignment plan, and then follow your plan. We have done this for you (because it requires some experience to know what to do), and the sketch in Figure 7 defines the optical layout including a plan for aligning the laser beam into the optical cavity. Whenever you set out to build something, planning is always the first step.

While the mirrors and lenses in Figure 7 look about like you would expect, the FPQFA-5 optical cavity is a bit of a small “black box”. What this housing contains is essentially just two cavity mirrors separated by about 5mm. The spacing can be changed using a piezoelectric transducer (PZT), which is driven by the attached cable. For now, just locate and roughly place all the optical elements in Figure 7. At this point you may not see any light hitting the camera or the photodetector, and this is normal. Making everything work requires additional alignment steps, and for that you will need to set up the electronics that goes with the optics.

Electronics setup

As with the optics, setting up the accompanying electronics means starting with a plan, and with this in mind we provide the layout diagram in Figure 8. For your convenience (and ours), note that all the BNC cables connected to the function generator are already connected. And the two cables that go from the function generator to the oscilloscope are behind the ‘scope (or nearby) and are labeled. That keeps those cables out of the way, so they are not draped over the optics knocking things over. If you want things to work in the lab, clutter is the enemy.

Note that there are two photodiode detectors in the Cavity Lab, and both are connected to the Diode Laser Controller (the large chassis on the right). One or both of the photodiodes will likely be stowed in the two post holders in front of the chassis. For this setup, be sure to use the Bare photodiode (the other has a narrow 658nm bandpass filter that blocks the 633nm He-Ne light). If you need help finding anything, or just have questions, do not hesitate to ask.

The purpose of the Thorlabs Piezo Controller is to provide a high-voltage signal that varies the length (i.e., the spacing between the two cavity mirrors) of the FPQFA-5 optical cavity. The PZT is very stable, but it only moves a few tens of nanometers per volt; hence the need for high voltage,

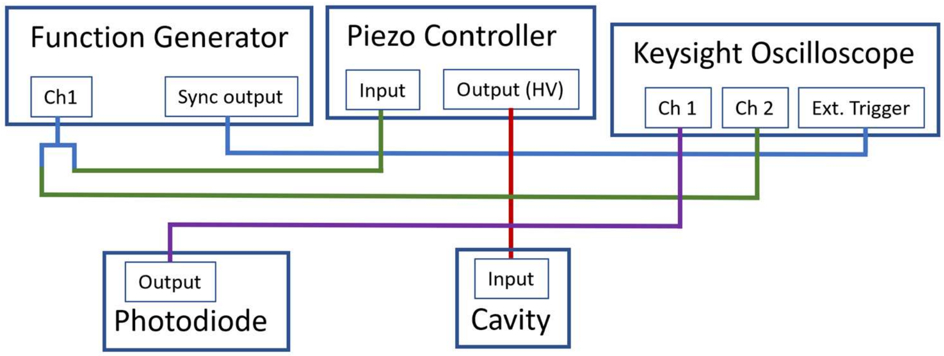

Figure 8. The electronics diagram that goes with setting up the He-Ne laser and FPQFA-5 optical cavity. For this setup, use the 2-channel 'scope already attached to the other components (not the 4-channel 'scope attached to the laser controller, which you will use later).

up to 150 volts in this case. When you turn this controller on, set the voltage to about 30 volts for now.

Next set up the function generator to produce a ramp-wave function with a frequency of 10 Hz, a LoLev of 0V (press the Amplitude button to toggle between Amplitude/Offset and HiLev/LoLev). Note that the Piezo Controller wants a positive-voltage input, so set LoLev to 0V and leave it there. Start with HiLev = 2V for now. Also on the function generator, press Utility/Sync On to turn on the Sync signal that triggers the 'scope. The function generator will probably wake up with some of these settings already in place.

On the 'scope, set the Trigger to External and have a look at ch2. Adjust the settings until you see a stable triangle wave that looks something like that shown in Figure 9. (Ignore ch1 for now.) This 'scope trace shows one sweep of the cavity length, which is all you need; more sweeps just gives you redundant information. Once you have the triangle wave centered on the screen and stable, you next need to align the cavity optics.

Separating degrees of freedom

It is useful at this point to examine why we have chosen the optical layout in Figure 7. Sending light into an optical cavity is not a simple point-and-shoot operation … the input laser beam must be aligned along the cavity axis with considerable precision. Achieving this alignment involves adjusting four **degrees of freedom** (dof). First, the incoming laser beam must strike the center of the first cavity mirror (which involves two dof: the x and y positions of the beam), and second, the angle of the incoming beam must point along the optical axis of the cavity (another two dof: the θ and φ input angles). To adjust these four dof, we use four adjustment knobs: two on the M2 mirror mount and two on the cavity mount.

With the layout in Figure 7, note that M2 mainly changes the position of input laser beam on the first cavity mirror, while the cavity mount mainly changes the angles of the cavity relative to the laser beam. This is what we mean by "separating the degrees of freedom." In practice, first you adjust M2 to center the beam on the cavity mirror. Then you adjust the cavity mount to get the entrance angles right. Of course, changing M2 changes both xy and $\theta\varphi$, so the dof are not completely separated, but it works.

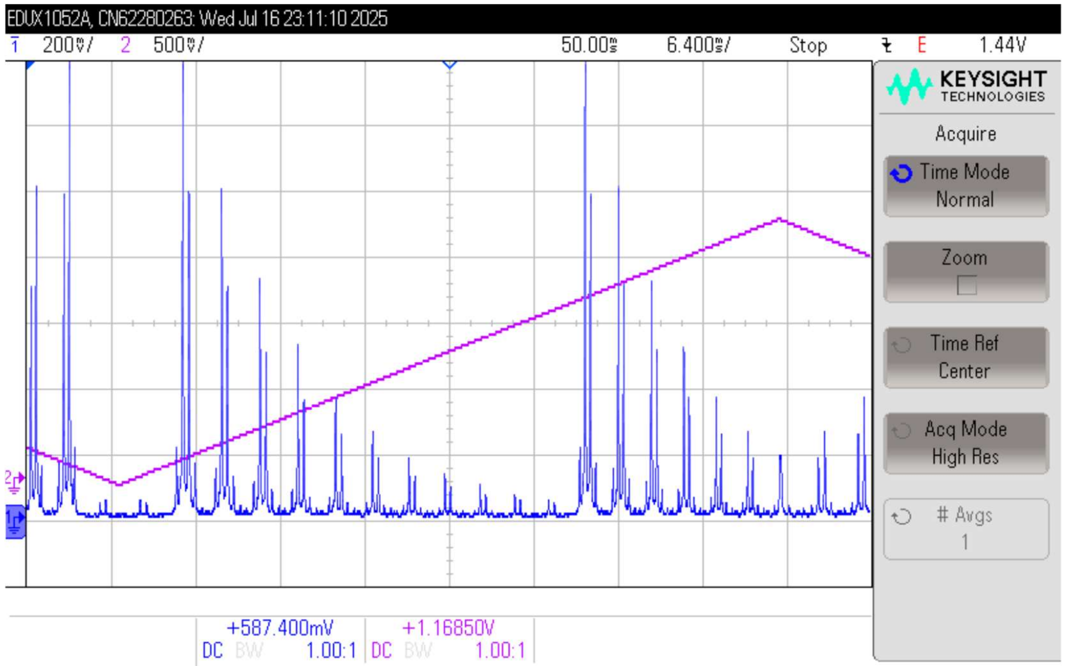

Figure 9. This oscilloscope trace shows the triangle-wave signal used to modulate the cavity length, along with the cavity transmission signal from the photodiode.

You can see the centering just by looking at the first cavity mirror, or by using a white card to see the beam. Just *do not touch any optical surfaces in the process*. The entrance angles can be determined by looking at the light reflecting off the front cavity mirror, as this beam should be "retroreflected" back toward the direction of the incoming beam. To get this roughly right, send the beam though a hole in a white card and make sure that the retroreflected beam goes back through the same hole. When the alignment is close, you can see the retroreflected beam when it hits the back side of the optical isolator. Because the cavity mirror is curved, the retroreflected beam expands before it hits the back of the optical isolator. You will probably have to iterate the alignment process to get it right. So first use M2 to center the laser beam on the cavity, then use the cavity pointing to retroreflect the light. Repeat until the series converges.

[Aside: The purpose of the optical isolator is to keep that retroreflected light from going back into the laser. The isolator uses a special crystal in a strong magnetic field that passes laser light through the crystal in one direction, but the retroreflected light does not pass through in the other direction. This works because of the Faraday effect, which you can look up if you are interested. Laser optics uses a big back of tricks, and the optical isolator is one.]

When you have these alignment steps complete, you should see a blob of transmitted light on the TV, and you should be able to see a photodiode signal on the 'scope. If you see nothing from the camera, check that no filters are in place in front of the camera tube (neither the OD3 tube nor the 658 nm filter should be in place). Next continue tweaking the alignment until you see a single oscillating spot of light on the TV, along with a signal that looks something like that shown in Figure 9. If you want to move the transmission signal back and forth on the 'scope, adjust the PZT voltage offset on the PZT controller. To zoom in or out, change the HiLev setting on the function generator. As a general rule, do not zoom in using the Horizontal scale on the 'scope; you want to keep that triangle wave centered, filling the screen.

The first thing you see in Figure 9 is that the signal repeats during a single sweep. This is because you are observing more than one free-spectral-range (FSR) of the cavity. If you look at the tallest peak at the beginning of the sweep (where ch2 is low), that peak means the cavity is resonant and contains N half-wavelengths of the laser light. The same peak later in the trace means the cavity is a bit longer so now it contains N+1 half-wavelengths of the laser light. We don't know what N is exactly, but it is roughly N = 2*5mm/633nm ~ 16000.

If you zoom in on the tallest triplet, these peaks are all 00 modes of the FPQFA-5 cavity, and they are 00 modes of the He-Ne cavity also. The fact that there are three peaks tells you that the He-Ne laser running "multi-mode" at three different frequencies. The scanning FPQFA-5 cavity is acting like an optical spectrum analyzer, as described in Figure 6.

To verify this, zoom in on that 00 triplet and you can see the peaks drifting slowly in time. To see this better, touch the He-Ne laser tube and you should find that it is warm but not so hot that it would burn you. If it's safe, wrap your hand around the tube (*gently*, so you do not move the laser pointing) to extract some heat, and you should see the peaks move more quickly. When you take your hand away, the peaks slowly move back as the laser warms up again.

This is a good illustration of how you can use a cavity as an optical spectrum analyzer. The FPQFA-5 cavity can resolve peaks separated by as little as about 100 MHz, compared with the laser frequency of $c/\lambda = 477$ THz, so a frequency resolution of about 0.2 ppm.

As for the other triplets, those arise because the He-Ne laser excites higher-other Laguerre-Gaussian modes in the FPQFA-5 cavity, and we will get into that more below.

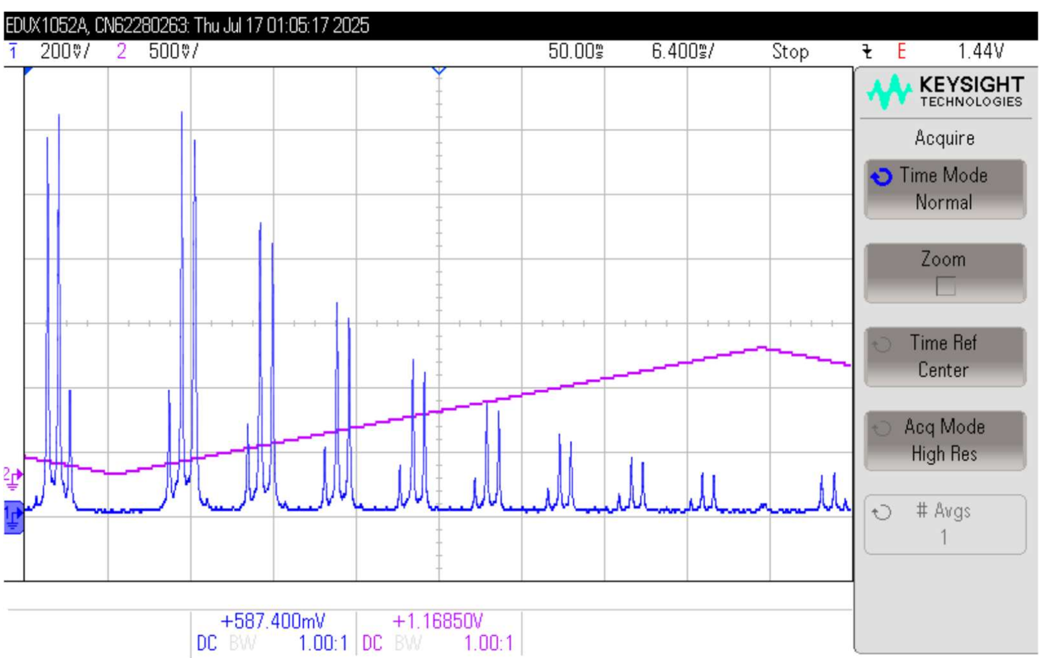

Figure 10. When the cavity alignment is good, you should see a series of "triplet" peaks in the cavity transmission signal.

Exercise 2. Before discussing more about the physics of what's going on, produce two screen shots like those in Figure 9 and Figure 10 for your e-notebook. Optimize the signal by increasing the height of the tallest peaks, and by reducing the smaller peaks you see between the triplets. Getting a nice spectrum requires iterating the fine adjustments of all four knobs. Having a good signal on the 'scope makes it easier to understand the theory, because you can relate the two in real time.

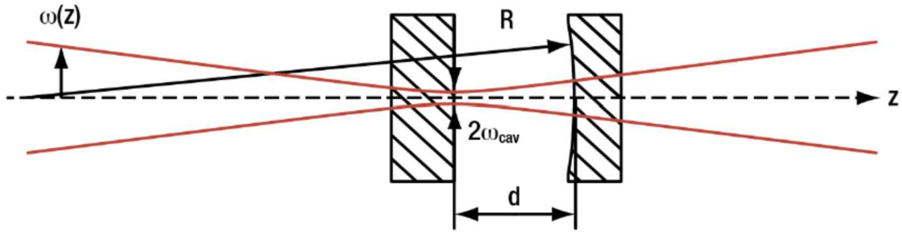

Figure 11. This diagram shows the plano-concave setup for the FPQFA-5 optical cavity. The specifications tell us that d = 5mm, R = 250mm, and the free-spectral range (FSR) is c/2d = 30 GHz. The finesse is >300, the cavity resolution is ~100 MHz, and the cavity waist is $w_{cav} = 83\ \mu m$ *at 633nm.*

Understanding the physics

Now that you have a clear signal on the oscilloscope, it is time to focus on understanding the physics behind the signal. To begin, Figure 11 shows the geometry of the FPQFA-5 optical cavity, along with some of its parameters.

When you sweep the cavity voltage, you are sweeping d, which sweeps the frequencies of all the FPQFA-5 cavity modes. Looking at just the cavity TEM $_{00}$ mode (the tallest triplet on the 'scope), this mode frequency changes linearly with the applied voltage during the sweep. When the cavity 00 mode frequency equals one of the laser frequencies in the triplet, then there is a spike in the cavity transmission. If the TEM $_{00}$ mode was the whole story, then you would just see the tallest triplet in the spectrum. The three peaks mean that the laser is running at three distinct frequencies (which we will discuss more below). The other triplets correspond to higher-order modes in the FPQFA-5 cavity, which adds some complexity to the problem. One step at a time.

Before you leave the lab, make sure you know how to do the following exercises. Better still, do these exercises while you are in the lab, so you can get help if needed.

Exercise 3. Save a .csv spectrum like that shown in Figure 9 and plot it for your e-notebook. Use these data to measure the finesse and resolution of the FPQFA-5 cavity. Define the finesse as $F = f_{FSR}/\Delta f_{FWHM}$, where f_{FSR} is the free spectral range and Δf_{FWHM} is the full-width-half-maximum of a sharp spectral peak (also called the spectral resolution). You may assume that the $f_{FSR} = 30$ GHz, as this comes directly from the cavity length d. Pro tip: When you have a nice spectrum set up on the 'scope, do an 8-trace average and then press Run/Stop to freeze the signal. Now you can zoom in using the Horizontal Scale and measure the various spectral features using the 'scope cursors.

Exercise 4. Using this same spectrum, what is the spacing between adjacent peaks within one of the triplets, in MHz? Assume that these peaks are adjacent modes of the He-Ne laser cavity. Use this information to calculate the length of the He-Ne gas tube (specifically, the spacing between the mirrors in the He-Ne cavity). [Hint: it is certainly not longer than the outer aluminum laser tube, and it's not a lot shorter than that either.] Show your work.

Exercise 5. Measure dx/dV for the PZT in the FPQFA-5 cavity, which is how much the cavity length changes for a change of 1V on the PZT. How do you make this measurement? Use your understanding of optical cavities to figure out a way! Hint: use the voltage output seen on the Piezo Controller. Show your work.

Exercise 6. Assume the gas atoms in the He-Ne laser have a kinetic temperature of around 350 K. How much is the Doppler broadening of the 633nm transition in neon atoms at this temperature? Pro tip: This is a terrific question to ask for favorite AI; it's a simple question, but it involves some tedious math that you don't want to bother with. Just make sure you understand the answer you get back. Sometimes your AI just doesn't know what it's talking about.

If you understand these exercises, then you can see why there are three peaks in the 00 triplet, no more and no less. At any given instant in time, there is an ensemble of neon atoms moving at just the right velocity to interact with one of the laser cavity modes. These Doppler-shifted atoms emit light at the same frequency as that particular cavity mode, and the cavity was engineered so only 00 modes will lase. (The linewidth of a Neon atom at rest is very narrow, so you can assume that the atomic transmission is monochromatic… but the Doppler shift changes the frequency in the rest frame of the cavity mirrors.

At the same time, there are two other ensembles of atoms that form independent lasers operating at the neighboring cavity modes. Although these three lasers exist in the same space, they act independently because each laser uses different atoms moving at different velocities. And three is just the number of cavity modes that fit into the Doppler width. (If you found that the Doppler width is not about three times the laser cavity FSR in your calculations, then clearly something is wrong.)

When you took some heat out of the laser tube, this changed the He-Ne cavity length slightly, which in turn changed the He-Ne mode frequencies. But the Neon atomic transition remained constant. So, what you saw was a series of cavity peaks passing through the Doppler width of the Neon atoms. Put your hand on the He-Ne tube again and you will see that the peaks move while the envelope remains fixed. Lots of good physics to unpack here!

Exercise 7. Zoom in on the tallest triplet by setting HiLev to 0.1 volts (which you may want to do in steps, so you can stay with the tallest triplet). Do not zoom in using the 'scope… make sure you still see the triangle wave you see in Figure 9. Looking at the TV screen, you should see a clear TEM $_{00}$ mode shape. Add the ND3 filter to the end of the camera tube so the image is not so badly saturated. Move the spectrum over to look at the other triplets and you should be able to make the correspondences you see in Figure 12. Take pictures of the different modes and combine them with a .csv file of the 'scope data to recreate Figure 12. (It's okay to take photos of the TV screen using your cell phone. The camera image on the SC card will be better, but identifying modes does not require a high-fidelity image.) [Related Pro Tip: you will save a lot of time if you put everything in your e-notebook now, while you collect the data in the lab, together with at least some brief notes also in your e-notebook. Otherwise you have to sort through the jumble of plots, images, and notes later, which takes more time than you think it will.]

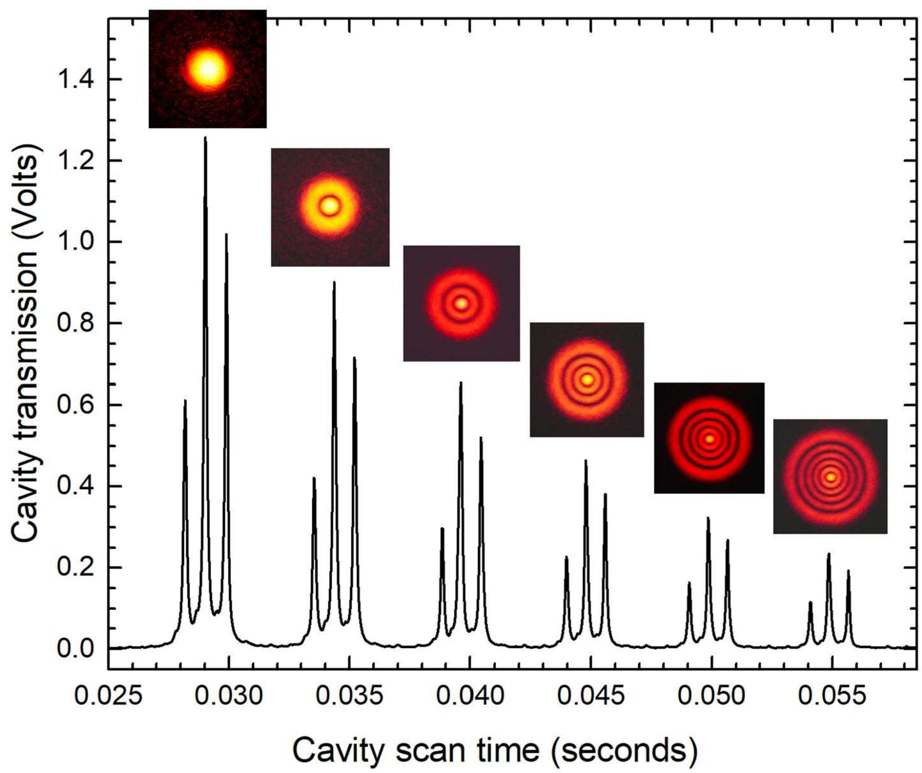

Figure 12. This diagram illustrates how the different triplets in the cavity sweep correspond to different modes in the FPQFA-5 cavity. As you zoom in and scan over just one triplet at a time, you can see which mode is being excited.

Mode matching

The appearance of these higher order modes is related to the topic of "mode matching", which is an important concept whenever optical cavities are used. The basic idea is shown in Figure 13, which is meant to represent the He-Ne laser beam shining on the short FPQFA-5 cavity (with the cavity mirrors not shown in the sketch). If the He-Ne wavefronts curved in such a way that they matched perfectly into the cavity mode wavefronts, then the laser would only excite the 00 cavity mode. But there is a mismatch, so the laser excites other higher-order modes in addition to the 00 mode. Note however, that everything is radially symmetric in the sketch, so only radially symmetric modes will be excited … as you see in Figure 12.

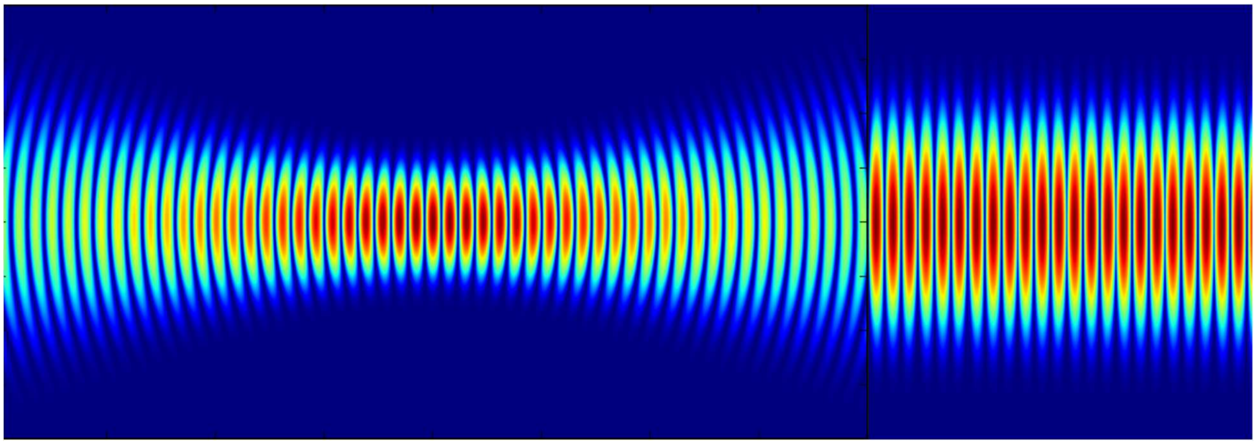

Figure 13. The sketch above represents a He-Ne laser beam (on the right, roughly a series of flat wavefronts) shining on a cavity (on the left, where the 00 mode shown has curved wavefronts). There is clearly a mismatch between these two different wave trains. As a result, the He-Ne beam excites the 00 cavity mode as well as numerous higher order cavity modes, as seen in Figure 12. But if you add a lens (sketch on the right), you can curve the He-Ne wavefront to match the cavity. If done properly, the He-Ne beam will mainly excite the 00 cavity mode, and not the higher-order modes.

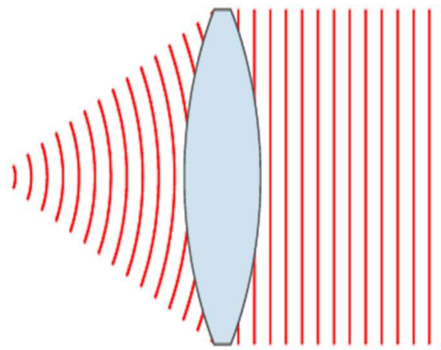

Exactly where the different triplets appear in your cavity sweep on the oscilloscope is complicated. In general, the higher-order cavity modes have different frequencies, all different from the 00 mode and different from each other. This means that the cavity length will be different when these modes hit the He-Ne laser frequency. The He-Ne only produces a triplet of 00 modes, but these couple to different modes of the FPQFA-5 cavity. Thus the triples all occur at different places in the cavity scan. It makes sense if you think about it long enough.

If you are on board so far, then you are ready to take the next step in "mode matching" the He-Ne laser to the FPQFA-5. The hard part is figuring out the right focal-length lens to use, and where to put that lens in the beam. Lucky for you, we have done that part for you, which makes it easier to just see the results in the lab.

Your starting point is a well aligned cavity, which you just set up. Then you simply place a 100mm-focal-length lens (should be on the optical table) about 80-85mm ahead of the front of cavity. This new lens is not shown in Figure 7, but there is room for it.

Before placing this lens or adjusting any mirrors, note where the spot is on the TV, and then position the lens up/down and left/right so the spot on the TV stays where it was, while keeping the lens position at about 80-85mm. Keeping the spot fixed on the TV means the laser beam is not deflected by the lens, so the beam is going through the center of the lens. Next adjust the four knobs like you did before, trying to make the 00 triplet peaks higher. When you get started, tweak the knobs to "push" all the peaks to the left of the 'scope screen. Then tweak more to make other modes as small as possible.

Exercise 8. Keep going with the mode-matching process until you produce a spectrum that looks something like that in Figure 14. Save a .csv file and make a plot for your e-notebook. Congratulations! You have now aligned and mode-matched an optical cavity… a nontrivial introduction to modern laser optics. And now the spectrum is much simpler – 00 modes go in, and these all couple to the cavity 00 modes. So only one triplet remains.

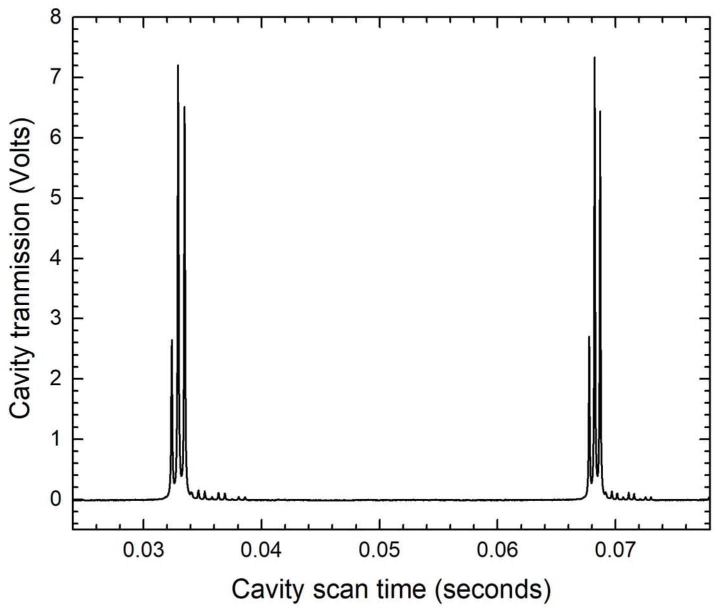

Figure 14. With good mode matching, you can nearly eliminate all but the 00 cavity modes. Now you mainly see just two triplets separated by one cavity FSR.

Using a tunable laser

Single-frequency gas lasers like the Helium-Neon laser have their uses, but a great many scientific and technical applications require tunable lasers, often for various forms of optical spectroscopy. Your next task, therefore, is to start working with the tunable diode laser you see under the plastic case on the optical table. You will see (or have seen) more about how these lasers work in the Diode Lasers experiment in the Optics Track, so we will not repeat that discussion here.

Laser frequency scanning

With the He-Ne laser, you basically had two settings to choose from – on and off. With our tunable diode laser, however, you can adjust the laser power and sweep the laser frequency, which means more control parameters to set. To get you started, Figure 15 shows the electronics layout for the diode-laser controller and Figure 16 shows the front of the laser controller chassis. Assuming you are new at this (although perhaps not, if you already did the Rubidium lab), let's walk through the settings to see what the different controls are for.

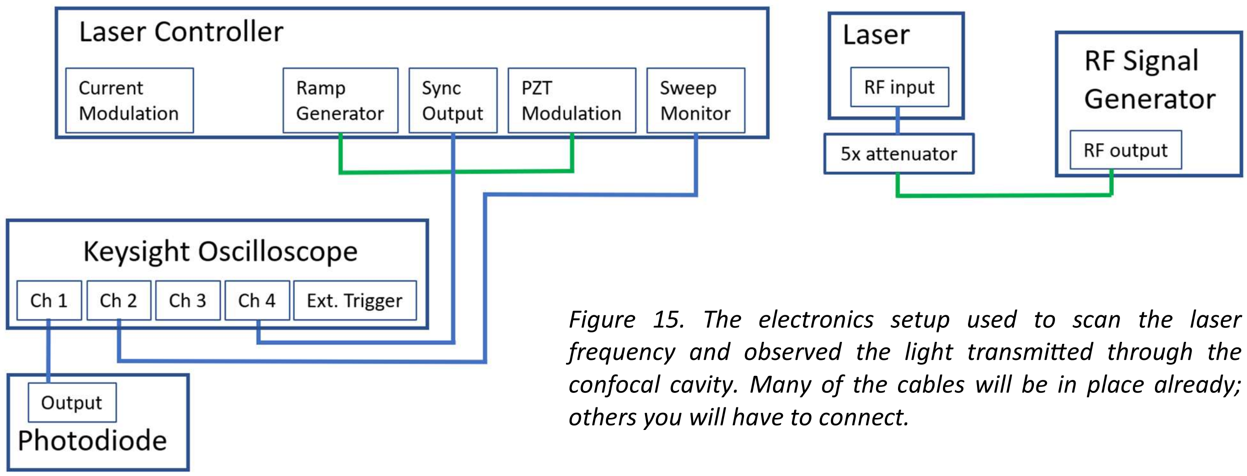

Figure 15. The electronics setup used to scan the laser frequency and observed the light transmitted through the confocal cavity. Many of the cables will be in place already; others you will have to connect.

1) *Current Set* (100mA Full Scale) – this 10-turn potentiometer is used to change the laser current. Any you may have to tweak this knob now and then to get the laser to operate well, as you will see below. Leave the Attenuator knob (below the 10-turn pot) set to maximum.
2) The *Ramp Generator* produces a triangle-wave signal for scanning the laser frequency, and this voltage output goes to the *Piezo Modulation* input, as shown in Figure 15. The piezo (or PZT, short for piezoelectric transducer) moves the external grating right next to the diode laser tube (you can see this small diffraction grating inside the small acrylic laser box; the piezo only moves the grating a few microns). The physics behind all this is described in the *Diode Lasers and LIDAR* handout (part of the Optics Track), but for now just remember that the main purpose of this setup is to change the laser frequency by an amount that is proportional to the applied voltage. With a triangle-wave voltage, the laser frequency then scans back and forth continuously in time. For a reasonable starting point, set the *Ramp Generator Amplitude* knob to 4, the *Piezo Attenuator* knob to 0.4, the *Ramp Generator Offset* to zero, and set the *Piezo Controller Output Offset* to zero.
3) You will often want to monitor the laser-scanning signal using the *Piezo Monitor* output, so send this signal into ch2 of your oscilloscope. Send the *Sync Output* signal into Ch4 on the oscilloscope,

Figure 16. The laser controller front panel.

trigger on ch4, and adjust the 'scope settings so you see a stable triangle-wave signal like you see in Figure 17. Increase the *Piezo Controller Output Offset* to bring the minimum ch2 signal down to around 0V. You may have to use the Normal trigger mode if the triangle-wave signal frequency is too slow for Auto mode.

4) Once you see the Piezo signal, use the Measure feature on the 'scope to measure the signal frequency, and set the *Ramp Generator Frequency* so the 'scope measures 10 Hz. Then get the 'scope set up so you see the red trace in Figure 17, with a single half-cycle of the triangle-wave visible, with a negative slope. A half-cycle is all you need, because the other half-cycle just shows the laser scanning in the opposite direction, which provides no additional information. The negative slope is just a convenient convention. Note the zero-voltage flags (left side of screen, with the small "ground" symbols) and the trigger point (solid triangle on the top-left of the screen). Use these to help you adjust the 'scope settings. This sweep trace has become our standard laser scanning setup, and you can probably leave it unchanged from this point forward. The important knobs you need to remember at this point are:
5) *Piezo DC Offset* – Turn this knob and watch what happens to the sweep trace on the 'scope. You should see a constant offset adjust, as the name suggests. Adjust this to change the center frequency of the laser sweep, but also make sure the *Piezo Monitor* voltage always remains positive.
6) *Ramp Generator Amplitude* – Again, turn this knob and you should see the triangle-wave increase in amplitude. This allows you to adjust the range of the laser frequency scan. The laser frequency (if all goes well) will be proportional to the piezo voltage, which is equal to the ch2 signal on the 'scope.

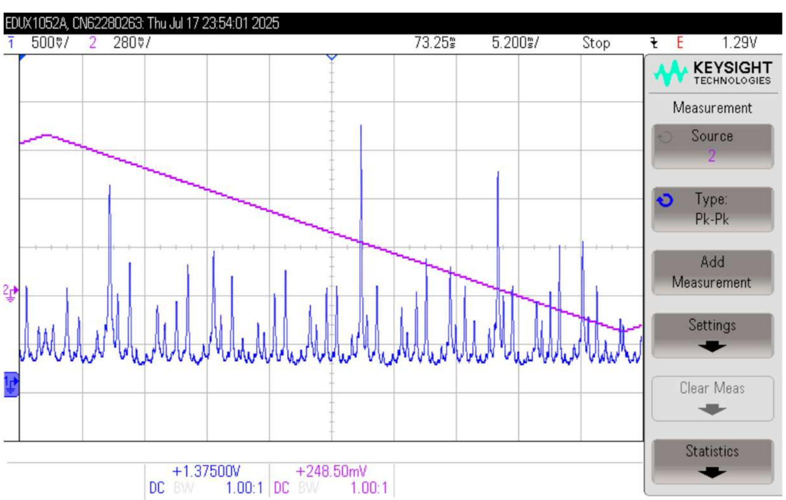

Figure 17. This screenshot shows the Diode Laser sweep signal (red trace) and a typical beginning transmission signal through the short cavity.

Next turn on the laser current (switch on the far left) and set the current knob to be in the 8-9 range. You should see a red laser beam coming out of the plastic case, as this laser operates at a wavelength of 658 nanometers. Once again, you do not need eye protection to use this laser (the warning on the laser controller applies to the invisible 780nm laser we use in the Rubidium Lab, where laser goggles are required).

Before going further, temporarily place a 50mm lens in the beam and have a look at the intensity profile of the expanded beam, like you did with the He-Ne laser earlier. While gas lasers like the He-Ne typically have circular Gaussian beams, diode lasers often produce elliptical Gaussian beams that have different waist sizes along the two principal axes. This is not a desirable feature, but this is what you get from the semiconductor cavity. Again, this is discussed in more detail in the *Diode Lasers and LIDAR* lab in the Optics Track.

With the triangle wave frequency at 10 Hz, measure the Peak-to-Peak amplitude and set it to 1 Vpp. For convenience later, set the Ramp Generator Amplitude knob to 4 and then use the Piezo Controller Attenuator to produce a 1Vpp scan. With this set up, you can zoom in on a spectral feature by turning the RGAmp knob down, and then go back to 4 to get the full 1Vpp scan. Set it up now, and you will see why that is useful below. When done, center the triangle wave on the 'scope, as you did before with the FPQFA-5 cavity.

If all is well, the Laser Controller should now be scanning the laser frequency back and forth with a triangle-wave signal. This signal is made in the Ramp Generator, it gets amplified by the Piezo Controller, and it moves the grating back-and-forth in the laser housing to change the laser frequency. Lots of engineering is needed to make all this happen, but our focus today is on optical resonators, so let us move on to the optical layout.

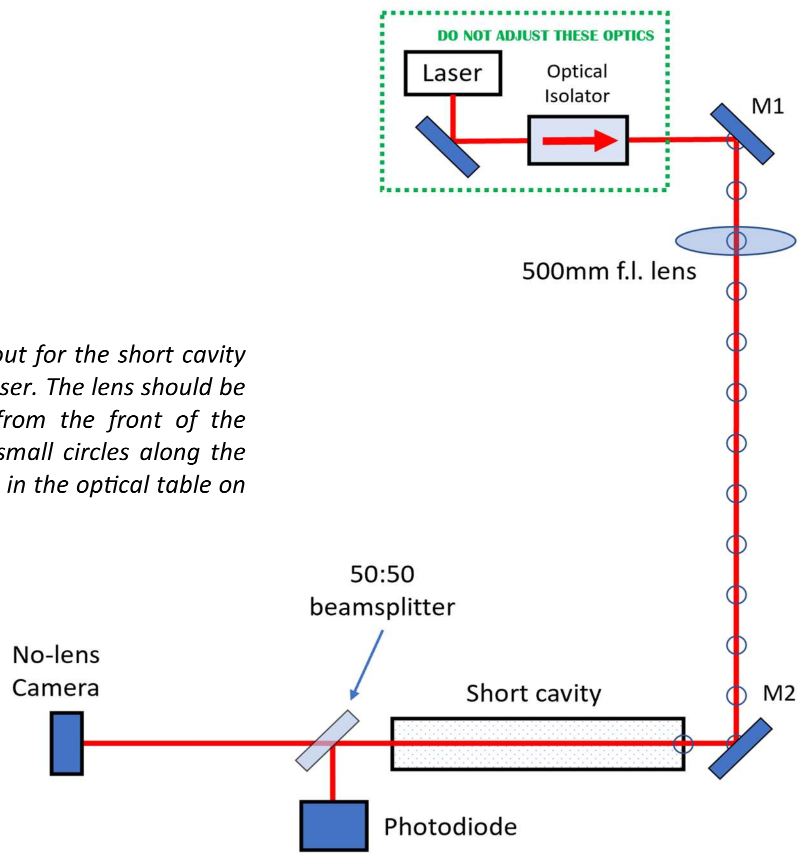

Figure 18. The optical layout for the short cavity with the tunable 658nm laser. The lens should be placed about 11 inches from the front of the cavity; in this sketch the small circles along the beams represent the holes in the optical table on a one-inch grid.

The optical setup

Figure 18 shows a layout that is not too different from what you used with the He-Ne laser. But instead of a fixed-wavelength laser and a tunable cavity, this time we are using a tunable laser and a fixed-length cavity. As before, the optical isolator has been set up for you, and this device allows laser light to pass through in one direction only. And again, the isolator is there to keep light reflecting off the cavity from producing instabilities in the laser operation.

The "short" cavity has a mirror spacing of 9cm, so much longer than the 5mm spacing of the FPQFA-5 cavity. Thus the free-spectral-range of this cavity is $c/2L \approx 1.67$ GHz. The two cavity mirrors both have radii of curvature of 20cm.

As always, for best results, set up the optics on rectangular grid, and remove the 500mm lens when doing your initial alignment. We again want to separate the alignment degrees of freedom, so leave very little space between M2 and the front of the cavity, as shown in Figure 18. And leave about 35cm between M1 and M2, so you have room to place the 500mm lens later. Attach the 658nm filter to the front of the camera tube; this blocks the room lights while passing the 658nm laser light to the camera. As before, there is no lens on the camera.

With the 500mm lens removed, place the short cavity and align M1 and M2 to send the beam through it as the laser frequency scans. Like before, you should use M1 to center the beam on the first cavity mirror, while using M2 to point the beam so it retroreflects. Iterate the mirror alignment steps as needed. Separate those degrees of freedom.

The beginning alignment can be done by eye – see that the beam hits the center of the cavity entrance and view the retroreflected beam on the optical isolator. If that looks good, then you should see a blob of light on the TV. Use the *fixed* 50:50 beamsplitter (not the one that includes adjustment screws), and make sure you use the photodiode with the 658nm filters built in. The transmitted light may be too faint to see by eye using a white card. With the He-Ne setup the laser was bright and the transmitted intensity was high. Not the case now.

You will probably need to turn up the photodiode (PD) gain to maximum (on the back of the photodiode box). Note that the PD sensor is quite large (about 5mm in diameter), but it is not centered in the PD box. The PD sensor is right below the large black dot on top of the PD box. Note also that you cannot see the PD behind the shiny bandpass filter. The PD and beamsplitter alignments do not need great accuracy, so you can just move the components around to maximize the signal you see on the ‘scope.

Observe Laguerre-Gaussian modes

As before, adjust M1 and M2 to bring the blob to a tight ball of light, pulsating as the laser frequency is scanning. Pro tip: When aligning M1 & M2, use one hand to tweak the M1 vertical while the other hand tweaks the M2 vertical. Optimize these settings, then do the same using the two horizontal settings, again in tandem with both hands. Play with the alignment until you get a tight spot of light on the TV with not a lot of "halo" light around it. Optimize as best you can.

Your ‘scope signal at this point should look something like that in Figure 17 – a forest of peaks with no obvious ordering. Most of these peaks are higher-order Laguerre-Gaussian modes of various types. Because the “short” cavity is so much longer than the FPQFA-5 cavity, it is easier to excite lots of different modes. Also, the diode laser beam is *not* a clean circular Gaussian beam like the He-Ne beam, so the laser itself is sending out a host of different L-G modes. The result is a mess of peaks. Try tweaking the alignment to produce a smaller number of higher peaks, as best you can.

You can see which peaks correspond to which L-G modes by zooming in on a single peak and seeing the mode on the TV screen. To do this, set Acquire/TimeRef/Center on the ‘scope and you should see a small downward pointing triangle at the top of the screen. Center this triangle as you see in Figure 17. Next use the Piezo Controller DC Offset on the Diode Laser Controller to center your tallest peak on the triangle. Then you can zoom in on that peak using the Ramp Generator Amplitude knob. (You will be doing this a lot, so make sure this zoom-in/zoom-out feature is fast and easy.) When you zoom in on the tall mode, you will see which mode you have selected on the TV screen. Adjust the PC-DC offset to center different peaks on the ‘scope and see some different modes.

Exercise 9. Take several pictures of the TV screen to sample some low-order and high-order Laguerre-Gaussian modes, and add your collection of pictures to your e-notebook. Figure 5 shows some examples, but clearly there are many more modes not shown. Identify and label the modes with their L-G indices if you can (it’s not always possible with high-order modes). If you misalign M2 a bit, you can observe some quite whacky mode shapes, so play with this a little while and see what you can find.

Mode matching

Your next task is to mode-match the laser to the cavity (as best you can) using the 500mm lens shown in Figure 18. Start with the lens omitted and the cavity aligned, as you were. Then place the lens at the correct position (close to 11 inches in front of the cavity entrance) and align the lens up/down and left/right to see the same nice blob of light on the TV screen. Again, this lens up/down and left/right position guarantees that the light is going through the center of the lens, so the beam is not zig-zagging all over the table. This lens should be placed close to 11 inches in front of the cavity entrance, as this satisfies the mode-matching criterion fairly well.

Once you have a good lens placement, lock the lens down and tweak the M1 & M2 knobs to optimize the overall cavity alignment, as you did before. Your mode-matched cavity sweep spectrum will never be as clean and simple as what you achieved with the FPQFA-5 cavity, but do the best you can. While you are doing your alignment, locate the 00 mode and tweak the alignment so this is one of the tallest modes.

Exercise 10. Figure 19 shows a fairly good mode-matched spectrum, with L-G modes identified for the tallest peaks. Make a diagram like this once you are satisfied with your mode-matching, and add it to your e-notebook.

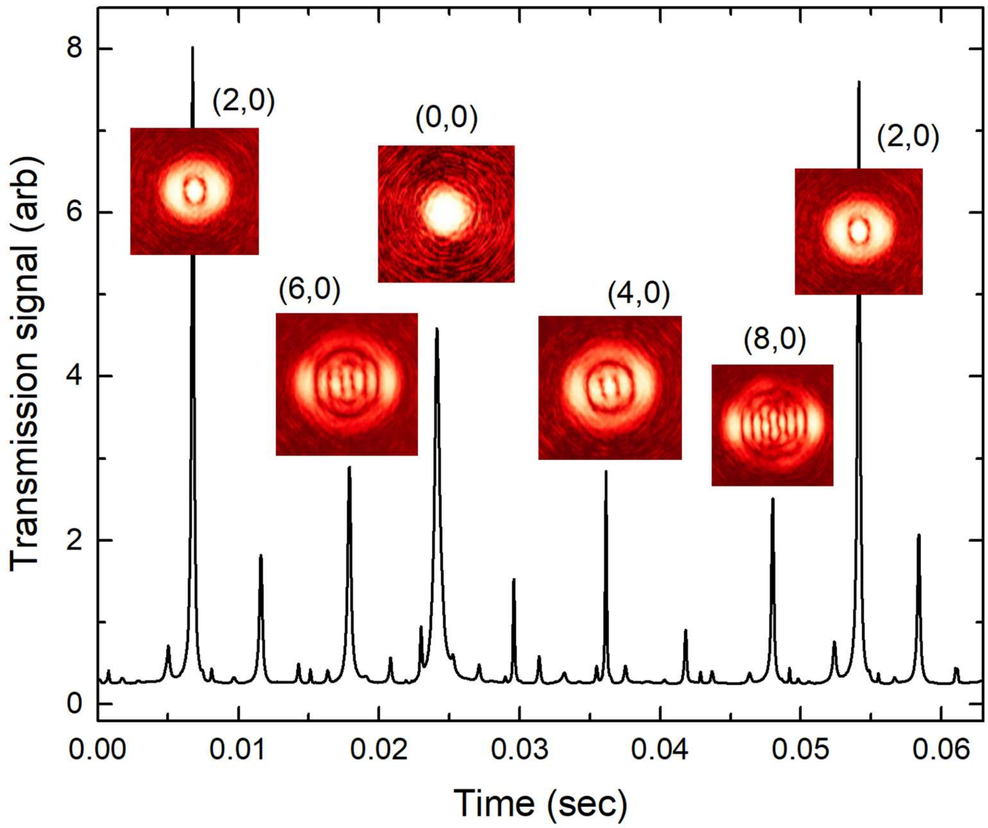

Figure 19. After careful mode matching, the largest peaks in the spectrum are dominated by (2i,0) modes. This reflects the fact that the diode laser beam is elliptical in shape.

Bring in the 00 modes

Now try to isolate the 00 mode peaks as much as possible. Hopefully you already found the 00 mode among your tallest modes in the previous exercise. While scanning at 1Vpp, tweak the alignment more until you make the 00 mode the tallest peak. Here are some additional tricks to try:

1) Try changing the laser current a bit. You started at about 8-9 on the dial, so you might try hunting around in that range. (This usually does not matter much, but sometimes….)

2) Try changing the DC offset to look a bit outside where you are scanning. (Again, this usually does not matter much, but sometimes….)

3) Try turning the scan down to near zero and then slowing searching for the 00 mode using the DC offset. That should show you the 00 mode even if the peak amplitude is low, and it gives you a place to start.

4) Place an OD 0.3 filter in the beam in front of the 500mm lens. This reduces the amount of retroreflected light getting back to the laser, which can make the spectrum less jittery. But it also reduces the light hitting the photodetector, giving you less signal to look at.)

5) Ask for help. Cavity alignment is difficult when you have not done much of it before.

Exercise 11. Try to reproduce a spectrum that looks like the one in Figure 20, showing two 00 modes separated by a cavity FSR. Make sure these are both 00 modes by zooming in on both to view the Gaussian profiles on the TV screen. [Note: when you are scanning at about 1Vpp, then one FSR will fill about half the total scan, as shown in the figure.]

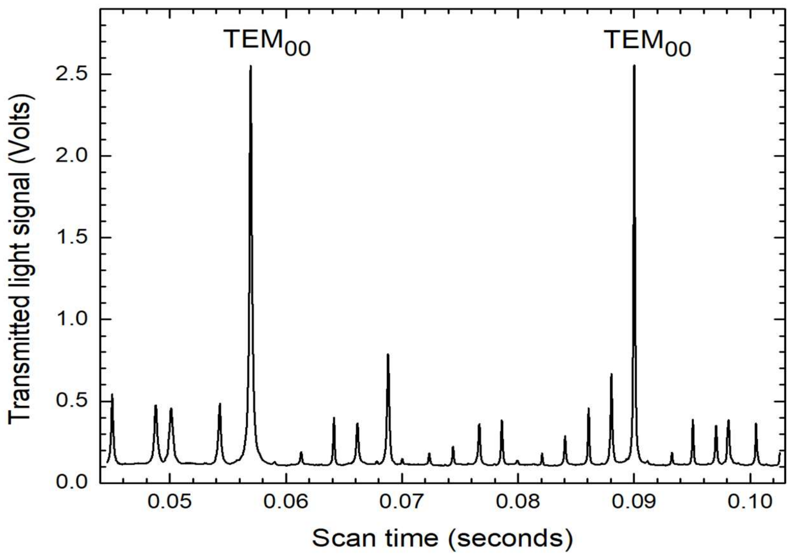

Figure 20. With careful mode-matching, you should be able to produce a spectrum showing two 00 modes separated by one free-spectral-range of the cavity.

Exercise 12. From your spectrum, measure the finesse and spectral resolution of the 9cm cavity. You may assume that the cavity length is exactly 9cm. You may find that the peak widths vary quite a bit with alignment. If that is the case, add that fact to your estimated measurement uncertainties.

In case this is not glaringly obvious by now, the diode-laser setup is much more difficult to align than the He-Ne setup. There are a number of reasons for this:
1) The He-Ne laser has a near-perfect Gaussian beam, which makes it relatively easy to couple to the 00 mode. In contrast, the diode laser's elliptical beam will never couple perfectly efficiently into the 00 mode.
2) Shorter cavities are generally easier to align, giving the 5mm FPQFA-5 cavity a considerable advantage over the 90mm short cavity. As you have seen, however, longer cavities provide higher spectral resolution, which is desirable.
3) Tunable diode lasers are generally more finicky than fixed-frequency gas lasers, both in beam quality and frequency stability. Not surprisingly, lasers with variable frequencies are less stable than lasers with fixed frequencies.

In spite of their many disadvantages, diode lasers are cheaper than any other laser options, and usually by a large margin. Plus, diode lasers are frequency tunable and can be modulated at GHz frequencies. For these reasons, diode lasers have wide applicability in many technological fields, including experimental physics. There are many trade-offs to be considered when choosing the correct laser for a specific task. Our goal in this lab is to give you some exposure to a range of technologies and techniques, and tunable diode lasers are certainly an ever-more-useful technology in professional physics labs.

Part 2. Confocal cavities & FM spectroscopy

For the next phase of the lab, you will be switching to a *confocal cavity* (see below), which is a special cavity geometry that is especially easy to align (yay!), even with relatively long cavity lengths. You will set up the cavity as an optical spectrum analyzer and use it to observe what happens when you modulate the laser frequency. This introduces you to frequency-modulation (FM) spectroscopy, which has countless applications spanning the electromagnetic spectrum from radio frequencies to the optical domain. You will then use these techniques to "lock" a laser to an optical cavity, which again ties into to a broad range of applications in experimental physics.

Confocal cavities

For a typical optical cavity with some R_{roc} (mirror radius-of-curvature) and L (spacing between mirrors) with $R_{roc} > L$, the frequencies of the Laguerre-Gaussian modes are scattered throughout frequency space with no obvious discernible pattern. Of course, the 00 modes are uniformly separated by Δf_{FSR}, but shining a laser beam into a cavity does not excite just these modes. Unless the incoming laser beam has a near-perfect Gaussian profile and is inserted into the cavity with nearly perfect alignment, a host of other L-G modes will be excited as well. However, a special case exists when $R_{roc} = L$, which is known as a *confocal cavity*, illustrated in Figure 21. When this special condition is met, about half the L-G modes are degenerate (in frequency) with the 00 modes, while the other half are similarly degenerate with frequencies that lie halfway between the 00 mode frequencies.

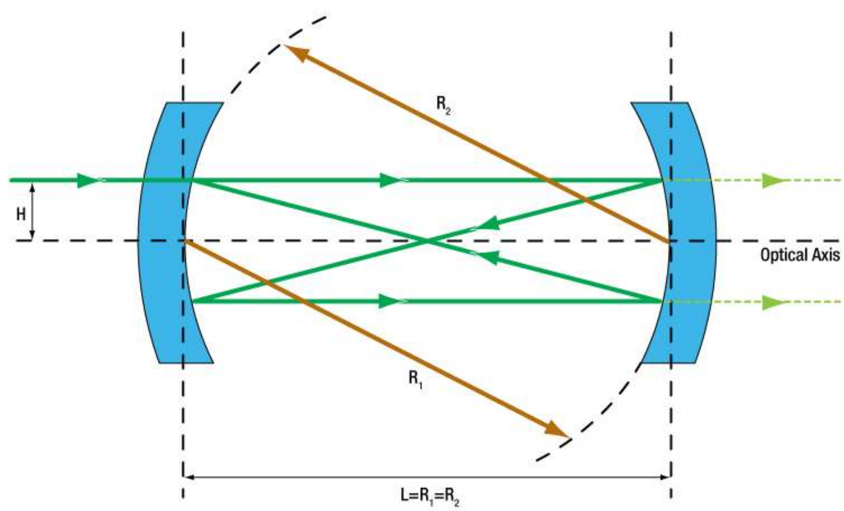

Figure 21. The basic optical layout of a confocal optical resonator made from identical mirrors. The mirrors are separated by the mirror radius of curvature, so a light ray (green) follows a "bow-tie" route inside the resonator. The otherwise-complex frequency spectrum of Laguerre-Gaussian modes magically reduces to a set of equally spaced peaks in this cavity, where each peak contains a large number of degenerate modes.

The remarkable merging of many L-G peaks into two twin peaks means that the transmitted light looks just like the 1D theory in Figure 3 (which you also calculated in Exercise 1), except that the spacing between the peaks is now

$$\Delta f_{FSR,confocal} = \frac{c}{4L} \tag{12}$$

For our cavity, $L \approx 20$ cm is the cavity length, giving $\Delta f_{FSR,confocal} \approx 375$ MHz.

Remarkably, the transmission peaks through a confocal cavity are also quite *insensitive* to alignment, which is very much not the case for non-confocal cavities (as you saw with the short cavity). One does not often talk about Laguerre-Gaussian modes with a confocal cavity, because usually dozens or hundreds or degenerate modes combine to form the observed transmitted peaks.

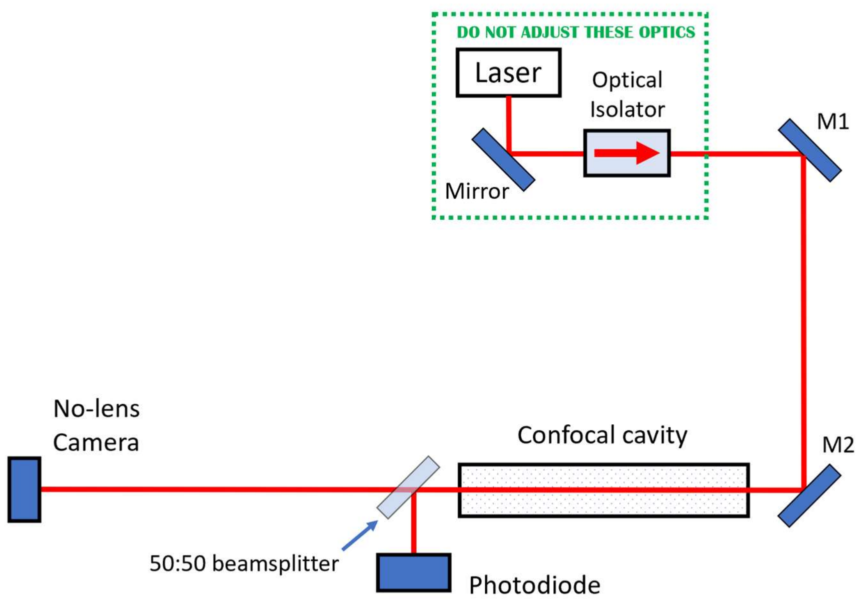

Figure 22. The optical layout for observing transmission peaks through a confocal cavity.

The optical setup

Figure 22 shows the optical setup for the confocal cavity, and you know the drill at this point. Follow the optics rules, separate those degrees of freedom, and follow your plan. As you did with the previous setup, keep about 30cm between M1 and M2 for good dof separation, and keep a short distance between M2 and the cavity entrance. Again use the fixed 50:50 beamsplitter (not the one that includes adjustment screws), and make sure you have 658nm filters on the photodiode (PD), and the camera. Do the same rough alignment you did previously and then tweak the alignment until you get something that looks like Figure 23. Pretty easy, right?

Exercise 13. Produce a screenshot that looks like Figure 23 and add it to your e-notebook.

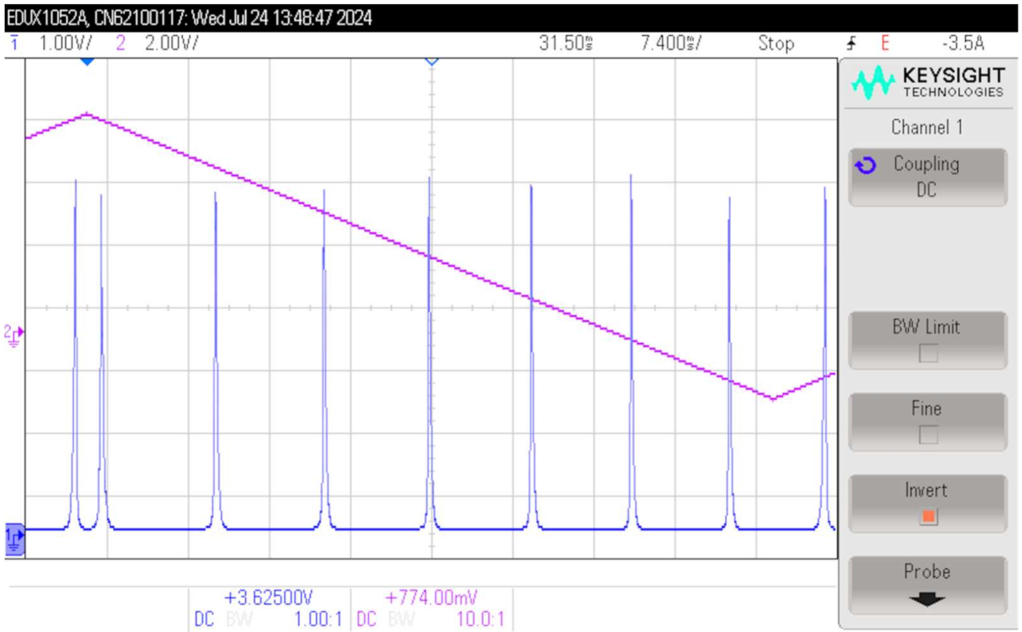

Figure 23. The top trace in the oscilloscope screen shot shows the Piezo Monitor signal as it sweeps the laser frequency. The lower trace shows the cavity transmission signal with several cavity peaks.

When your spectrum is looking good, try adjusting the laser *Current Set* knob to see what happens. You will see that there are some laser-current values where you see additional peaks, indicating that the laser is not running single-mode at those currents. This is normal for tunable diode lasers. We always want the laser to run single-mode, so try to avoid the undesirable laser-current values. If you turn up the sweep (using the *Ramp Generator Amplitude* knob), you may find that the laser usually has trouble running single-mode over a large frequency sweep. Avoid this problem by adjusting the laser current, or by running with a smaller sweep, or both. If all is working well, then turning the laser current up and down a bit should just shift the cavity peaks back and forth nicely. This indicates that the laser is running stably in a single mode at these settings.

Add some sidebands

Now let us look at how a confocal cavity is especially useful as an optical spectrum analyzer by adding some frequency structure to the laser output. The setup for this is included in Figure 15, so turn on the signal generator (Siglent SDG 2122 X) and set the frequency to 40 MHz and the amplitude to 4 Vpp. What this signal does is modulate the laser current slightly at high frequencies, yielding a pair of small "sidebands" on the cavity transmission signal. In a nutshell, the laser that had been operating at a single frequency $f_0 \approx 456$ THz (equivalent to a wavelength of 658 nm) now has a small amount of optical power at $f_0 \pm f_{mod}$, where $f_{mod} = 40$ MHz is the modulation frequency.

If you turn the RF frequency down to lower values, you will see additional peaks in the spectrum, all separated by the RF frequency. If you turn the frequency above 40 MHz, you can see a single set of sidebands with reduced amplitudes at higher frequencies. Consider these to be empirical observations for now; we will discuss the origin of these sidebands later when we will get into the math behind frequency-modulation (FM) spectroscopy.

Exercise 14. Use these sidebands to measure the free-spectral range of the confocal cavity. Turn the RF frequency up to 100 MHz and increase the RF amplitude until you see small second-order sidebands at 200 MHz. Zoom in on these and adjust the RF frequency until the peaks overlap. Use this to estimate the cavity FSR and the cavity length, with measurement uncertainties. [Hint: look for the point of maximum peak overlap, and do this several times to estimate the uncertainty.]

Exercise 15. Reproduce Figure 23 again, but this time with some small 100MHz sidebands (and negligible second-order sidebands). Save a .csv file to USB and add a plot of the digital data to your e-notebook. From this data, measure the cavity finesse, and measure the FSR once again. If you include several FSRs on the ‘screen (like shown in Figure 23), you can easily see the PZT nonlinearity in the digital data. (To avoid ambiguity, define the confocal finesse using $\Delta f_{FWHM} = \Delta f_{FSR}/F$, where $\Delta f_{FSR} = c/4L$.)

Exercise 16. Use the Cursors on the ‘scope to measure df/dV_{PZTmon}, where f is the laser frequency and V_{PZTmon} is the Piezo Monitor signal. (Make sure ch2 probe is set to 1:1 on the ‘scope for this measurement: Press 2, then Probe on the ‘scope to set this ratio.) [Pro tip: set the Cursor mode to Track and set the X1 and X2 Sources to ch2. Then the ‘scope will measure Δ(ch2) when you place the cursors on different peaks.]

Precise confocal alignment

Next turn off the sidebands, turn off the cursors, and reproduce Figure 23 once again. Then turn your attention to the transmitted light seen on the TV. Place the camera about 30 cm behind the second cavity mirror, so the blobs of transmitted light are nice and large on the screen. Turn the scan down until you see four peaks on the ‘scope.

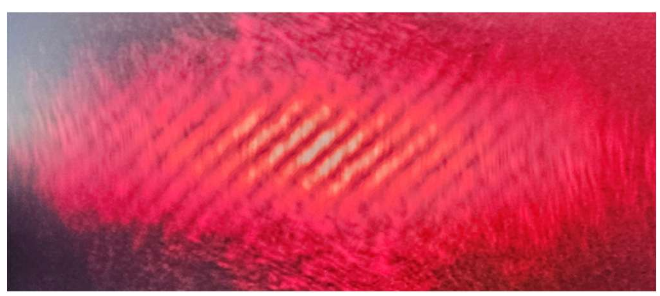

If you align M1 and M2 to overlap the two main ovals of light, you should start to see some interference fringes (bright/dark bands on the ovals). The image on the right shows an example, but your fringes will probably be more closely spaced. If you cannot see any fringes at all, try tweaking M1 and M2 with the same care you used with the 9cm cavity. The fringes can be quite faint, so look carefully.

Once you see some fringes, change the M1/M2 alignment to make the fringes more widely spaced. The trick here is you must change both M1 and M2 to make progress. For example, if you see vertical fringes, then change the M1 left/right knob and the M2 left/right knobs to explore the full 2D parameter space. (Make a small change in M1, then adjust M2 to bring the fringes back. If the fringes are now broader, do it again. If the fringes are narrower, go in the other direction.) The same procedure applies to horizontal fringes, except now you use the two up/down knobs.

Iterate this process, alternating between the two axes, until you make the fringes as broad as you can, with the ovals overlapping. This takes some patience. When you hit the right spot, you will notice that your transmission spectrum no longer consists of four peaks of equal height; instead two peaks will be shorter. Once you see the asymmetry in the peak heights, tweak M1 and M2 to minimize the heights of the shorter peaks. Note that the asymmetry is very sensitive to alignment now, and even a small misalignment will bring you back to having peaks of equal heights.

Exercise 17. Produce a screenshot with your optimal alignment, similar to that shown in Figure 24.

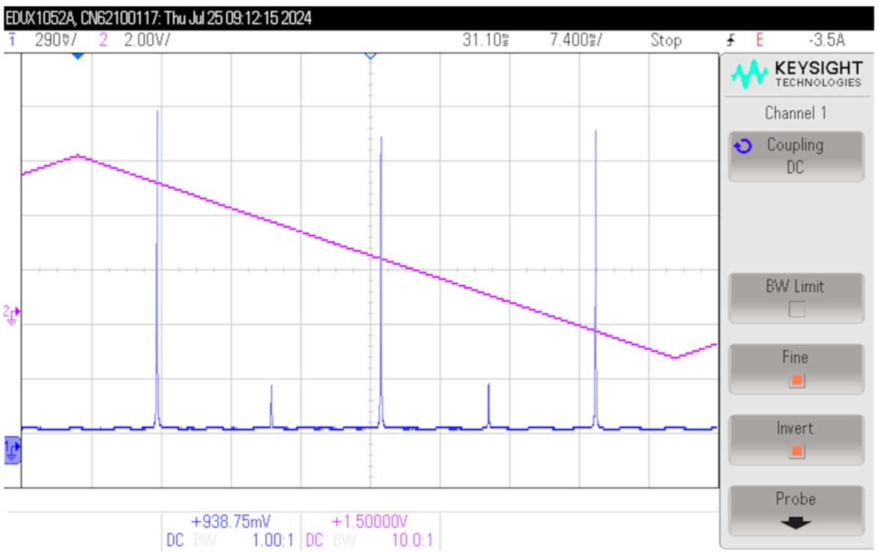

Figure 24. When the input beam to a confocal cavity is carefully aligned, the transmission peaks are no longer of equal height with Δf=c/4L. Instead, the cavity starts to look like a "normal" optical resonator with the primary peaks separated by Δf=c/2L.

One takeaway message from this exercise is that confocal cavities are peculiar instruments. With a slight amount of misalignment, the bow-tie beam geometry shown in Figure 21 produces a set of uniform transmission peaks separated by $\Delta f_{FSR,confoca} = c/4L$. This spectrum is not explained by the plane-parallel model you calculated in Exercise 1. Instead it relies on a remarkable set of frequency degeneracies of the Laguerre-Gaussian modes, where the modes are nicely divided into two classes, giving two transmission peaks equal amplitude (per the normal $\Delta f_{FSR} = c/2L$). There is no simple description of this phenomenon, as the math does not lend itself to easy analysis.

Another take-away is that the confocal cavity is king if you just want an optical spectrum analyzer. The peaks are sharp and the alignment is not difficult. In fact, as you have just seen, a bit of misalignment makes the spectrum easier to interpret. The cavity length must be precisely set to be equal to the radius of curvature of the mirrors to make a confocal cavity, but this is generally easy to achieve in practice.

Sideband spectroscopy

In the radio-frequency domain, there exists a substantial technology built up around amplitude-modulation and frequency-modulation of an electromagnetic carrier wave (which first got into the big leagues with AM and FM radio broadcasting), and these techniques are widely applied in the radio, microwave, and optical domains. And that includes laser modulation. Modulating the injection current to our diode laser is a very simple way to modulate the laser output, both in frequency and amplitude. (Using non-linear crystal modulators is another way to modulate a laser beam.) The basic idea is that one drives the laser with an injection current consisting of a large DC component and a small high-frequency AC component. The current modulation produces both amplitude and frequency modulation of the laser, but we will ignore the smaller amplitude-modulation part for now.

For pure frequency modulation, we can write the electric field of the laser beam at some fixed location as

$$\vec{E}(t) = \vec{E}_0 exp(-i\omega_0 t - i\phi(t)) \tag{13}$$

where $\phi(t)$ is the modulated phase of the laser output. We always assume that $\phi(t)$ is slowly varying compared to the unmodulated phase change $\omega_0 t$, because ω_0 is at optical frequencies and our modulation will be at radio frequencies. If we apply a simple sinusoidal phase modulation we have

$$\phi(t) = \beta \sin(\Omega t) \tag{14}$$

where Ω is the modulation frequency and β is called the *modulation index*. If we note that the instantaneous optical frequency is given by the instantaneous rate-of-change of the total phase, we have

$$\begin{aligned} \omega_{instant} &= \omega_0 + d\phi/dt \\ &= \omega_0 + \beta\Omega\cos(\Omega t) \\ &= \omega_0 + \Delta\omega \cdot \cos(\Omega t) \end{aligned} \tag{15}$$

where $\Delta\omega$ is the maximum frequency excursion. Note that $\beta = \Delta\omega/\Omega$ is the dimensionless ratio of the maximum frequency excursion to frequency modulation rate.

It is useful to expand the above expression for the electric field into a carrier wave and a series of sidebands weighted by Bessel functions, and doing the math gives

$$\begin{aligned} \vec{E}(t) &= \vec{E}_0 \exp[-i\omega_0 t - i\beta\sin(\Omega t)] \\ &= \vec{E}_0 \sum_{n=-\infty}^{n=\infty} J_n(\beta)\exp[-i(\omega_0 + n\Omega)t] \end{aligned} \tag{16}$$

$$= \vec{E}_0 \left\{ J_0(\beta)\exp(-i\omega_0 t) \right.$$
$$\left. + \sum_{n=1}^{n=\infty} J_n(\beta)[\exp[-i(\omega_0 + n\Omega)t] + (-1)^n \exp[-i(\omega_0] \right\}$$

This transformation shows that our original optical sine wave has now developed a comb-like structure in frequency space. The J_0 term at the original frequency ω_0 is the optical "carrier" frequency (in analogy with radio terminology), while the other terms at frequencies $\omega_0 \pm n\Omega$ form sidebands around the carrier. The sideband amplitudes are given by $J_n(\beta)$, which rapidly becomes small in the limit $n > \beta$. (John Hall and Theodor Hänsch received the Nobel prize in 2005 for developing methods that yield extremely high β in the optical domain, making what are now called *optical frequency combs.*)

Note that the total power in the beam is given by

$$\bar{E} \cdot \bar{E}^* = E_0^2 \left[J_0^2(\beta) + 2 \sum_{n=1}^{n=\infty} J_n(\beta)^2 \right] = E_0^2 \tag{17}$$

which is independent of β, as it must be for pure frequency modulation. Often one wishes to add two small sidebands around the carrier, for which one wants $\beta \ll 1$, and the sideband power is then given by $\sim J_1(\beta)^2 \approx \beta^2/4$. Evaluating the above sum and convolving with a Lorentzian laser+cavity spectrum gives an output power

$$I(\omega) = J_0^2(\beta) L(\omega; \omega_0) + \sum_{n=1}^{n=\infty} J_n(\beta)^2 [L(\omega; \omega_0 + n\Omega) + L(\omega; \omega_0 - n\Omega)] \tag{18}$$

where $L(\omega; \omega_0)$ is a normalized Lorentzian function centered at ω_0.

Exercise 18. Evaluate and plot the above optical spectrum, as you might expect to see it using your Fabry-Perot optical spectrum analyzer (remember that a photodiode measures optical power, not electric field amplitude). Plot optical power versus frequency $f = (\omega - \omega_0)/2\pi$. Assume a Lorentzian laser+cavity linewidth of $\Delta f = 5$ MHz. Plot three curves with: 1) $\Omega/2\pi = 100$ MHz, $\beta = 0.5$; 2) $\Omega/2\pi = 30$ MHz, $\beta = 1.5$; and 3) $\Omega/2\pi = 1$ MHz, $\beta = 30$. Note for the last plot you will have to evaluate the sum up to fairly high n, at least to $n > \beta$. For $\beta \gg 1$, note that the spectrum looks much like what you would expect for a time-average of slowly scanning the laser frequency from $\omega_0 - \beta\Omega$ to $\omega_0 + \beta\Omega$. Show your work (for example, include a screenshot of your Mathematica notebook in your e-notebook).

For the next phase of this lab, you will be using the optical layout shown in Figure 25. Begin by assembling the confocal cavity optics like you did previously, except using a 50:50 beamsplitter in place of M2 as shown (this time use the beamsplitter on an adjustable mount). You can ignore the

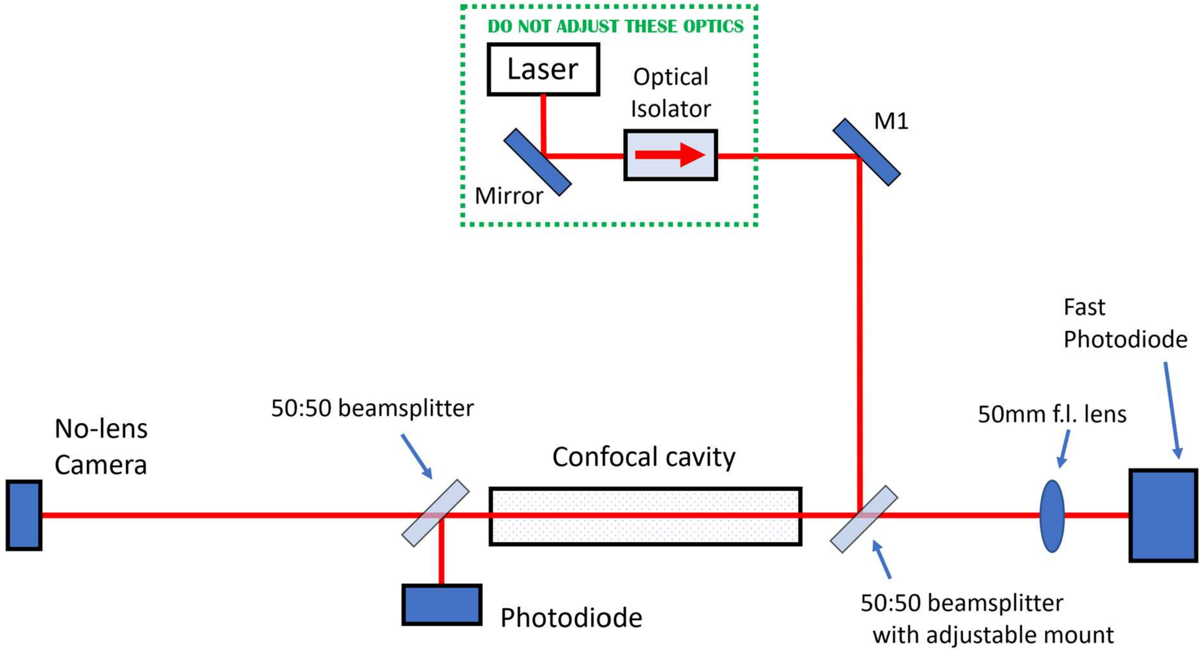

Figure 25. The optical layout for locking a laser to a cavity using the Pound-Drever-Hall method. The fast photodiode is a NewFocus/Newport model 1801 with a 125 MHz bandwidth (compared to about 1 MHz for the slow photodiode).

fast photodiode for now, and the electronics setup is also the same as you used previously in Figure 15. Follow the same steps you did before, and once again produce a 'scope screen that looks like that in Figure 23. Just for completeness, set things up to produce the negative slope in the piezo sweep signal as shown, and make sure ch2 is not set to Invert. Optimize the mirror alignments to produce tall, sharp cavity peaks. Next apply RF modulation to the laser current using the RF signal generator, like you did before.

Exercise 19. Try to obtain sideband patterns that are similar in shape to the three patterns you calculated. You will need to adjust both the RF amplitude and frequency. Zoom in on a single cavity peak to see more detail. When the RF frequency is high, you will see distinct sidebands at the RF frequency, so you can use that to calibrate your spectrum, converting sweep time (recorded in your .csv files) to laser frequency. Moreover, if you calibrate at high RF frequency when you have distinct sidebands, the calibration should remain unchanged as long as you leave the laser sweep parameters unchanged. (Except you can change the Piezo DC Offset without affecting the calibration.)

Document each of your three RF patterns in your e-notebook and note your settings. Compare with your calculations, plotting each as $I(\Delta f)$, the transmitted light as a function of $\Delta f = f - f_{carrier}$. Even better, plot the data and theory on the same graph for each spectrum, thus allowing a direct comparison. Note that the RF frequency in your theory should equal that of the RF signal generator, but the Lorentzian linewidth and strength parameter β can be adjusted, as these change with cavity alignment and other factors.

The Pound-Drever-Hall method

In many precision optical measurements, it is desirable to have a laser with a well-defined frequency. For example, atomic physics experiments often require lasers with frequencies fixed at or near the frequencies of atomic resonance lines. For tunable lasers, it is therefore necessary to have a means of controlling the laser's operating frequency, stabilizing it and "locking" it at a desired value. In this laboratory task you will investigate a popular method for achieving this known as the Pound-Drever-Hall method (named after R. V. Pound, who first used the technique to stabilize microwave oscillators in the 1940's, Caltech professor R. W. P. Drever, who extended the ideas to the optical domain in the early 1980's, and John Hall, who added important pieces of RF technology) [1].

The topic of laser frequency locking is itself embedded in the much larger field of control theory, which is a branch of engineering that deals with methods for controlling physical systems. A thermostat is a classic example, and you can immediately see what is needed in this system: 1) a sensor that measures the temperature (which is what we want to control); 2) an actuator (in this case some kind of heating/cooling system) that can change the temperature; and 3) a controller that reads the sensor and drives the actuator to achieve the desired result - in this case a stable room temperature. Controlling physical systems is a big part of experimental physics, and you may have heard some of the terminology, including things like "servosystems," "servomechanisms," "feedback control systems," "feedback loops," and the like. We go into all this in more detail in the Electronics Track of Ph77.

When working with lasers in physics, one common goal is to lock the laser frequency to a 00 mode of an optical cavity. As you have seen, however, the 00 modes can be quite difficult to work with, so we will focus instead on locking the laser to one of the degenerate modes of a confocal cavity. The cavity transmission seen by the photodiode provides a possible sensor signal, and our locking system will then want to maximize this PD voltage. An immediate problem arises, however, because what should the controller do if the PD voltage drops below its maximum value? The PD voltage alone does not indicate whether the laser frequency should be increased or decreased, because we want to sit at the peak. The solution to this problem is to "dither" the laser frequency back and forth, as this gives a time-dependent signal that now has enough information for the controller to act. Researchers have found that dithering at high frequencies works best, and this is the essence of the Pound-Drever-Hall (PDH) technique.

The best way to begin understanding the PDH method is to see it in action, and for that you need to set up the optics as shown in Figure 25 and the electronics as shown in Figure 26. With the 50mm lens removed for now, leave about 10cm of space between the beamsplitter and the fast photodiode, and then place the lens about 1cm in front of the photodiode sensor for now. The fast photodiode is powered by two 15-volt power supplies to the right of the main electronics chassis, so turn both these on (in either order) and make sure each is set to 15 volts.

Set the RF Signal Generator to a frequency of 20 MHz and an amplitude of 2 Vpp to get started. Note that ch2 on the 'scope should first be used to monitor the sweep voltage while you set up the cavity and get a nice set of transmission peaks. Once these peaks look good, you can change ch2 so it looks at the low-pass-filter output as shown.

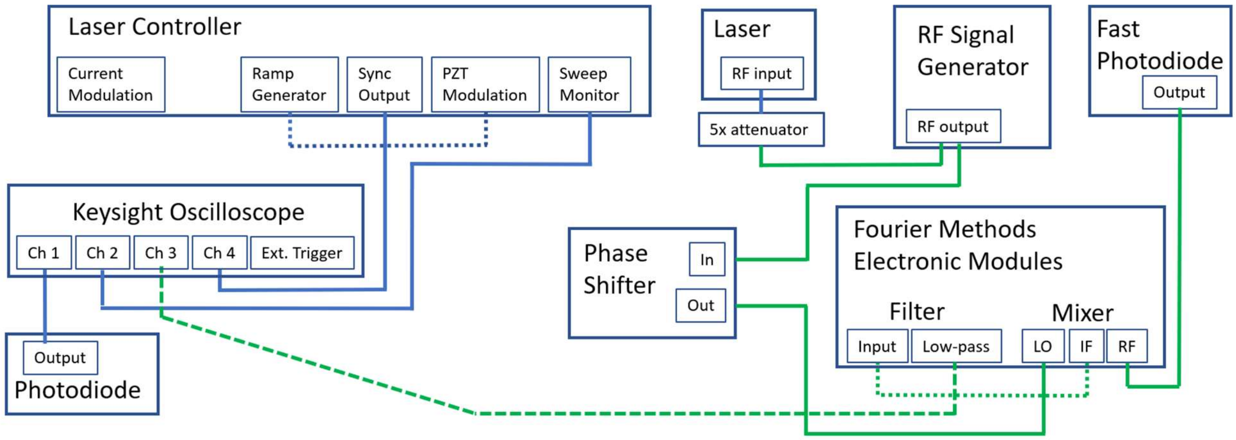

Figure 26. The connection diagram for observing the PDH error signal. Several of these connections may already be in place and need not be removed. Note that two signals go into ch2 on the 'scope at different times. Start by sending the Sweep Monitor to ch2; later you will remove this and plug in the Low-pass filter instead.

One key element in this electronics setup is the *mixer* (specifically a doubly balanced mixer), which is a passive device consisting of a set of transformers and diodes. Diodes are inherently nonlinear devices, and nonlinearity is what gives the mixer its superpowers (but nonlinear physics can be super complicated, so we don't teach much about it in undergrad classes; this is your chance to dip your toes into those waters).

In a nutshell (skipping over the actual electronic schematic), the mixer takes two inputs – the Local Oscillator (LO) and the Radio-Frequency signal (RF) – and the device essentially multiplies them together. The LO and RF ultimately derive from the same signal generator, so they have the same frequency, but they typically have different amplitudes and phases. The output (IF) is then proportional to

$$\mathrm{IF} \sim \cos(\varOmega t)\cos(\varOmega t + \eta) = \frac{1}{2}\cos(\eta) + \frac{1}{2}\cos(2\varOmega t + \eta) \tag{19}$$

which has terms at DC and at 2Ω. The mixer IF output is sent to a low-pass filter to remove the high-frequency term, and the smoothed output becomes an *error signal* in our locking system.

To calculate the error signal, start with the cavity reflection amplitude you calculated above:

$$E_r = \frac{(1 - e^{i\delta})\sqrt{R}}{1 - R\, e^{i\delta}} E_{in} \equiv A[\delta, R] E_{in} \tag{20}$$

where R is the mirror reflectivity and $\delta = 2\pi \varDelta\omega / \varDelta\omega_{FSR}$. A plot of the reflected intensity $|E_r|^2$ shows dips corresponding to the peaks in cavity transmission (because the incoming laser light has to go somewhere, so it is either transmitted or reflected from the cavity).

For $\beta \ll 1$ we can write the input electric field as a carrier and two weak sidebands

$$E_{in} = E_0 \{ -M e^{i(\omega - \varOmega)t} + e^{i\omega t} + M e^{i(\omega + \varOmega)t} \} \tag{21}$$

which gives the reflected amplitude

$$E_r = E_0 e^{i\omega t}\{-MA_- e^{-i\Omega t} + A_0 + MA_+ e^{i\Omega t}\} \tag{22}$$

where

$$A_0 = A\left[\frac{2\pi\Delta\omega}{\Delta\omega_{FSR}}, R\right] \tag{23}$$

$$A_\pm = A\left[\frac{2\pi(\Delta\omega \pm \Omega)}{\Delta\omega_{FSR}}, R\right]$$

The reflected signal at the photodiode is then $I_{phot} \propto |E_r|^2$, which you can verify is a real function containing DC terms and terms proportional to $\cos(\Omega t)$ and $\sin(\Omega t)$. Because I_{phot} is a real function, the effect of the mixer is essentially to multiply the photodiode signal by $\exp(i\Omega t + \zeta)$ and take the real part, where ζ is a constant phase factor depending on the relative phase of the photodiode and local oscillator signals at the mixer. Putting all this together, after low-pass filtering the IF signal, the error signal becomes

$$\varepsilon \propto Re\{e^{i\zeta}(A_0^* A_+ - A_0 A_-^*)\} \tag{24}$$

Exercise 20. Calculate and plot $\varepsilon(\Delta f)$, where $\Delta f = f_{laser} - f_{cavity}$ in MHz, using a variety of values for the phase factor ζ, assuming a sideband frequency $\Omega/2\pi = 25$ MHz, $\Delta f_{FSR} = 375$ MHz, and a cavity reflectivity of $R = 0.995$. Adjust ζ to produce a plot that has an overall shape like the experimental error signal in Figure 27. The change ζ again to make a set of plots showing the different signal shapes. In the lab, you will generate $\varepsilon(\Delta f)$ as the laser frequency is slowly scanned around a cavity resonance and display it on the oscilloscope. If all goes well, the resulting error signal should match your calculations.

Note that the error signal equals zero when the laser is locked to the cavity mode (at $\Delta f = f_{laser} - f_{cavity} = 0$), and it becomes positive/negative if the laser frequency is too high/low. These are the properties you want in an error signal, because then the controller can use this signal to adjust the laser frequency higher or lower as needed, so that it remains locked at $\Delta f = 0$.

As you can see, the math behind laser FM spectroscopy is nontrivial, especially when you consider the interaction of a frequency-modulated laser with a resonant cavity. It takes a while for all this to sink in, so consider this a brief introduction. We present the material here because the PDH technique is a staple of laser optics, appearing in a broad range of applications where one wants to precisely control laser frequencies. If your career direction takes you into laser physics or related technologies, you will likely become quite familiar with laser FM spectroscopy.

Observing the error signal

Once you have everything set up as shown in Figure 25 and Figure 26, you can begin to see some signals like those shown in Figure 27. There is a lot to unpack here, so you should not be surprised

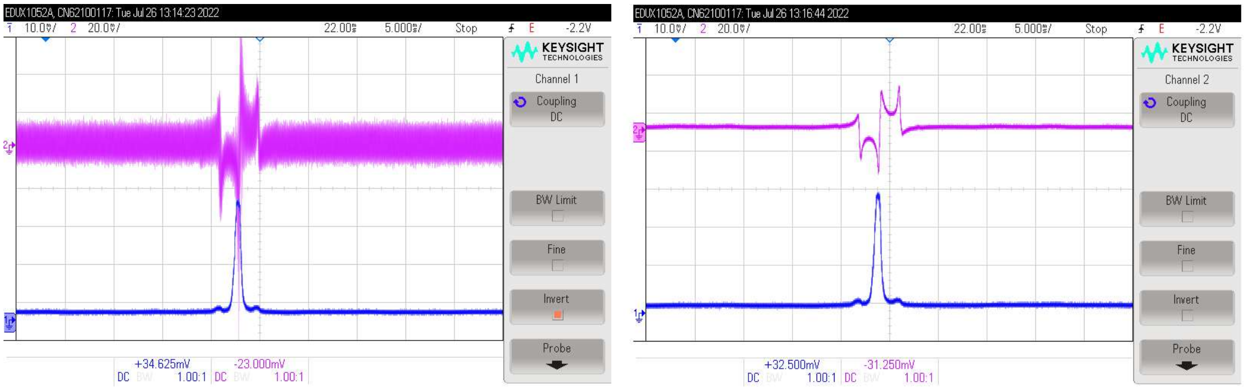

Figure 27. Oscilloscope screenshots showing the PDH error signal, both before (left) and after (right) passing the mixer output through a low-pass filter.

if this does not immediately work perfectly! The first thing to look for is sidebands on the transmitted signal. With the Signal Generator at 20 MHz and 2 Vpp, the sidebands might be small, but they should be visible.

If sidebands are present but the error signal is absent, try moving the 50mm lens around in front of the fast photodiode. This lens likes to sit about 1-4 cm in front of the PD sensor, perhaps with the light not so bright that it saturates the sensor. Once a tiny signal is visible on the ’scope, move the lens around in all three dimensions to optimize the signal. Ask your TA for help if needed.

Once you have a solid error signal on the ’scope, zoom in so you get a better look at its structure. Optimize the size of the error signal and it should have a peak-to-peak amplitude in the 50-100 mV range. Optimizing the lens position (in all three dimensions) is most important for obtaining a strong error signal, but you should also look at tweaking the laser current, the RF frequency, and maybe even the cavity mirror alignment.

Exercise 21. Document some error-signal screenshots in your e-notebook. Try varying the phase factor ζ using the phase shifter in Figure 26. (This device consists simply of short lengths of BNC cables that add time delays to the signal.) Try to connect your observed error signals with those you calculated in the previous exercise.

Adjust ζ so the error signal is mostly positive on the right side and mostly negative on the left side, as illustrated in Figure 27.

Laser locking

When you have a tunable laser, the first thing you want to do is scan the laser frequency and look at the resulting spectroscopic signals, as you have been doing. But it soon becomes desirable to “lock” the laser at some specific frequency, for example a cavity resonance or an atomic transition. This kind of precise laser frequency control is the main application of the PDH technique.

To see this in action, the next step is to add the Servo Controller that closes the loop and locks the laser. We use a commercial controller for this purpose (mostly designed for optical systems; its power switch is on the back) with the setup shown in Figure 28. If you previously set up the ’scope to view a downward sweep voltage as shown in Figure 23, and if your error signal has a positive slope on resonance as shown in Figure 27, then you should connect the low-pass output to input A on the controller, as shown in Figure 28. But if your optimal error signal has a negative slope, no problem, just connect the low-pass output to input -B on the controller. This sign is important, of course, because the controller has to have the correct output sign to lock the laser.

With the Error Monitor now plugged into ch2 of the 'scope as shown, the error signal should look much like it did before on the 'scope, and it should have a positive slope at $\Delta f = 0$. But you will need to adjust the controller's Input Offset carefully so that the baseline signal is quite close to 0V. The controller's job is to make the error signal go to zero, and this works best when the baseline voltage is also at zero volts. Set the P-I corner to INT and turn the Gain up to 7 on the dial.

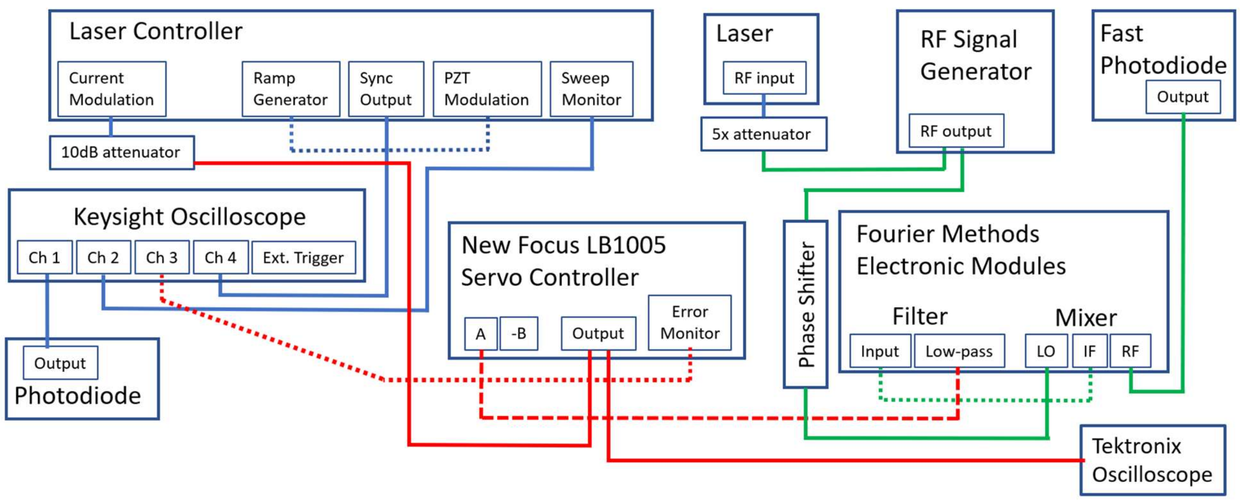

Figure 28. The full electronics diagram for locking the laser to the confocal cavity using the PDH method.

Now one last checklist before you attempt a lock:

1) With the laser doing a fairly short sweep (about like that shown in Figure 27), tweak the piezo DC offset up and down a bit. This should just make the cavity peak and error signal both rock back and forth on the 'scope. If these signals get distorted during this process, then the laser may be near a mode hop, so tweak the laser current until both signals are nice and stable.
2) Next tweak the laser current up and down a bit, and again you should see the cavity peak and error signal both rock back and forth on the 'scope, with no mode hops or odd distortions.
3) On the Servo Controller, make sure the Sweep IN port is not connected to anything, and turn the Sweep Span to maximum. Also, set the Current Attenuator knob to maximum on the Laser Controller. Then turn the Center knob on the Servo Controller back and forth. This sends a DC voltage to the Servo Controller output port, and this goes into the laser current modulation input on the Laser Controller. So this test is just like the previous one; it moves the laser current up and down. And again the cavity peak and error signal should both rock back and forth. (Sign check; if your error signal has a positive slope on the 'scope, then turning the Center knob up should make the peak move to the right.) This servo output also goes to the Tektronix oscilloscope, so you can see this voltage go up and down there as well. If all looks good, turn the

Center knob until the servo output is at zero volts. [If this is somewhat confusing at this point, that is just because there is a lot going on. Study Figure 28 a bit more, go through the checklist from the top, and talk it over with your lab partner and/or TA. Servo controls tend to be complicated beasts, and this one is no exception.]

4) One last time, make sure the background error signal is set to 0V using the Servo Controller Input Offset.

If all this looks good, try to do a "human" lock before engaging the controller. Do this by turning the Ramp Amplitude down to zero (so the laser is not sweeping) and use the Piezo DC Offset to find the cavity peak. If you move the Piezo DC offset as you watch the 'scope, you can keep the laser near the cavity resonance, albeit not very effectively.

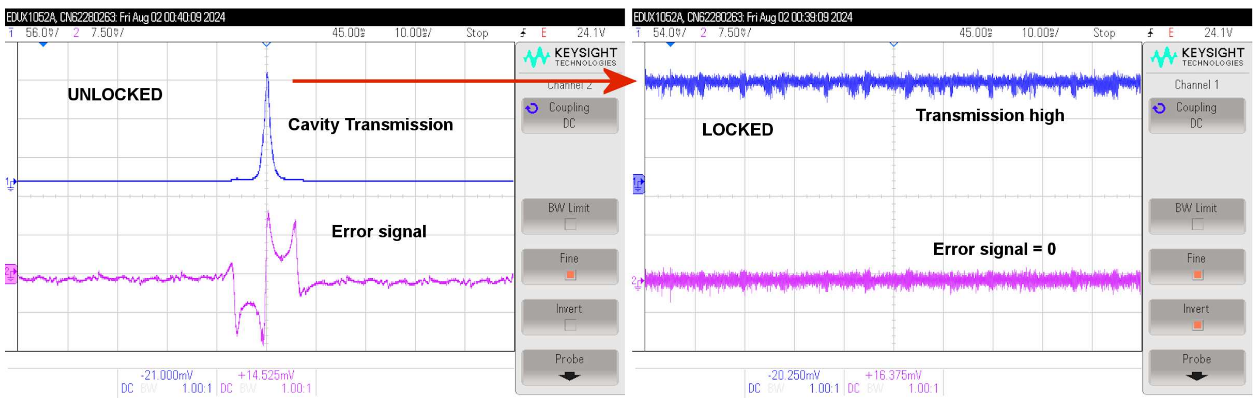

*Figure 29. (Left) The cavity transmission signal and the corresponding error signal while the laser is sweeping. (Right) The same signals with the laser locked. Note that **the cavity transmission remains high** when the laser is locked to a cavity resonance peak, because the servo controller keeps the error signal at zero.*

All this double-checking is to give the controller a fighting chance at acquiring lock. Just connecting the various elements and hitting the lock switch immediately rarely works. So, if all looks good, use the human lock to set the laser near the correct frequency (with the laser not scanning) and then turn the controller switch from Lock Off to Lock On. You should see the error signal go to 0V while the transmitted light signal stays at its maximum value. If the signals are both oscillating like mad, turn the Gain down on the Servo Controller. If the cavity transmission stays at zero with the lock on, go back to the human lock, with the laser not scanning, and again try the Lock On/Off. If still no luck, go through the checklist once again. If still no luck, ask your TA for help.

If the lock works, move the traces around on the scope so you can see both clearly in both the locked and unlocked states. Take screenshots of both for your e-notebook. If you try tapping on the table, you may find that the system falls out of lock pretty easily. This is normal. If you can give the table a good tap with your fingernail and it stays locked, then well done!

Next lock the laser and slowly turn the Servo Gain down. You will see the lock become "softer" with lower gain until eventually the laser will fall out of lock and will no longer lock at all. Turn the gain back up to get it locking again and then try turning the gain up slowly. Soon the error signal will begin to oscillate madly as the servo overcompensates to the point that is driving the system with positive feedback instead of the desired negative feedback. This behavior is not unlike the feedback you hear in audio systems when you place the microphone in front of the speaker.

Finally, watch the Servo Controller output on the Tektronix 'scope when you slowly adjust the laser current knob - just a tiny amount so you do not lose lock. Changing the laser current would normally change the laser frequency, but the servo senses this and applies an equal and opposite signal to the laser current modulation input, so the actual laser current does not change as long as the laser stays locked. Thus the servo is doing its job – you (the environment) try to make the laser run at some frequency that is not equal to the cavity resonance, and the servo reacts to not let that happen. If you look at your 'scope signals and put in the numbers, you will see that the locked laser frequency is now stable to about one part per billion!

Exercise 22. Document your laser locking setup and signals (above) in your e-notebook.

References

[1] Eric D. Black, An introduction to Pound–Drever–Hall laser frequency stabilization, Am. J. Phys. 69, 79–87 (2001).