

Ph 3 - INTRODUCTORY PHYSICS LABORATORY

– California Institute of Technology –

Laser Interferometry

1 Introduction

The purpose of this lab is for you to learn about the physics of laser interferometry and the science of precision measurements, and for you to gain some laboratory experience with optics and electronic test equipment along the way. You will build and align a small laser interferometer in the Ph3 lab, you will determine how accurately you can measure distance using this interferometer, and then you will use the interferometer to examine the small motions of a oscillating mechanical system.

1.1 A Brief Laser Primer

First, a few words about lasers and how they work. The word LASER is an acronym for Light Amplification by Stimulated Emission of Radiation, and there are many types of lasers in widespread use. Solid-state lasers include the ubiquitous laser pointer, powerful YAG (Yttrium-Aluminum Garnet) lasers, precision Ti:Sapphire (Titanium-doped Sapphire) lasers, etc. Gas lasers include the red Helium-Neon (He-Ne) laser, powerful blue/green Argon lasers, super-powerful industrial CO₂ lasers, etc. Laser physics is a huge and fascinating field, and there is much active research in this area today. Of course we cannot discuss all of laser physics here, but we can look at the few essential features that are common in all lasers.

Figure 1 shows a schematic diagram of a laser. Note that all lasers must include: 1) an amplifying medium that can emit light; 2) an external energy source to excite the amplifying medium; and 3) an optical cavity to confine and guide the laser light. In the case of the He-Ne laser you will be using in the lab, the amplifying medium is a mixture of helium and neon gases, held in a glass tube. The energy input is an electrical discharge in the tube, and the laser cavity is formed by mirrors on the ends of the tube. Light bounces back and forth between the two mirrors, forming a standing wave between them. The light intensity in this standing wave is amplified until it reaches some steady state,

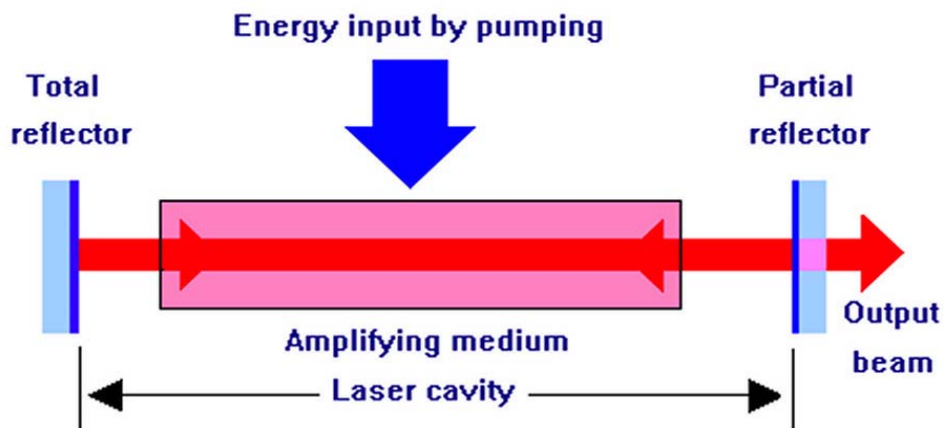


Figure 1. The basic features of a laser. [Image from <http://www.cartage.org.lb/en/themes/sciences/physics/optics/LaserTutorial/Laseroscillator/Laseroscillator.htm>.]

at which point a beam of laser light emerges from the partially reflecting output mirror.

It is useful to consider what happens when you first turn on the He-Ne laser in the lab, keeping Figure 1 in mind. Flipping on the power switch applies a high voltage (about 10 kV initially) to two electrodes inside the laser tube, one on each end of the tube. This causes an arc discharge, which just means that electrons flow through the gas, which is at fairly low pressure inside the tube. The flowing electrons collide with the gas atoms, sending them to higher energy levels. The details of what goes on inside the He-Ne laser tube are a bit subtle, involving the various atomic energy levels of the helium and neon atoms. In a nutshell, the electron collisions produce excited helium atoms, and these in turn produce excited neon atoms via helium-neon collisions. The neon atoms are excited to a variety of upper energy levels, and some of these decay to a particular *metastable* excited state, which does not quickly decay back down to the neon ground state. This population of metastable neon atoms is the amplifying medium.

If the density of metastable neon atoms is low, then eventually the atoms undergo ordinary spontaneous decay, each atom emitting a 633-nm (red) photon during its decay. The photons spray all over, and this red glow discharge makes what is essentially a neon lamp, not yet a laser. Some of the emitted light bounces back and forth between the two (highly reflecting) mirrors, and this light is trapped for a little while in a standing wave inside the tube. If we keep increasing the density of metastable neon atoms, more and more light is emitted in all directions, and again some fraction of this light becomes trapped in the standing wave.

When the neon metastable density reaches a certain point, called the laser *threshold*, something interesting and very quantum-mechanical happens. The presence of the standing wave light causes the atoms to decay faster, through a quantum process called *stimulated emission*. Importantly, the standing-wave light stimulates the atoms to emit right into the same standing wave, and not in a random direction. This results in a positive feedback effect: stimulated emission causes atoms to emit into the standing wave, and this increases the standing wave intensity, which in turn cause more stimulated emission. Once the laser is above threshold, almost all the metastable neon atoms emit via stimulated emission, so nearly all the light inside the cavity gets injected into the standing wave. Some of this light leaks out of the output mirror (which is highly reflecting, but not totally reflecting), producing what we see as a laser beam. Lasers are fascinating objects, and understanding them in detail takes many years of study. Laser physics has been a career for many scientists and engineers over the years, and it is still an active area of research and development.

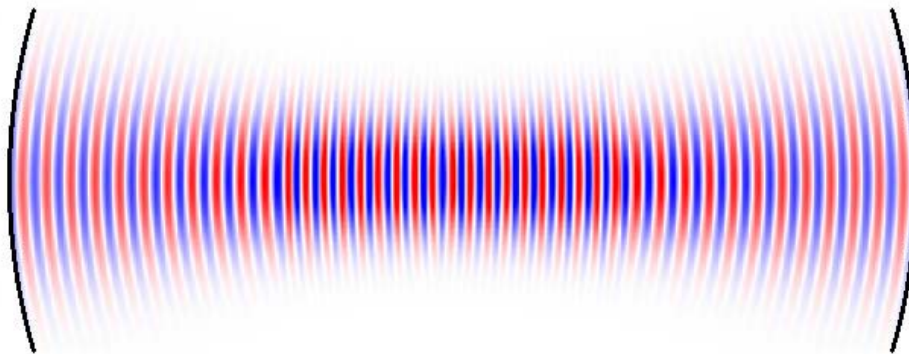


Figure 2. A schematic picture of a (very short) laser cavity, showing two curved mirrors with a standing wave between them. Note that the wavefronts are not flat, and the resulting output laser beam is necessarily diverging. [Image adapted from http://www.rp-photonics.com/resonator_modes.html.]

1.2 Laser Optics

A somewhat lesser-known feature of lasers is that the optical cavity is made from two curved mirrors, shown schematically in Figure 2. This picture represents a very short cavity, only some tens of wavelengths long. The He-Ne cavity you will be using is roughly 20 cm long, so of order $20\text{cm}/633\text{nm} \approx 300,000$ wavelengths long. An important

feature of these curved-mirror cavities is that the light intensity has a Gaussian profile, so $I(r) = A \exp(-2r^2/w^2)$, where r is the distance perpendicular to the propagation direction and w is a constant. This Gaussian intensity profile is not intuitively obvious, and it took people quite a while to work out all the math. Curved-mirror cavities like this have many wonderful properties that make lasers work (for example that the standing waves are not terribly sensitive to minor mirror misalignments, unlike flat mirrors), and the details make up a field known as *Gaussian optics*.

For our He-Ne laser, the beam divergence is small and the wavefronts are nearly planar, so the laser beam behaves much like a simple plane wave (except with the aforementioned Gaussian intensity profile), and we are okay using a plane-wave approximation. We will therefore not mention Gaussian optics much further, but you should at least know that it exists.

1.3 Laser Interferometry

In this experiment you will be building a basic Michelson interferometer (named after Albert Michelson, who invented it in the late nineteenth century). While optical interferometry can be done using broad-band light sources (as Michelson did long before lasers were invented), it is much easier with a nearly monochromatic light source like a laser. You will be using the optical layout shown in Figure 3.

The basic idea is fairly simple: the initial laser beam is split into two equal-intensity beams by the beam splitter, these beams both reflect off their respective end mirrors, and the beams recombine once again at the beamsplitter. How much light ends up going into the detector depends on the relative phase of the recombining beams. If the beams are in phase, then we say they combine *constructively*, and in this case the beams add and the full initial laser beam intensity lands on the detector. But if the beams are out of phase, then they combine *destructively*, so no light lands on the detector. If we translate one of the end mirrors as shown in Figure 3, then small changes in the mirror position can be measured by the changing light intensity on the detector. Laser interferometry is a sensitive way of measuring small displacements, and it is a non-contact measurement that can operate over quite large distances.

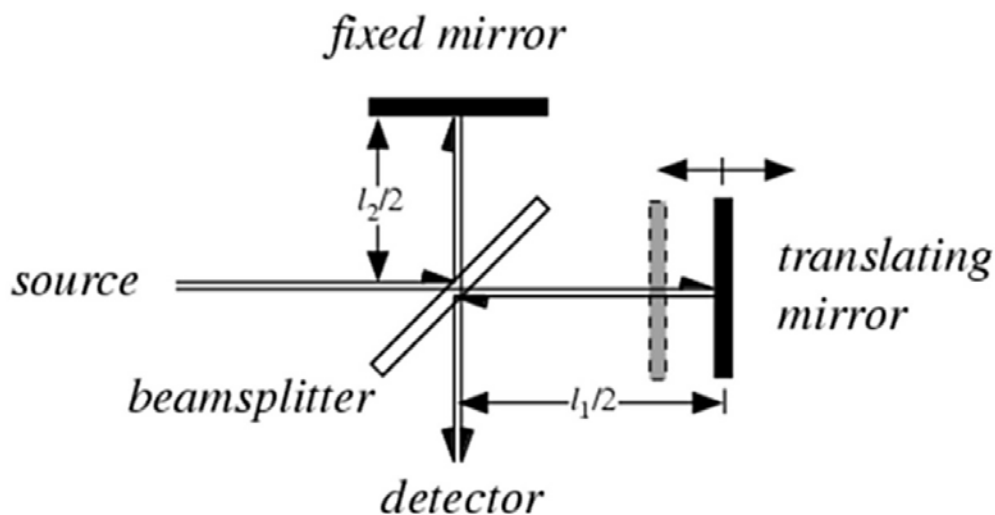


Figure 3. The optical layout of a basic Michelson interferometer. Light from the source (a He-Ne laser in our case) first strikes a 50:50 beamsplitter that splits the light into two equal-intensity beams. These beams travel down the arms of the interferometer, where they are then reflected back by two mirrors. The beams recombine back at the beamsplitter, so some of the light is directed toward the detector while some is reflected toward the source. The intensities of these two beams (one going toward the detector and the other shining back toward the laser) depend on the path length difference $l_1 - l_2$ of the two arms of the interferometer. (Image from <http://scienceworld.wolfram.com/physics/FourierTransformSpectrometer.html>.)

To understand how the interferometer works in detail, we have to look at the math. If we approximate our

Gaussian laser beams as plane waves, then we can write the electric field of a propagating light wave as

$$E(z, t) = E_0 e^{i(\omega t - kz)}$$

where E_0 is a constant, $\omega = 2\pi c/\lambda$ is the angular frequency of the light, $k = 2\pi/\lambda$ is the wave number, and λ is the light wavelength. Here we have defined our coordinates so that the laser beam is propagating in the z direction. For $\lambda = 633$ nm, the oscillation frequency of the electric field is $\nu = \omega/2\pi = c/\lambda \approx 4.7 \times 10^{14}$ Hz.

Since we have no electronics that can operate at 470 THz, we do not measure any electric fields when we are working with visible light. Instead we measure the power in a beam, which is proportional to the square of the electric field averaged over time, and we write this as $I_0 = CE_0^2$ where C is a proportionality constant. For the laser you will be using, the initial power of the laser beam is about $I_0 = 2$ milliWatts, which is about what you get from a weak laser pointer.

A brief word about laser safety. A good rule-of-thumb to remember is that staring directly at the sun puts about 1 mW into each of your eyes, so a 2 mW laser beam is enough to cause eye damage, at least in principle. But this only happens if you stare directly into the beam for at least a second or two, and this is all but impossible to do by accident. Thus a 2 mW visible laser (like a laser pointer) is quite safe to work with.

Now a key point in interferometry is that when the beams recombine at the beamsplitter, they do so by adding their electric fields, and adding electric fields is not the same as adding powers. To see this, write the electric field of the recombined beam going toward the detector as (leaving off some of the extraneous factors for now)

$$E_{\text{det}}(l_1, l_2) \sim \frac{E_0}{2} e^{-ikl_1} + \frac{E_0}{2} e^{-ikl_2}$$

The first term in this expression represents the beam that traveled down the first arm and back, a total distance l_1 , picking up the phase factor e^{-ikl_1} , while the second term represents the second beam that traveled a distance l_2 . Rewriting this

$$E_{\text{det}} \sim \frac{E_0}{2} [1 + e^{-ik\Delta l}] e^{-ikl_1}$$

where $\Delta l = l_2 - l_1$ is the path difference between the two arms, we then take the absolute square to get the power hitting the detector

$$\begin{aligned} I_{\text{det}} &\sim |E_{\text{det}}|^2 \\ &\sim \frac{E_0^2}{4} |1 + e^{-ik\Delta l}|^2 \\ &\sim \frac{E_0^2}{4} (1 + e^{-ik\Delta l})(1 + e^{+ik\Delta l}) \\ &\sim \frac{E_0^2}{4} [1 + \cos(k\Delta l) - i \sin(k\Delta l)][1 + \cos(k\Delta l) + i \sin(k\Delta l)] \\ &\sim \frac{E_0^2}{4} [2 + 2 \cos(k\Delta l)] \end{aligned}$$

If we do this carefully to get all the factors of two right, then the power hitting the detector becomes

$$I_{\text{det}} = \frac{I_0}{2} [1 + \cos(k\Delta l)] \quad (1)$$

We could do the same calculation for the light that ends up going back toward the detector (the other path through the beamsplitter; see Figure 3), and this would give

$$I_{\text{back to source}} = \frac{I_0}{2} [1 - \cos(k\Delta l)]$$

Adding these two paths gives I_0 , which just means that energy is being conserved.

Figure 4 shows a plot of I_{det}/I_0 as a function of x/λ for this ideal interferometer, where $x = \Delta l/2$ is the displacement of the translating mirror. (The path length change Δl is equal to $2x$ because the beam travels to the mirror and back again.) Note that I_{det} goes from a minimum of zero, when the cosine term equals minus one, to a maximum of I_0 when the cosine equals one. Figure 4 is probably the most important plot in this experiment, so make sure you understand what it means. If not, ask someone about it.

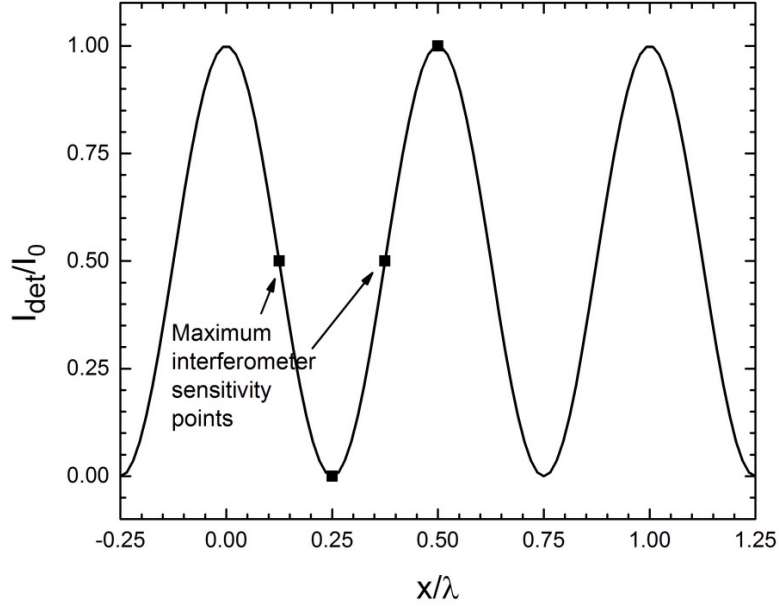


Figure 4. A plot of the power hitting the detector for an ideal interferometer, divided by the initial laser power I_0 , as a function of the displacement x of the translating mirror divided by the light wavelength λ . The interferometer is most sensitive to small changes in x at the points shown, when $I_{\text{det}} = I_0/2$.

Before you go into the lab and start playing with optics, let's take a closer look at how the interferometer shown in Figure 3 gets translated into a real experiment. We use a photodetector (a photodiode and some amplifying electronics) to convert the incident optical power to a voltage V_{det} that is proportional to I_{det} , converting the optical power I_{det} (in milliwatts) to a measured voltage V_{det} (in volts). Also, in a real interferometer the two beams do not recombine perfectly, so the light intensity does not go all the way to zero as given in Equation 1. Putting all this together, the actual detector voltage becomes

$$V_{\text{det}} = \frac{1}{2}(V_{\text{max}} - V_{\text{min}}) [1 + \cos(2kx)] + V_{\text{min}} \quad (2)$$

where we have written this so that V_{det} goes from a minimum V_{min} to a maximum V_{max} .

If you want to measure large displacements with $x \gg \lambda$ using an interferometer, then you can do so by counting dark/light cycles as you translate the mirror. Our focus here will instead be on measuring very small displacements $x \ll \lambda$, and in the lab we will see just how small. In this case, the interferometer will be most sensitive to changes in mirror position x when dI_{det}/dx is maximum, as shown in Figure 4. Doing the math, and using the measured V_{det} instead of I_{det} , we get

$$\begin{aligned} \frac{dV_{\text{det}}}{dx} &= \frac{1}{2}\Delta V (-2k) \sin(2kx) \\ &= -k\Delta V \sin(2kx) \end{aligned}$$

and

$$\begin{aligned} \left| \frac{dV_{\text{det}}}{dx} \right|_{\text{max}} &= k\Delta V \\ &= \frac{2\pi\Delta V}{\lambda} \end{aligned} \quad (3)$$

where we defined $\Delta V = V_{\text{max}} - V_{\text{min}}$.

Putting in $\lambda = 633 \text{ nm}$, we find that a displacement of the mirror by δx at this most sensitive point gives a change

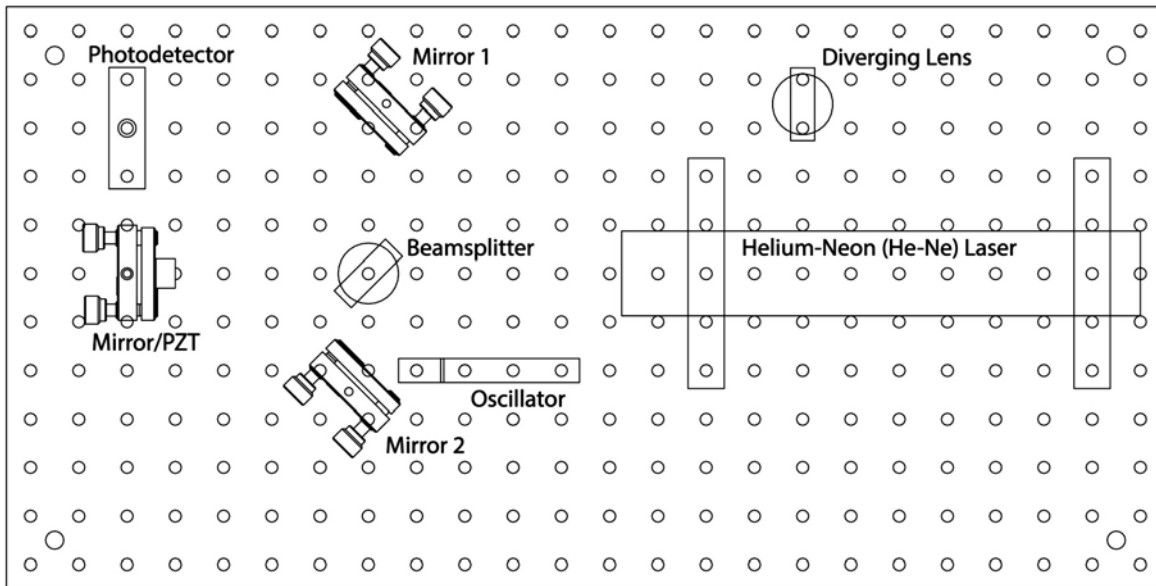


Figure 5. The optical breadboard (aluminum plate with holes) and the various optical elements. Not shown are two beam blockers (just paper cards on small stands) and the clamps that hold Mirror 1, Mirror 2, and the beamsplitter in place.

in V_{det}

$$\begin{aligned} \delta V_{\text{det}} &= \frac{2\pi\Delta V}{\lambda} \delta x \\ &\approx \frac{\delta x}{100 \text{ nm}} \Delta V \end{aligned} \quad (4)$$

Let's put this in terms of atoms, assuming an atom is about 0.1 nm in size. Then a mirror displacement of one atom gives an interferometer signal of $\delta V_{\text{det}} = \Delta V/1000$. In the lab you will measure that ΔV is about 6 volts. So if you can see δV down to 6 millivolts, then you can see a mirror displacement of just one atom (0.1 nm). If all goes well, you will in fact be able to measure displacements substantially smaller than that.

2 Lab Procedures – Week One

Your main objectives for the first lab session are to build and align a simple Michelson interferometer, and then determine how well you can measure small displacements with this interferometer. What follows are some step-by-step instructions on how to accomplish these objectives. Each paragraph has a small task, and in series they will guide you through the lab. As you can see, there are a lot of steps, so you will need to work efficiently. If you find that one step is taking quite a while, you might want to ask for help before it chews up too much time.

- Before you turn on the Laser Interferometry apparatus, take a look at the optical elements on the breadboard. Figure 5 shows a diagram of how the previous group probably left things, with the various optical elements labeled. Note there are three optical elements held down with large, grey, fork-like clamps. These pieces (Mirror 1, Mirror 2, and the Beamsplitter), along with the Diverging Lens and the two beam blockers (white cards in stands), are yours to move around during this experiment.

The remaining pieces (the He-Ne Laser, the Photodetector, the Mirror/PZT, and the Oscillator) are not meant to be moved, so **please do not try to move these pieces**. (More on what the PZT stands for later.)

- Another important point: **please do not touch any optical surfaces**. The mirrors are front-surface mirrors, plus the beamsplitter and lens have dielectric coatings on both sides. All these surfaces are quite fragile and easy to damage. Fingerprints leave oils and acids that damage the coatings over time. If you see any fingerprints on

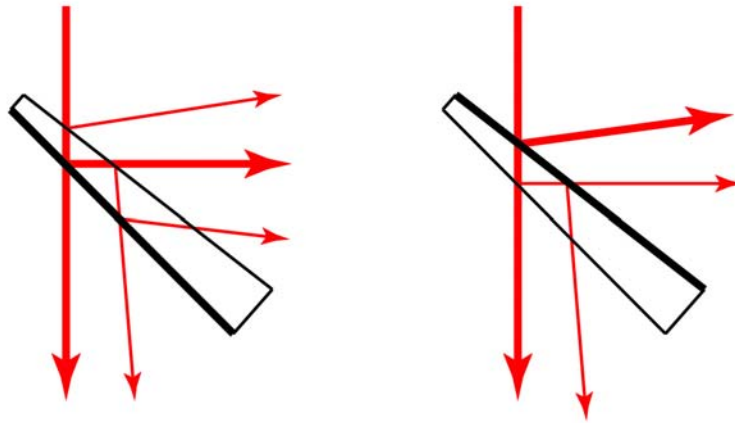


Figure 6. This drawing shows various reflections from the wedge-shaped 50:50 beamsplitter. The thick black line is the 50:50 coated surface, while the thinner line has a much lower reflection. Note that the drawing includes all beams that reflect at most once from the low-reflectance surface.

any optical surfaces, or if the surfaces are dirty or damaged in any way, please tell someone. Do not try to clean the optics yourself; it is very easy to scratch the coatings while cleaning if you do not use the correct materials. Optical components are expensive, so please handle them with care.

- With all that in mind, begin by removing the three fork clamps and placing the bolts and clamps in the tray provided. Be careful not to knock any of the pedestals over when you remove the clamps. Note how the clamps can hold down the optics at any position on the breadboard, and at any rotation angle. Set all these optic elements behind the laser (with the Diverging Lens in Figure 5) so they are out of the way. Turn the chassis power on, and in a few seconds (sometimes up to a minute) the laser beam should appear. The beam hits the Mirror/PZT and then reflects back toward the laser. Use a white card to see the laser beam coming out of the laser. On the front panel of the Laser Interferometry chassis, set the Lock/Scan switch to the Scan position, and remove any cables connected to the front panel.
- Next place the Beamsplitter in the beam, as shown in Figure 5. You don't need to clamp the beamsplitter down yet, just set it on the breadboard. Then place a beam blocker in front of the Mirror/PZT to block the transmitted beam. Follow the reflected beam down a few meters using a white card, and you will see the beam expand slowly with distance, with a divergence of roughly one millimeter per meter. If you were to measure the intensity profile of the beam, you would find that it follows a Gaussian $I(r) = A \exp(-2r^2/w^2)$, as mentioned above. If the beam shape is not nice and round, perhaps the beam is hitting the edge of the Beamsplitter and getting corrupted. If so, reposition the Beamsplitter.
- Try putting the Diverging Lens in the beam where it exits the laser (again no clamp needed), and you will see the beam diverge much faster, still maintaining its Gaussian profile. If you look carefully you may see some imperfections in the beam profile, probably from dirt on the optics.
- With the Diverging Lens removed, you should be able to see two or three beams reflecting from the beamsplitter – one bright beam and one or two faint beams. The extra beams come from the fact that the beamsplitter has two sides. One side has a dielectric coating that causes about a 50% reflection. The other side as an anti-reflection (AR) coating that results in about 1% reflection (a bare glass surface would give about 4% reflection at each surface). Also, the beamsplitter is wedge-shaped, so the reflected beams come out at different angles. Figure 6 shows where the different beams come from. This drawing does not include any beams that undergo two 1% reflections, as these are quite dim and hard to see. Try the beamsplitter in both orientations and observe the different reflected beams shown in Figure 6. These are easiest to see with a white card about 30-50 cm down from the beamsplitter. If you cannot see all the faint beams, probably one is hitting the edge of the beamsplitter mount, so move the beamsplitter to keep this from happening. Yes, you can do all of this without ever touching the optical elements, so please don't touch any glass or mirror surfaces.

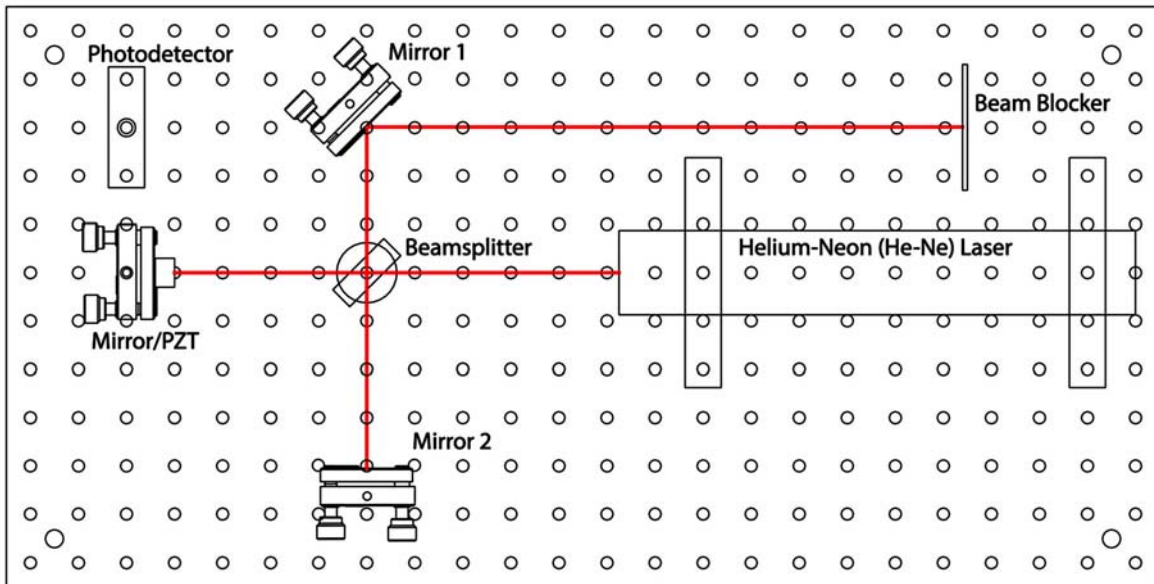


Figure 7. An initial set-up to examine various aspects of the interferometer optics. The beam blocker is a white card on a stand.

- In Figure 6 you see that fewer faint reflections are present if the 50% surface is nearest the laser (second orientation in the Figure), so orient the beamsplitter like this from now on.
- Next you should learn the five rules of aligning laser optics:
 - 1) Keep all laser beams at a fixed height. In your case, the optics are already set up to satisfy this rule.
 - 2) Keep the beams on a fixed right-angle grid. That is, make all beam reflections be right angles. This keeps things neat and orderly, and some optical elements (the 50:50 beamsplitter, for example) are designed for right-angle reflections.
 - 3) Make sure each beam hits each optical element close to its center. The optical elements are not large, and it is surprisingly easy to have a beam hit the mounting hardware instead of the intended optic. This is called *vignetting*, and you obviously don't want this to happen.
 - 4) Make especially sure that the beam hits the center of the photodiode. This is basically an extension of Rule 3, but it deserves special treatment because the photodiode is so small that there is little room for error.
 - 5) Keep all adjustment screws near the middle of their range. If an adjustment screw is near the end of its range, then soon it will lose its function. The adjustment screws are only needed for that last fine tweaking. So put the screws in the middle of their range and do the gross alignment before you tighten the clamps.
- With these rules in mind, you will next set up the optics as shown in Figure 7. One of the keys to working with optics is to place each element down carefully and methodically, so the alignment goes smoothly later. Begin by placing the beamsplitter 10 cm from the Mirror/PZT surface, as shown in the Figure. The holes are on a 1-inch (2.54 cm) grid, so four holes is about right. When placing the beamsplitter, you want to accomplish several things at once: 1) the beam should be close to the center of the beamsplitter; 2) the distance from the beamsplitter to the Mirror/PZT should be close to 10 cm; and 3) the reflection should be at a right angle (here again you can use the holes in the breadboard to guide you). If you place the beamsplitter with care, then you can accomplish all these things.

You will likely find that the fork clamps take a bit of getting use to. A good strategy is to first place the beamsplitter about where you want it and decide which hole you will use for the clamp. Then place the clamp in loosely while not worrying about the beamsplitter alignment. Finally align the beamsplitter just how you want it and tighten the clamp. Use the Allen wrench to tighten the bolts pretty tight, but no more than what you can do easily by hand. No need to be a gorilla.

- Next place Mirror 1 as in the Figure (no clamp needed) and place the beam blocker as in the Figure, and you

should see several reflected beams of different intensities.

- Before putting any more optics down, align the Mirror/PZT so the initial reflected beam goes straight back into the laser, then misalign the mirror a bit, and you should see a fairly bright secondary beam near the main beam on the white card. Change the Mirror/PZT pointing to see how the beams move on the card. The origin of this secondary beam is as follows: the main laser beam comes out of the laser, reflects off the Mirror/PZT, goes back toward the laser, reflects off the front mirror of the laser, reflects again off the Mirror/PZT, and then reflects off the beamsplitter and onto the card. This is an unwanted beam we will have to deal with later. For now, align the Mirror/PZT so this secondary beam overlaps the main beam.
- Next place Mirror 2 so it is also 10 cm from the beamsplitter and clamp it down. Importantly, try to place Mirror 2 so the two arms of the interferometer have the same length, as close as you can anyway. Align Mirror 2 so its bright reflected beams overlap the other bright beams on the card. Now see what happens when you misalign Mirror 2 a bit. Quite the mess to contend with! This is one of the practical realities of laser optics – unwanted beams. Again align Mirror 2 so the brightest beams all overlap.
- Next put a beam blocker in front of Mirror 2, and then point the Mirror/PZT so the secondary beam spot is just above the main beam, then keep moving it up until the secondary beam disappears. This happens when the mirror is so misaligned that the first reflected beam no longer hits the front mirror of the laser.
- Next remove the beam blocker from in front of Mirror 2 and align Mirror 2 so the main beams overlap on the card. You should see the combined beams flash a bit as they interfere with one another.
- Place the Diverging Lens directly in front of the laser and observe what's on the card. You may have to move the lens around a bit to keep the beam, now quite large, from hitting the different optical mounts. Once you have a nice round spot on the card, you should be able to tweak the alignment of Mirror 2 to see a bulls-eye pattern. If you misalign Mirror 2, you see curved *fringes* on the card. The bulls-eye pattern is actually a bit complicated to explain, since it appears when a diverging beam hits the beamsplitter. One has to follow the curved wavefronts to understand this completely, but doing the math is not terribly enlightening.
- If you push gently down on different places on the breadboard, you can see the fringes move. This is because you are flexing the aluminum plate, which in turn changes the lengths of the interferometer arms. Half a wavelength is a pretty small distance, so it doesn't take much of a push to see the fringes go through one cycle.
- Now move the Diverging Lens so it is between the Beamsplitter and Mirror 1 (see Figure 7). Now as you misalign Mirror 2 you see straight fringes, which is what you expect to see from the interference of two plane waves. If you misalign Mirror 2 quite a bit, you will see two separate beams, so far apart that they no longer interfere, and thus no more fringes. If the beams partially overlap, then you see fine fringes, but only between the beams where they overlap. It is a simple thing, laser interference, yet there is quite a lot to see as you move the beams around.
- If you align Mirror 2 quite carefully to overlap the beams and eliminate the fringes as best you can, you will see that the laser spot goes from bright to dim as you press gently on the breadboard. In a perfect world, you could get a complete overlap of the two beams, so the beam intensity went all the way to zero. But nothing is ever perfect in the real world, alas. Try if you like, but you will not achieve a perfect null on the white card.
- Once you understand the interferometer fringes to your satisfaction, remove the Diverging Lens and realign the Mirror/PZT so that secondary beam reappears, then align so this beam overlaps the main beam. You probably want to block Mirror 2 during this realignment. Next unblock Mirror 2 and align it for full overlap of the beams, like you had before.
- Next set up the interferometer shown in Figure 8 and clamp Mirror 1 down. Check that all the clamps are secure, and that all the beams overlap at the photodiode.
- Almost there; just the final alignment left. Block Mirror 2 and put a second blocker right in front of the photodiode. Tweak the Mirror/PZT alignment so that secondary beam is about 3-4 mm above the main beam. You need to separate these beams, but at the same time you don't want to be too far from perfect retroreflection, so 3-4 mm is a compromise. This is just enough to keep the secondary beam off the photodiode. (As you may be noticing, experimental physics takes place in the real world – imperfections and compromises are part of the deal.)
- Next unblock Mirror 2, align it so the two main beams overlap on the white card (the two secondary beams will also overlap about 3-4 mm above), and align Mirror 1 so the overlapping main beams strike the center of the

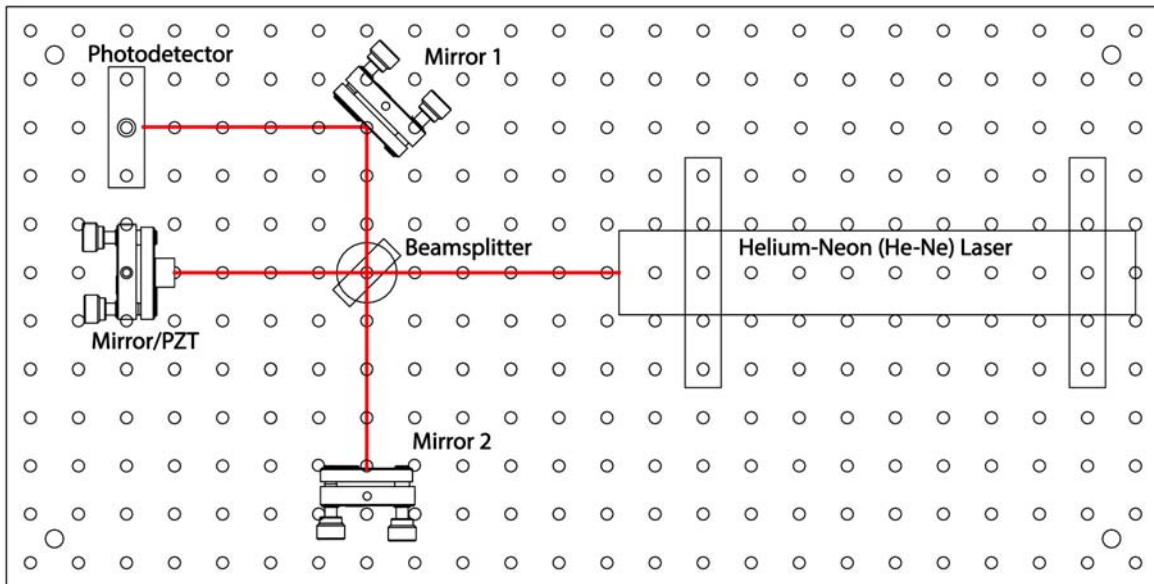


Figure 8. The optical setup for a basic Michelson interferometer. Note both mirrors are about 10 cm (4 inches) from the beamsplitter, and note how the beams follow right angles defined by the hole pattern.

photodiode. The two secondary beams should be above the photodiode (which has a 3.6 mm by 3.6 mm sensing area) where they will not contribute to the interferometer signal.

- Use a white card to make sure that the principal beam strikes the center of the photodiode. Note how the various degrees of freedom are used for different tasks: the Mirror/PZT pointing is used to define how closely the laser beam is retroreflected back into the laser, thus positioning the secondary beam just above the main beam. Then Mirror 2 is used to overlap the recombined beams. And finally Mirror 1 is there so you can center the beam on the photodiode.
- Once the optical setup looks good, connect the Signal Generator OUT-2 to channel 1 of the oscilloscope. Turn the Frequency Coarse Adjust knob all the way down, and turn the Amplitude Adjust knob all the way up. Trigger the oscilloscope on channel 1, and you should see a square wave or a triangle wave.
- The knobs and switches on the right-hand side of the front panel control a simple signal generator inside the main chassis. There is a switch to change between a triangle-wave or square-wave output, and knobs to control the signal generator amplitude and frequency. Flip the switch and turn the knobs to see what happens to the signal generator output on the oscilloscope.
- Next set the Frequency Coarse Adjust to make the lowest frequency triangle wave (the Fine Adjust doesn't matter here), and turn the Amplitude all the way up. Next connect the Signal Generator OUT-1 (which is exactly the same as OUT-2 as you see it on the oscilloscope) to the Scan IN, and verify that the Lock/Scan switch is set to Scan. Finally, connect the Photodiode OUT signal to channel 2 on the oscilloscope, and set things up so you see both signals.
- What's happening is this: By sending the triangle wave to the Scan IN port, you are sending a voltage to a piezoelectric transducer that sits behind the small mirror in the Mirror/PZT housing. (PZT stands for PieZoelectric Transducer; or for Lead-Zirconate-Titanate, the most commonly used piezoelectric material.) The PZT expands when a voltage is applied, so the mirror position x is proportional to the voltage you apply. Put another way, the signal you see on the oscilloscope channel 1 is proportional to the mirror position x . The triangle wave means that the mirror is simply moving back and forth in a linear fashion. Meanwhile, the Photodiode OUT signal gives a voltage proportional to the light power hitting the photodiode. By observing these two signals together on the oscilloscope, you can see the interferometer signal $V_{\text{det}}(x)$ as given in Equation 2, and shown in Figure 4. You should see an oscillatory signal, but one that goes back and forth as x executes a triangle wave. With the Amplitude Adjust turned all the way up, you should see 2-3 fringes on the oscilloscope. If you are not seeing

this, check the optical alignment or ask someone for help.

- At this point the interferometer is probably not aligned very well, so V_{\min} will not be very close to zero. Tweak the alignment of Mirror 2 to minimize V_{\min} as best you can. To do this, turn up the gain on the oscilloscope so you can see V_{\min} more closely. You should be able to get to $V_{\min} < 0.5$ volts, and < 0.2 volts if the alignment is good. V_{\max} should be around 7 volts or more. Press on the breadboard gently and see how the signal responds. If this doesn't make sense, ask someone.

- When you have minimized V_{\min} , measure the fringe *visibility*, sometimes called the fringe *contrast*, defined by

$$S = \frac{V_{\max} - V_{\min}}{V_{\max} + V_{\min}}$$

with possible values ranging from 0 to 1. The Measure feature on the oscilloscope is useful here, since it will show you V_{\min} and V_{\max} directly. You should be able to get $S > 0.9$. Write this value in your notebook, and do a screen shot of the oscilloscope (just press the print button if you have a networked oscilloscope) and tape this into your notebook. Record V_{\min} , V_{\max} , and $\Delta V = V_{\max} - V_{\min}$ as well.

- At this point, stop and pat yourself on the back for putting together a working laser interferometer. You can use the PZT to adjust x , and then you can see $I_{\det}(x)$ from the photodiode signal. From your observations, calculate dx/dV_{PZT} . That is, calculate how much the length of the PZT changes when you change the applied voltage. Everything you need for this is on the oscilloscope: the voltage going to the PZT is what you are sending to the Scan IN port, and the change in x can be seen from the interferometer response, and you already know $\lambda = 633$ nm. One method for getting the numbers you need is to freeze a good trace on the oscilloscope (by pressing the Run/Stop button), then use the cursors feature to make measurements. Compare your answer with the manufacturer's specifications for this PZT, which say it has a displacement of $10 \pm 1 \mu\text{m}$ at 150 volts, and record all this in your notebook.
- Just in passing, you should be aware that the piezoelectric effect is a wonderfully useful method for moving things around electronically. It is used to focus lenses in portable cameras, to position things in scanning probe microscopes, to actuate quartz crystal oscillators in clocks, and in a host of other applications. The study of piezoelectric materials is a substantial area of current research, and new applications are being found all the time. If you ever need to position small objects carefully over small distances, you should remember piezoelectrics.
- So, now that you have a working laser interferometer, you can begin investigating different strategies for measuring small mirror displacements. To proceed, the next step is to "lock" the interferometer at its most sensitive point, and this is done electronically by monitoring $V_{\det}(x)$ and then setting the PZT voltage to maximize dV_{\det}/dx as described in Equation 3. We do not have time to discuss all the details of how this happens, but you can see it in action by setting the Set Point Adjust knob to the center of its range and flipping the Scan/Lock switch to the Lock position. The LED near the switch should stay green.
- When in Lock mode, the interferometer output on the oscilloscope should become a nearly constant value (jumping around some as the interferometer is being jostled). The electronics tries to keep the signal at some set point V_{set} , determined by the Set Point Adjust knob. If you turn this knob up and down, the interferometer output should follow. Use the oscilloscope output in the Scan and Lock modes to adjust the Set Point Adjust so $V_{set} = V_{\min} + (V_{\max} - V_{\min})/2$ – that is, so the lock point is halfway between V_{\min} and V_{\max} . When this is done, then dV_{\det}/dx is maximized (consider Figure 4 to see why this is true). By setting x to maximize dV_{\det}/dx , it will then be easiest to see small changes in x .
- The electronic circuit behind all this is called a *feedback loop*, as the measured V_{\det} is fed back to change x , which in turn changes $V_{\det}(x)$. This is also called a *servo mechanism*, and the math that describes it is part of *Control Theory*, which is a branch of engineering. Servo mechanisms are used in your home (for example in the thermostat controlling the room temperature), in your car, in your electronic devices, and of course in practically every modern physics experiment. Nearly every time you want to control a temperature, a current, a voltage, a position, etc., you can do so using a servo mechanism. You can take whole courses describing Control Theory in detail if you are interested (see the ME section of the course catalog).
- Notice the LED next to the Scan/Lock switch. If the interferometer is in Lock mode and this LED is green, then all is well; the lock should be working. But if you tap on the breadboard a bit, the servo may eventually stop working and the LED will turn red. The reason is that if x_{lock} satisfies $V_{\det}(x_{lock}) = V_{set}$, then $x_{lock} + \lambda/2$ will

satisfy it also (again consider Figure 4 to see why this is true). The servo will pick the nearest x_{lock} when you turn it on, but it may later jump to a different x_{lock} requiring a larger or smaller PZT voltage. If this voltage gets too high or too low, then the servo can no longer function properly. The bottom line is that if you see a red LED, then the interferometer is no longer properly locked. If the LED turns red, switch the lock off and on again to reset it. If this does not turn the red light to green, ask for help.

- Now the fun begins; the next step is to modulate the mirror position x while the interferometer is locked. To see this happen, lock the interferometer, turn the Amplitude Adjust knob all the way down, and connect Signal Generator OUT-1 to Modulation IN (disconnecting from Scan IN). Leave channel 2 of the oscilloscope connected to Photodiode OUT, and connect channel 1 to Sync OUT. Set the Multiplier knob to 100, switch the signal generator over to a square wave, and turn the Servo Gain Adjust all the way down. Finally, turn up the Amplitude knob a bit until you see a small square wave modulation on V_{det} . That is, V_{det} will equal V_{set} plus a small square wave modulation.
- Two things are happening here. First, the servo is in Lock mode, so on average V_{det} should be equal to V_{set} . But in addition to this, sending the signal generator output to Modulation IN tells the PZT to add a modulation given by the signal generator. The two commands to the PZT are separated in frequency: at low frequencies the servo wins, keeping $V_{det} = V_{set}$ on average. But at high frequencies the signal generator wins, causing the PZT position to be modulated by the input square wave.
- Sometimes these two controls work against one another. To see this, turn the Servo Gain up and down and see what happens. When the servo gain is low, then the servo only operates at very low frequencies, so the square wave modulation is seen faithfully on V_{det} . But when the servo gain is high, then the servo works harder to keep $V_{det} = V_{set}$. When the square wave modulation causes a jump in V_{det} , the servo sees this and adjusts x to compensate, bringing V_{det} back down to V_{set} within one half-cycle of the square wave. As a result, the square wave in V_{det} becomes distorted. There is quite a bit going on at once here, so it can be difficult to understand at first. But once you see it, it is not so complicated and it actually makes sense. Ask someone at this point if you need help absorbing it all.
- Next turn the signal generator frequency to about 700 Hz (use the measure feature on the oscilloscope, looking at channel 1), and turn the Servo Gain to its maximum. Use AC coupling for channel 2 on the oscilloscope and turn up the gain knob on the oscilloscope. Turn the Amplitude Adjust knob so the square wave you see in V_{det} has a top-to-bottom amplitude of about 500 mV. If you turn the modulation frequency down, you will see the distortion that happens when the servo partially compensates for the modulation. If you turn the frequency up above 700 Hz slowly, you will see some resonances in the mirror housing. Shaking the mirror at these high frequencies causes the mirror housing to shake as well, and the amplitude of this motion peaks when you hit a mechanical resonance.
- Adjust the signal generator frequency so that the V_{det} signal (channel 2 on the oscilloscope) looks similar to the Sync signal (channel 1 on the oscilloscope). The Sync signal has been rounded a bit (by sending it through a low-pass filter) to mimic the V_{det} response that has been similarly filtered. So you should be able to find a frequency where the two signals both look like similarly rounded square waves, somewhere in the neighborhood of 700 Hz.
- Now verify that the Multiplier knob is set to 100, and adjust the Amplitude knob so that the top-to-bottom amplitude is equal to $\Delta V/10$, which should be around 700 mV. Again, use the measure feature on the oscilloscope for this. Since $\Delta V/1000$ corresponds to a mirror shift of one atom (see Equation 4), we see that $\Delta V/10$ corresponds to 100 atoms, or 10 nm. Although 10 nm is a pretty small distance (about 10,000 times smaller than the 100- μm diameter of a human hair!), it is easily seen with the interferometer. The signal you see on the oscilloscope should have a high signal-to-noise and be easy to measure. Print out a screen shot and add it to your notebook. Again, if this is not making sense, ask someone for help.
- Next turn the Multiplier knob down to 10, which then corresponds to a mirror displacement of just 10 atoms (1 nm). If you increase the gain on the oscilloscope accordingly, you will now see that V_{det} is quite a bit jumpier. The noise mostly comes from seismic noise and acoustic noise. You could reduce the seismic noise by adding some additional rubber pads under the interferometer chassis, but just not bumping the chassis or the table helps quite a bit. (Note that the seismic noise here has little to do with seismic waves in the Earth; it is mostly from human-caused shaking of the floor and table.) Acoustic noise shakes the optical mounts directly, and you can see this on the oscilloscope when you talk or clap your hands. Keeping quiet helps reduce this noise, of course.
- Now try using the average feature of the oscilloscope (under the Acquire menu) to average 64 or 128 traces. This

slows down the response time, but substantially reduces the jitter. This works because the modulation signal is tied to the Sync signal, which the oscilloscope is triggering on, while all the noise is random in time. Thus averaging traces reduces the noise but not the signal. With some averaging, the modulation signal should be quite clear even with only 10 atoms of mirror motion.

- So far you have used several methods to enhance the sensitivity of the interferometer:
 - 1) First, you are modulating the mirror position at high frequencies. This helps because most of the noise, especially the seismic noise, is at low frequencies. This is a fairly ubiquitous phenomenon, broadly called *1/f noise*, since for many noise sources the noise intensity scales roughly as one-over-frequency. Putting the modulation signal at high frequency (700 Hz) means lower noise. If you go back to sample mode on the oscilloscope (not averaging), you can see that most of the noise is just the trace jumping up and down. This is low-frequency noise, and it averages away nicely when you switch back to average mode on the oscilloscope.
 - 2) Second, you are modulating the mirror position in a known way, in this case a square wave modulation. Because you then know how the V_{det} signal should be responding, you can look for just this response in the signal.
 - 3) And finally, you are averaging data. For random noise, the noise amplitude will be reduced by $1/\sqrt{N}$, where N is the number of averages.
- So far, so good. Next turn the Multiplier knob to 1, meaning a modulation of just one atom, or 0.1 nm. Keep the averaging on, and increase the gain on the oscilloscope to 20 mV/division. There is a lot of noise now, but you should still be able to see the square-wave signal in the V_{det} trace on the oscilloscope, and the peak-to-peak modulation should be roughly $\Delta V/1000$.
- We have one more trick up our sleeve, and that is to increase N still further. At 700 Hz, one cycle of the square wave takes a bit over 1 millisecond, so in 10 seconds you can average nearly 10,000 cycles, so then \sqrt{N} gives a factor of 100 – not a small gain in signal-to-noise. This averaging has been done for you in the interferometer electronics using the following steps:
 - 1) a running average V_{ave} is subtracted from V_{det} (by running the signal through a high-pass filter), which is similar to the AC coupling on the oscilloscope.
 - 2) the Sync signal is turned into a square-wave switching signal going from -1 to +1.
 - 3) the detector signal ($V_{\text{det}} - V_{\text{ave}}$) is multiplied by the switching signal and multiplied by 100.
 - 4) the resulting signal is averaged with a 10-second time constant, and the resulting signal is sent to the Amplitude OUT port on the front panel.

If you think about this a bit (maybe draw some sketches), you will see that the output from this process will be a new constant signal that is proportional to the modulation amplitude of V_{det} . In particular, the Amplitude OUT signal, call it δV_{mod} , will equal about 50 times the peak-to-peak square wave modulation signal in V_{det} . In other words, you expect $\delta V_{\text{mod}} \approx \Delta V/20$, since the modulation signal is $\Delta V/1000$. In fact, the averaging process is not perfectly efficient, so δV_{mod} will be more like $\Delta V/30$, mainly because the various square-wave signals are not perfect square waves. (Imperfections and compromises, once again.)

- You should now use a digital voltmeter to measure δV_{mod} (the Amplitude OUT signal) directly. This gives you a way to make a *quantitative* numerical estimate of the sensitivity of your interferometer, beyond just seeing a trace on the oscilloscope. You should see that δV_{mod} is clearly nonzero, and furthermore it is not terribly jumpy. To quantify this, read the voltmeter about once every five seconds (once the interferometer has settled down and the modulation signal is reasonably stable on the oscilloscope) and write these numbers down. Once you have accumulated 10-20 numbers, calculate the mean and standard deviation. The mean $\langle \delta V_{\text{mod}} \rangle$ corresponds to a mirror displacement of one atom, the calculated standard deviation σ_{mod} gives how well you can measure this mirror displacement in the 10-second averaging time. How many atoms of displacement are indicated by your calculated σ_{mod} ?
- Next turn the Multiplier knob to 0, which means there is no modulation being sent to the PZT. In principle, δV_{mod} should go to zero in this case, but the signal-averaging is not done perfectly in the electronics, so you will see a non-zero δV_{mod} . Record this signal every five seconds, like you did before, to determine $\delta V_{\text{mod}}(0)$ with a one-sigma error estimate.
- Finally set the Multiplier knob to $M = 0.1$, meaning a modulation of 0.1 atoms. Follow the same measurement

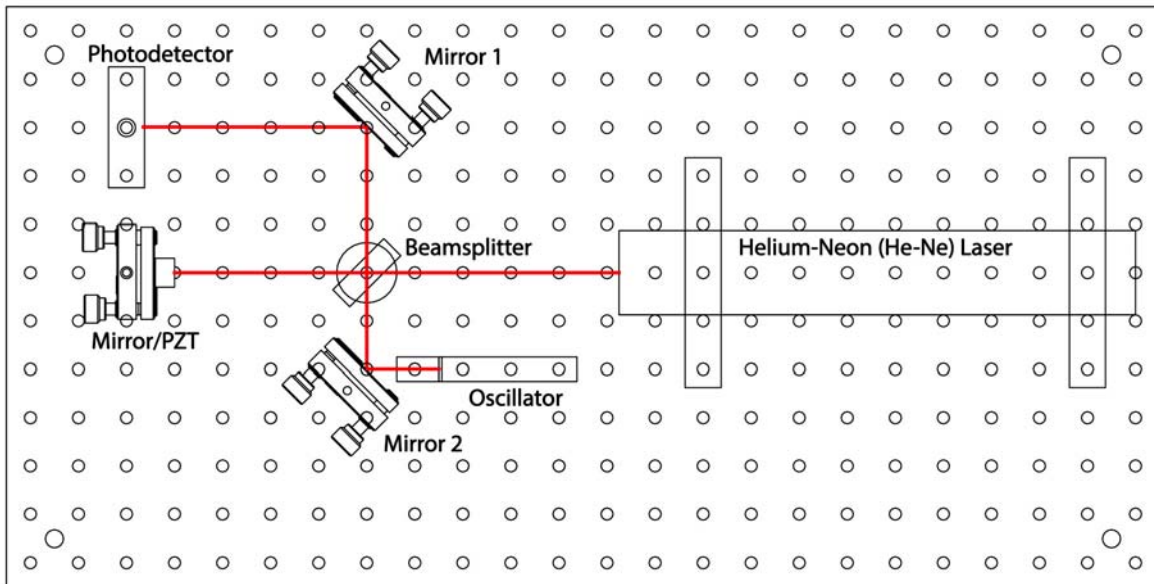


Figure 9. The interferometer set-up for the second week. Here Mirror 2 is a folding mirror, while the small mirror on the oscillator retroreflects the beam.

procedures as above. On a good day, when the external noise sources are not bad, both the $M = 1$ and $M = 0.1$ measurements should give you 1-sigma errors of about 0.01 atoms, or one picometer – not bad for a table-top device!

- Note that your 1-sigma error is for measuring changes of $x = dL$ in the length of the interferometer arm L , which is about 10 cm. So what is the dimensionless dL/L , called the *strain sensitivity* of the interferometer, if dL is 0.01 atoms? If you wanted to measure the distance from New York to Los Angeles (4500 km) with the same strain sensitivity, about what dL would you need?
- So, what should you take away from this experiment so far? First of all, you now know a bit about setting up a laser interferometer. You have seen the interferometer fringes with your own eyes, and you have seen the detector signal. You are now aware that piezoelectric devices can be used to push things around. And you have seen that laser interferometry can be used to measure very small displacements, and that there are a number of tricks for improving the measurement sensitivity.
- People at Caltech and elsewhere have worked for many decades to push laser interferometry to ever greater sensitivities, by carefully reducing seismic noise (by building vibration isolating mirror suspensions) and acoustic noise (by putting the interferometer in a vacuum), until finally one is dealing with quite fundamental noise sources, for example shot noise (the fact that one is making a measurement using a finite number of photons) and even radiation pressure from the laser beams themselves. The strain sensitivity of the LIGO (Laser Interferometer Gravity-wave Observatory) interferometers is currently below 10^{-22} for $L = 4$ kilometers, and there are plans to improve this by another order of magnitude or more.

3 Lab Procedures – Week Two

The next step is to examine the behavior of a mechanical oscillator in some detail by placing it in one arm of the interferometer. If you have already done the Magneto-Mechanical Harmonic Oscillator experiment, you will see some familiar concepts here (oscillating systems are common in physics). If not, no problem; this experiment will serve as an introduction to some oscillator physics, and you will see more by doing the MMHO experiment.

If you take a closer look at the oscillator (see Figure 9), you will see what is basically an aluminum post with a

small mirror attached to one side and a magnet attached to the other side. The post behaves like a rigid body that rotates about the fixed attachment point at the base of the post (like one arm of a tuning fork). In a nutshell, this device behaves like a simple harmonic oscillator; if moving freely, the post oscillates back and forth with some resonance frequency ν_0 . If you send an oscillating current through the nearby coil, the resulting magnetic field pushes and pulls on the magnet, and you have a driven harmonic oscillator.

Since the post pivots about its base, this system is best described as a torsional harmonic oscillator. If you displace the angle of the post by a small amount θ from its equilibrium position, then this flexes the aluminum at the base and results in a restoring torque $\tau = -\kappa\theta$. The equation of motion in angular coordinates becomes $\tau = I\ddot{\theta}$, where I is the mass moment of inertia.

Although these are the best equations for describing the oscillator motion, the interferometer is not very sensitive to θ , but instead measures the mirror position x . So it is convenient to change variables to $x = R_{post}\theta$, where R_{post} is the length of the aluminum post. Doing this gives us the new equations of motion

$$F = m_{eff}\ddot{x}$$

and

$$F = -k_{eff}x$$

where $m_{eff} = I/R_{post}^2$ can be thought of as an effective mass of the oscillator, and $k_{eff} = \kappa/R_{post}^2$ is an effective spring constant. (You can easily do all the math to verify these relationships; all we did here was a change of variables.)

Putting these two equations together gives the harmonic oscillator equation

$$m_{eff}\ddot{x} + k_{eff}x = 0$$

and solving this gives oscillatory solutions with the resonant frequency

$$\omega_0 = \sqrt{\frac{k_{eff}}{m_{eff}}}$$

where $\omega_0 = 2\pi\nu_0$ as usual (ν_0 is the oscillation frequency in Hz, or cycles/second, while ω_0 is the angular frequency in radians/second).

Now that we have the mathematics we need, we proceed with more step-by-step instructions:

- Begin in the lab by setting up the interferometer as shown in Figure 9. The set-up and alignment procedure is the same as you did last week, so we do not repeat it here. You will probably find that the fringe contrast is lower than you obtained last week. This is because some light is lost with the extra reflection, and because the two arms of the interferometer are not as close to being the same length as you had before. Measure the fringe contrast S as you did before, and write it in your notebook.
- Once the interferometer is set up and aligned, lock it so $V_{set} \approx (V_{max} + V_{min})/2$, same as you did last week. Make sure the Servo Gain Adjust knob is all the way up.
- Connect Signal Generator OUT-2 to channel 1 of the oscilloscope, and set the signal generator to produce a low-frequency triangle wave signal (remember this requires that you trigger on channel 1). Connect Signal Generator OUT-1 to Coil Drive IN, which sends a current through the coil near the oscillator. Verify that the Photodiode OUT signal is connected to channel 2 of the oscilloscope.
- With all this, the interferometer should be locked while you are monitoring the photodiode signal on channel 2 of the oscilloscope. At the same time you are driving the coil with the triangle wave you see on channel 1. If you are not driving the oscillator on resonance (probably the case), then not much will be happening at this point.
- Now sweep the signal generator frequency while watching the oscilloscope. When you are near the resonance frequency of the oscillator (around 700 Hz; use the measure feature on the oscilloscope), you should see the oscillator respond. The response will peak when you are at the resonance frequency. Use the Frequency Fine Adjust knob to maximize the oscillator response, and set the Amplitude Adjust knob so the peak-to-peak response of the interferometer is about $\Delta V/10$. With this as the peak response, the interferometer response will still be in the linear regime, so the voltage modulation will be proportional to the mirror position modulation.



Figure 10. A series of weights you can add to the oscillator, here arranged from lightest (on the left) to heaviest. You should find a pair of each of these in the tray.

- Maximize the oscillator response while using the oscilloscope to measure the drive frequency (of channel 1 on the oscilloscope). You should be able to measure ν_0 to about 1 Hz accuracy this way. Treat this as a precision measurement, so try to measure ν_0 carefully.
- Next you should add some weight to the oscillator and see how this changes the resonance frequency. Figure 10 shows a series of weights you should use for this purpose; you should have a pair of each of these in your tray. The weights screw into the holes near the top of the oscillator post, and for balance you should put identical weights on either side of the post. Turn off the servo when you change weights, and of course do not touch the mirror surface at any time. Screw the weights in snugly, but only with a modest two-finger-only tightening; again, no gorilla hands.
- Weigh each pair of weights before or after you measure the new resonance frequency. You should have a small electronic scale for this (a black box labeled AWS). Please be gentle with this scale, as it is a precision instrument. Set it down gently on the table, and set the weights down gently when you weigh them. If you blow gently on the scale, you can see how sensitive it is. The scale reads to 0.01 grams, but the actual accuracy is probably closer to 0.03 grams at best. If you weigh one mass, then the other, and then both together, you can get some idea about the actual accuracy.
- Go through all seven sets of weights, recording the mass of each pair m_{added} and the resonance frequency of the oscillator with the weights attached. Plot your data; specifically, plot ν_0 as a function of m_{added} and add this plot to your notebook.
- You expect your data to be described by

$$2\pi\nu_0 = \sqrt{\frac{k_{eff}}{m_{eff} + m_{added}}}$$

so from a fit to the data you could extract both k_{eff} and m_{eff} . In this form, however, the equation would require a nonlinear fit. While this is possible, a linear fit is quicker and easier. This can be accomplished by rewriting the above as

$$4\pi^2\nu_0^2 = \frac{k_{eff}}{m_{eff} + m_{added}}$$

$$\frac{1}{4\pi^2\nu_0^2} = \left(\frac{1}{k_{eff}}\right)m_{added} + \frac{m_{eff}}{k_{eff}}$$

With this reformulation in mind, plot $1/4\pi^2\nu_0^2$ as a function of m_{added} . This time you should get a straight line. Apply a linear fit to the data (using the software tool of your choice), and use the fit slope and intercept to

determine m_{eff} and k_{eff} , and write both quantities in your notebook. (Make sure the numbers have units.)

Now you should not report a measurement in physics without some kind of uncertainty estimate (although you would not be the first person to do so!). Unfortunately, a proper uncertainty estimate here is a bit tricky. In your classes (including this one), you will be taught mathematical methods for how to estimate uncertainties, and when you do a fit to your data the computer will spit out something for uncertainty estimates. You should know right now that the mathematical methods are not always terribly accurate or reliable, because they make all sorts of assumptions about your data, and these assumptions are not always valid. Moreover, the error estimates given by the computer fitting routines are *often* just terrible. Except in ideal cases, estimating uncertainties in measured quantities is not trivial. With this in mind, we will now look at some shortcuts for estimating uncertainties.

- Your uncertainties in k_{eff} and m_{eff} will depend on the uncertainties in both your measured ν_0 and m_{added} . At this point it is not even obvious which is more important. Probably one dominates, but sometimes both are important. Let's say that your typical ν_0 uncertainty is 1 Hz, while your typical m_{added} uncertainty is 0.04 grams. Go back to your plot of ν_0 as a function of m_{added} , and measure a slope of $\nu_0(m_{added})$ in the middle of its range. That is, determine a typical $d\nu_0/dm_{added}$. Compare this with $\Delta\nu_0/\Delta m_{added} = (1 \text{ Hz})/(0.04 \text{ grams})$. You will probably find that the two numbers are within a factor of two, meaning that the uncertainties in both ν_0 and m_{added} are about equally important, so in principle you should not ignore either one. Nevertheless, you probably found that $\Delta\nu_0/\Delta m_{added}$ was a bit greater than $d\nu_0/dm_{added}$, so for expediency we will neglect the errors in m_{added} for now.
- Here is a relatively simple trick that will allow you to make “quick-and-dirty” error estimates with any data, including your present data, called the *Monte Carlo method*. As the name suggests, you change some data at random, then go through your analysis and see how the result changes. Throw the dice many times, and you start to see some trends.

To see the Monte Carlo method in action, change your highest frequency data point by adding 2 Hz (don't forget to change it back later!), which is probably a bit worse than any of your individual measurements. Just change this one data point, then refit the data and generate a new set of fit (k_{eff}, m_{eff}) , and write these numbers down. Next set that frequency measurement equal to its correct value minus 2 Hz, and again refit the data. Doing these refits should take about a minute or two per trial, since you already have the file set up. So do this a total of six different times: ± 2 Hz on the highest frequency point (which you just did), ± 2 Hz on the lowest frequency point, and ± 2 Hz on a middle point.

- Finally collect all 7 sets of (k_{eff}, m_{eff}) (including the original analysis with the uncorrupted data), and calculate means and standard deviations for k_{eff} and m_{eff} (if you do not know how to do this, search online for “standard deviation calculator”). And there you go – this give you respectable estimates for $\sigma_{k_{eff}}$ and $\sigma_{m_{eff}}$ without a lot of effort. And you didn't even have to use any of those tedious error propagation formulae. Rewrite your original measurements (fits) of k_{eff} and m_{eff} in your notebook, this time with your error estimates.

Now, if you really wanted to do this well (for publishing a paper, for example), you would write a short program to change every data point at once by a random amount, picking the amounts from a suitable Gaussian distribution. And you would add random errors to both ν_0 and m_{added} , since both are important. Once your program has added random errors to all the data points, it would then recalculate (k_{eff}, m_{eff}) like you did above. Have your program do this a thousand times, each time with a different set of random errors, and you would obtain excellent uncertainty estimates.

Okay, that was somewhat long-winded, but remember these things: 1) When you report a measurement, it needs to include an uncertainty estimate, otherwise it's just not good science; 2) The uncertainty estimates your computer fitting routine spits out are often rubbish, and 3) the Monte Carlo method gives you reasonable uncertainty estimates without a huge amount of effort, and without making a lot of assumptions about your data.

- For your next set of measurements, remove all the added masses from the oscillator and return them to the tray. Please do a count to make sure you have one pair of each of the masses you see in Figure 10. (If not, let someone know.)
- Now drive the oscillator on resonance ν_0 and measure the peak-to-peak amplitude of the response on the oscilloscope. Switch to a square-wave drive signal and adjust the signal generator amplitude so the peak-to-peak signal is about $\Delta V/10$. Then do not change the Amplitude Adjust knob for the remainder of the measurements. Use the measure feature on the oscilloscope to measure the peak-to-peak amplitude when you are right on resonance.

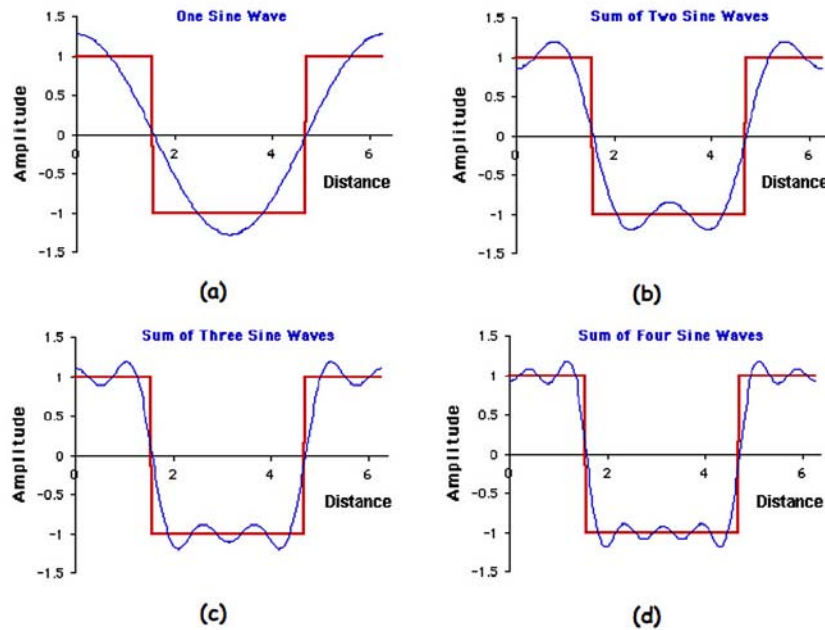


Figure 11. An example showing how a square wave can be expressed as the sum of a series of sine waves, as given in Equation 5. As more terms are summed (in plots a, b, c, and d), the result is a better approximation of the square wave.

That is, maximize the peak-to-peak amplitude as best you can using the Frequency Fine Adjust knob. Write this number down in your notebook along with your newly measured ν_0 (in case it drifted compared to your earlier measurement).

- Average 64 traces on the oscilloscope and repeat the measurement. You may get a smaller number this time, since you are averaging as you tweak the drive frequency, but do the best you can.
- Now drive the oscillator with frequencies $\nu_0/3$, $\nu_0/5$, $\nu_0/7$, and $\nu_0/9$ and again measure the peak-to-peak amplitude. Average 64 traces for each of these measurements, as this is the only way to get a reasonable signal-to-noise. You should see on the oscilloscope that the oscillator continues to oscillate at ν_0 , regardless of the drive frequency.

What's happening here is as follows. A square wave signal at a frequency ν_0 can be written as a sum of sine waves

$$f_{\text{squarewave}}(t) = \frac{4}{\pi} \sum_{n=1,3,5,\dots}^{\infty} \frac{1}{n} \sin\left(\frac{n\pi t}{T}\right) \quad (5)$$

and an example showing the first few terms in this series is shown in Figure 11. This is known as a Fourier series, and you will learn all about these in ACM 95. You do not need to know the whole theory to appreciate this one example.

When you drive the oscillator at $\nu_0/5$, for instance, then the oscillator doesn't respond to the drive signal at $\nu_0/5$; this is so far off resonance that it does little to excite the oscillator. The next term in the series has a frequency of $3(\nu_0/5)$, and again this does little to excite the oscillator. But the next term is at $5(\nu_0/5) = \nu_0$, and this is right on resonance. So this sine wave does excite the oscillator, and it excites it at ν_0 . But the amplitude of this term is only $1/5$ of the first harmonic, so the excitation amplitude should be $1/5$ of what it was when the signal generator was set to ν_0 . If doesn't make sense, ask someone. When you see it, it becomes clear.

- For your measurements, make a table comparing the measured oscillator amplitude A_n with A_1/n , where A_1 is the amplitude when the signal generator frequency was ν_0 . In this table you may note that systematic errors dominate over random errors. It becomes progressively more difficult to measure A_i as the amplitude decreases, and this systematic effect probably dominates over any other measurement uncertainties. If all goes well, you should see the data match the theory reasonably well.