Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

- Optics Track -

Diode Lasers & LIDAR

This handout describes two somewhat separate topics intended to be covered during a two-week period as part of the Optics Track in Ph77. The first experiment describes some properties of semiconductor diode lasers, including how they can be used to make frequency-tunable laser light for spectroscopy. These types of lasers are used in the Optical Resonators part of the Optics Track and in the Rubidium Spectroscopy part of the Atomic Track. Frequency-tunable lasers also play important roles in many areas of modern science and technology, and inexpensive diode lasers have been replacing larger and pricier gas-based and crystal-based lasers in an ever-increasing number of applications – a trend that is likely to continue in the future.

The second experiment demonstrates some basic concepts in time-of-flight LIDAR (LIght Detection and Ranging). As semiconductor diode lasers have become cheaper and more robust, new LIDAR applications have arisen throughout experimental physics and engineering, with autonomous cars being a recent popular example. This experiment introduces you to this rapidly emerging technology, including a discussion on how to calculate its effectiveness in different applications.

Part 1. Diode Laser Spectroscopy

Figure 1 illustrates the basic elements of any laser system: a power source (also called the *pump*), a lasing medium, and an optical resonator. The laser medium can be thought of as a gas of simple 2-level atoms, and the main job of the power source is to put these atoms into the upper excited state. As these excited atoms spontaneously decay, they emit light with $h\nu = E_{2level}$, and some of this light builds up in the optical resonator, where it bounces back and forth between the two mirrors in a standing wave. When this trapped light strikes the excited atoms, it induces *stimulated emission* that causes the atoms to emit light more readily than they would via *spontaneous emission* alone.

And this is where a key, purely quantum-mechanical effect comes into play: a stimulated-emission event sends the emitted photon into the same quantum state as the photon causing the stimulated emission. In practice, this means that the standing wave soon builds to high intensity, so nearly all the light emission is from stimulated emission, with only a tiny fraction being lost to spontaneous emission. When the laser reaches steady state, the output light is continuously replaced by stimulated emission in the laser medium.

The output power coming from a laser always shows some kind of threshold behavior, meaning that the output power is quite low when the pump power is low, but the output suddenly begins to increase dramatically above some threshold. The cause of this threshold behavior comes from the transition from spontaneous to stimulated emission. Looking at our toy model of a simple 2-level-atom laser in Figure 1, a low pump power means that not many atoms are driven into the exited state.

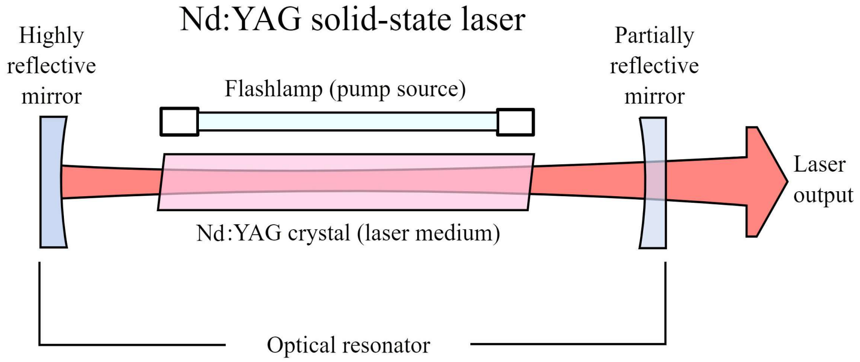

Figure 1. This sketch illustrates the essential components of any laser system. Here a flashlamp excites neodymium ions trapped in a Nd:YAG crystal (neodymium-doped yttrium aluminum garnet), sending them into a metastable excited atomic state. Laser light builds up in the optical resonator as this light induces stimulated emission from the Nb ions at 1064 nm. Some of this light leaks out through a partially reflective mirror to provide the laser output beam. The faces of the Nd:YAG crystal are cut at Brewster's angle to reduce scattering losses within the optical resonator. Gas lasers (such as helium-neon lasers, argon-ion lasers) have a similar construction, except the crystal is replaced by a gas-filled tube. In semiconductor diode lasers, atomic energy levels are replaced by electron-hole excitations in a semiconductor diode.

These atoms decay mostly via spontaneous emission, with emission in all directions. Most of this light is lost or absorbed, so there is not much light buildup in the optical resonator. The laser output power is then just that light leaking out the partially reflecting mirror in Figure 1, which is quite low.

As the pump power is increased, more atoms are driven into the excited state, and the light buildup in the resonator begins to become substantial. At some point, there is a rather abrupt transition (the laser threshold) when the atomic decay becomes dominated by stimulated emission. This stimulated emission is driven mainly by the standing wave in the cavity, so those photons are not lost like those from spontaneous emission. Above threshold, essentially all the atomic emission is into the cavity mode, where it contributes directly to the laser output power. Below threshold, note that the production of light is quite inefficient: a lot of input power is needed to produce little output power in photons. Above threshold, the conversion of electrical power to optical photons is higher (although still far below unity)

Laser technology is a fascinating mix of optical physics, atomic physics, semiconductor physics, and materials science, and popular laser types include gas lasers, solid-state crystal lasers, and semiconductor lasers. Even something as commonplace as a green laser pointer involves a remarkable mix of different technologies, as shown in Figure 2. Laser physics continues to be a very active area of current research, driven by countless applications in telecommunications, display technology, spectroscopy, and many other areas.

Semiconductor diode lasers

Semiconductor diode lasers rely on stimulated emission like all lasers, but the gain medium involves electron-hole pairs rather than atomic transitions. In a nutshell, sending a current through a semiconductor with p-type and n-type doping (see Figure 3) causes a recombination of electrons and

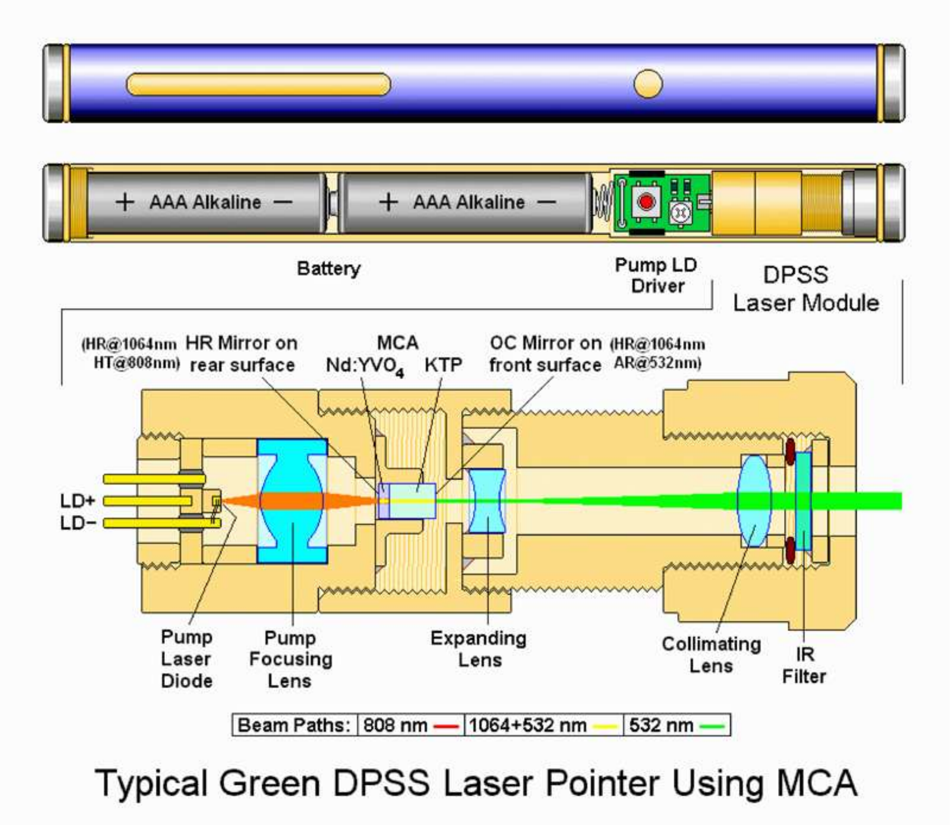

Figure 2. A common green laser pointer illustrates the complexities that are often present in commercial laser systems (DPSS = diode-pumped solid-state; MCA = multiple-crystal assembly). Here a near-IR semiconductor diode laser serves as the power source, optically "pumping" the ions in a neodymium-doped yttrium orthovanadate (Nd:YVO4) crystal. Note that the input mirror (a multilayer dielectric coating) on the crystal allows 808-nm pump light to enter the crystal but reflects the 1064-nm laser light inside the crystal. The output beam then passes through a KTP (potassium titanyl phosphate) nonlinear crystal that converts pairs of 1064-nm photons into 532-nm photons, thus "frequency doubling" the infrared light into green light. Finally, the beam exits through an IR filter that removes any residual (invisible) 1064-nm light. Someday this complex assembly will likely be replaced by a single green diode laser, but such a simple solution does not yet exist.

holes that emits light, the same process that makes LEDs. The vast subject of semiconductor physics is certainly beyond the scope of this handout but suffice it to say that the wavelength of the emitted light is largely determined by the bandgap within the semiconductor.

Creating a diode laser involves encasing the semiconductor gain medium within a microfabricated optical waveguide/resonator made using semiconductor materials with different refractive indices. Everything is fabricated on a chip and driven by sending an “injection current” through the diode. Once light leaves the tiny laser cavity, an external lens is needed to collimate the strongly diverging laser light into a narrow laser beam. (Alternatively, the semiconductor cavity output could be immediately coupled into an optical fiber.)

Our primary focus in this lab will be on using diode-laser technology to make tunable, single-frequency laser beams that we can use for atomic spectroscopy and other applications. This focus on

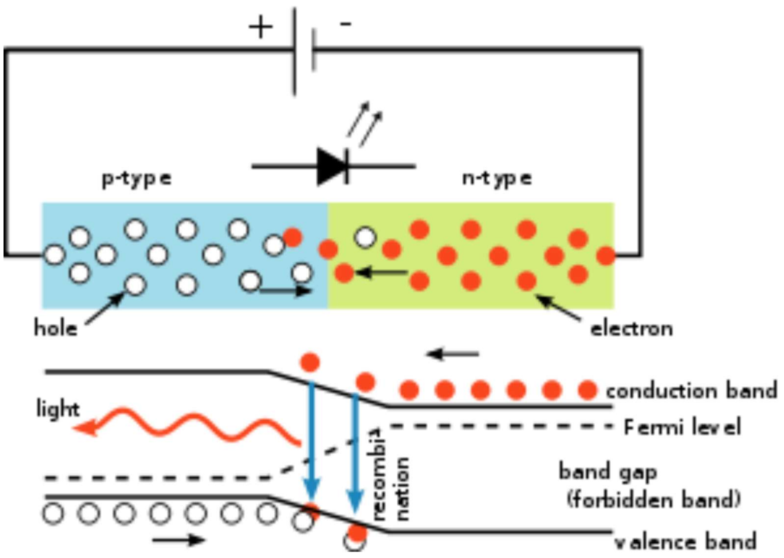

Figure 3. An LED consists of a p-type semiconductor (with an electron depletion, forming “holes” in the valence band) bonded to an n-type semiconductor (with an electron excess, populating the conduction band). When a current is driven through the junction, the electrons and holes recombine, emitting light. The photon energy is essentially equal to the energy band gap in the semiconductor.

spectroscopy is interesting in its own right, but it also provides some useful pedagogy, as it demonstrates many aspects of laser physics with a relatively simple laboratory setup.

Our story begins with a diode laser in a can, like the example shown in Figure 4. The chip itself is roughly 1mm in length, and of course this diagram glosses over the semiconductor physics, the optical layout, the electrical contacts, and all the vital details that make this device actually work. Unfortunately, bare diode lasers like this have a host of undesirable properties:

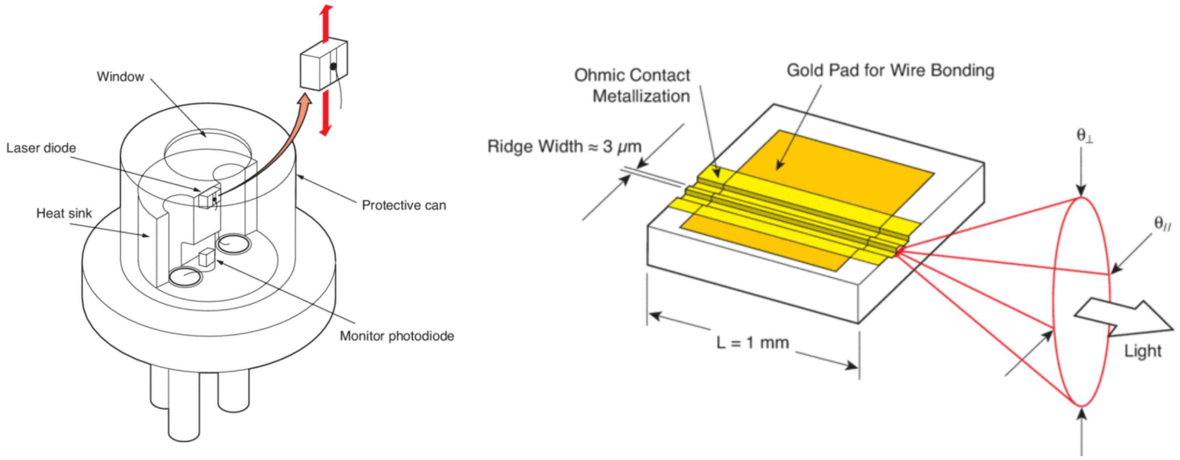

Figure 4. (Left) Individual diode lasers are often mounted in several-mm-diameter "cans" that include the laser itself on a heat sink, a protective window for the output beam, and perhaps a photodiode to monitor the optical power. (Right) The laser optical cavity is etched into a block of semiconductor material, using different materials to confine the cavity modes. Images from https://www.newport.com/t/laser-diode-technology.

- They have relatively low output power.
- The output wavelength is often ill-determined.
- The wavelength is difficult to tune.
- The laser beam quality is relatively poor.

On the other hand, diode lasers are quite inexpensive, and the output frequency can easily be modulated at GHz frequencies. The low price is an especially attractive feature, of course, so people have found clever ways to make diode lasers work in spectroscopic applications.

In this lab, we focus on a device called an External-Cavity Diode Laser (ECDL), which involves using a diffraction grating to send a small amount of controlled optical feedback into a bare diode laser, as is shown in Figure 5.

To begin to understand what is happening with this ECDL setup, we define a fictitious "generalized gain" that has the components shown in Figure 6. The curves in this figure provide a greatly simplified picture of a very complex phenomenon, but they can be thought of as the probability that lasing occurs at a particular frequency. The individual curves describe different physical effects, and collectively they contribute to the laser's operation. Semiconductor laser technology is an enormous field, of course, so our modest goal here is to examine just this basic, qualitative model of the underlying physics.

The *Semiconductor-medium* curve illustrates the semiconductor material itself, and the broad peak occurs at the electron-hole recombination energy. This curve is essentially the spectrum of a bare LED, which peaks at some wavelength and has spectral width of order 1-2 nanometers – much broader than the other curves in Figure 6. The only way to change this curve appreciably is by changing the semiconductors materials being used, although the peak wavelength will shift with temperature. There is a whole industry devoted to this curve, of course, and one can purchase diode lasers with a large (but finite) selection of different wavelengths.

The *Internal cavity* curve comes from the optical resonator (a.k.a. optical cavity) inside the semiconductor material, and here the peaks correspond to standing waves (a.k.a. resonant modes) inside the resonator. As mentioned above, all lasers involve optical resonators, and the laser is far

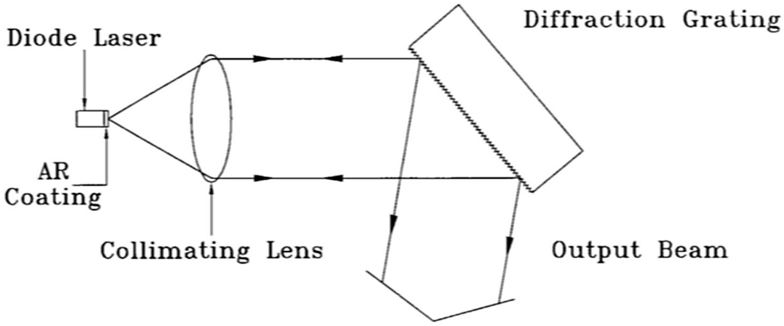

Figure 5. This illustration shows the basic parts of the grating-stabilized diode laser we use in the Ph77 lab. The Diode Laser chip contains the laser internal cavity and incorporates an antireflection (AR) coating to better couple the internal and external cavities. Most of the laser light exiting the semiconductor is collimated and then undergoes a simple reflection off the Diffraction Grating. But some light is diffracted from the grating and goes back into the semiconductor portion to form an external cavity.

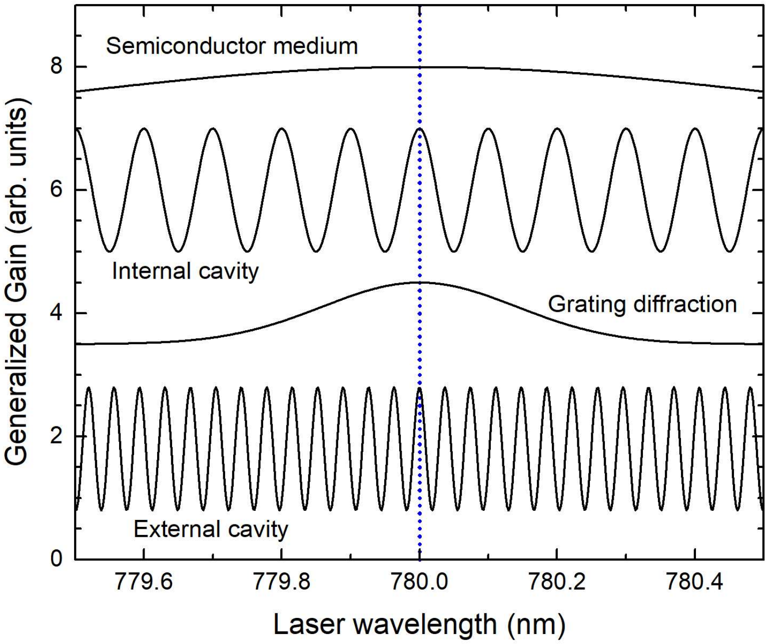

Figure 6. Tuning an External-Cavity Semiconductor Laser (ECSL) to a specific wavelength λ_0 requires a careful alignment of its various components. First (top curve), the semiconductor temperature is adjusted so lasing at λ_0 is preferred. Second (2nd curve), the semiconductor temperature and current are adjusted slightly so the internal cavity has a resonance at λ_0. Third (next curve), the grating angle is positioned so light at λ_0 is retroreflected back into the semiconductor laser. And finally (bottom curve), the grating position is adjusted so the external cavity also has a resonance at λ_0. Thus, achieving stable laser operation at λ_0 requires the simultaneous setting of the laser temperature, laser injection current, grating angle, and grating position. Note that the vertical axis is not to scale in any way in this diagram, and the curves are only meant to illustrate the underlying physics.

more likely to operate when the wavelength equals one of the resonant wavelengths of the cavity. Thus the generalized gain is highest at these points, and the laser will always operate at one of the cavity modes.

What are the cavity modes? Physically, each mode corresponds to a standing wave inside the cavity, and these have optical wavelengths λ_m given by

$$m\frac{\lambda_m}{2} = nL \tag{1}$$

where L is the length of the cavity, n is the index of refraction inside the cavity, and m is an integer index. This expression reflects the fact that there are an integer number of half-wavelengths for each resonant mode.

Exercise 1. Show that the spacing between cavity modes is

$$\delta\lambda \approx \frac{\lambda^2}{2nL} \tag{2}$$

or equivalently

$$\delta\nu \approx \frac{c}{2nL} \tag{3}$$

when $\lambda/L \ll 1$.

In the absence of the external grating, the laser behavior is driven by just the top two curves in Figure 6. The bandgap dictates that the laser will operate near the peak of the semiconductor-medium curve, and the resonant cavity forces the laser to operate at one of the cavity modes. Thus, as you might expect, the output wavelength of the laser is most likely at the cavity mode frequency nearest the top of the semiconductor curve. Again, Figure 6 is just a qualitative toy model, but it provides a useful physical picture of the laser physics.

Now imagine slowly sweeping the temperature of the laser. Both the bandgap and the index of refraction of the medium depend on temperature, so the top two curves in Figure 6 both shift in wavelength as the temperature changes. However, these two curves do not shift at the same rate. The result is that diode lasers tend to "mode hop" as the temperature is changed, as illustrated in Figure 7 (which shows the specific laser we are using in this lab). If you imagine shifting just the top curve in Figure 6, you can better see why this mode-hopping behavior comes about: the operating wavelength follows the peak of the semiconductor-medium curve, but it is also constrained to sit at one of the cavity modes. Material impurities and other laser imperfections make the mode-hopping behavior somewhat unpredictable and erratic, however, so the line in Figure 7 is not quite a simple sawtooth.

Interestingly, it is not an accident that the laser typically operates at just one internal-cavity mode. The reason arises from how stimulated emission drives the lasing process described above. Once the energy in one of the cavity modes rises a bit above the others, stimulated emission driven by the photons in this mode adds more optical energy to that same mode. This process reduces the density

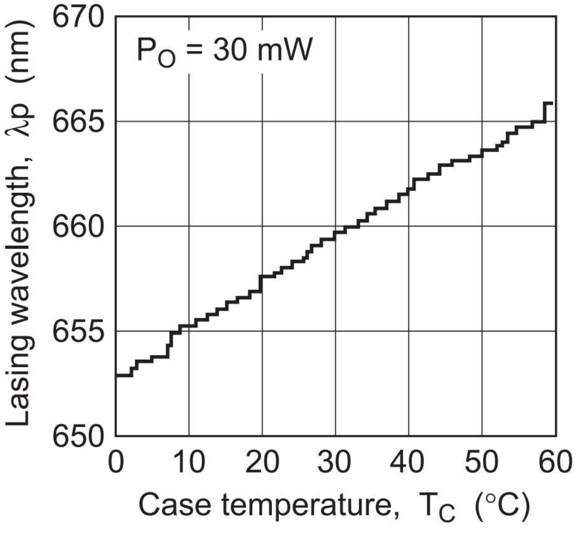

Figure 7. Diode lasers typically exhibit "mode hopping" that causes the laser wavelength to jump suddenly as a function of temperature. The reason for this lies in the semiconductor-medium and internal-cavity curves in Figure 6, as these two curves shift at different rates with changing temperature. (Figure from the HL6501MG data sheet.)

of excitations in the laser medium, thus robbing energy from other modes before they get started. Thus, stimulated emission produces a winner-take-all *mode-selection instability* in the laser operation: once one cavity mode grows in amplitude, it gobbles up nearly all the available optical energy in the system, depriving the other available modes. [Alas, little is simple in the real world, so diode lasers do frequently run "multi-mode"... but that is a complication we can ignore for now.]

Adding the external diffraction grating in Figure 5 changes things, as now the lower two curves in Figure 6 come into play. The *Grating diffraction* curve arises because laser light reflected back into the laser by the grating (which only happens for a specific wavelength) enhances the light buildup in the laser's internal cavity, and this enhances the overall gain at that specific frequency. Thus lasing is more likely to happen at this specific wavelength, which is set by the angle of the grating relative to the output laser beam. One can shift the Grating curve simply by rotating the grating, and this "knob" gives us a way to tune the laser wavelength, which you will see for yourself in this lab.

Lastly, the *External cavity* curve arises from the additional optical cavity that augments the internal cavity on the laser chip. This cavity includes the diode laser, the collimating lens, and the air gap out to the grating. The tilted grating surface acts like a "distributed" mirror that forms the end of the external cavity. Once again, lasing is most likely to occur at the highest generalized gain, so at one of the standing waves in this external cavity.

Often one wants a laser to operate at a specific pre-determined wavelength. For example, in the Rubidium Spectroscopy lab in Ph77, the laser must operate at the precise wavelength of the Rb atomic transition near 780 nm. To make this happen, all the curves in Figure 6 must combine to form a peak at this wavelength, as described in the figure caption. To sweep the laser frequency around this wavelength, one then shifts all these curves in unison (as much as this is possible). With our current level of tunable laser technology, some degree of meticulous engineering is required to make a tunable single-frequency diode laser function reliably. The technology continues to improve, but not as quickly as one would like. The market for single-frequency tunable lasers is relatively small compared to most laser applications (telecom, LIDAR, bar-code scanners, etc.) that just require a bright light source.

Exercise 2. Diode laser characterization. To better understand the steps needed to create a narrow-bandwidth tunable laser, let us walk through the different curves in Figure 6 for the case of the red laser (HL6501MG) used in the Optics Track. The Semiconductor medium curve mainly shows the spectrum of the underlying LED if there were no optical cavities to guide the light. For the AlGaInP semiconductor in the HL6501MG, the top curve in Figure 6 is centered at $\lambda_0 \approx 658$ nm and has a width of roughly $\delta\lambda_{semiconductor} \approx 2$ nm.

Expressing this in terms of the frequency bandwidth of this light source, we have

$$\delta\nu_{semiconductor} = \frac{\delta\lambda}{\lambda}\nu_0 \approx 1.4\ THz \tag{4}$$

where $\nu_0 = 4.55 \times 10^{14}$ Hz = 455 THz for a wavelength of 658 nm.

(a) For the Internal cavity curve in Figure 6, the free-spectral range (FSR; equal to the spacing between the internal-cavity peaks) is given by $\delta\nu_{internalcavity} = c/2nL_{int}$, where n is the index of refraction of the semiconductor material and L_{int} is the length of the optical cavity. Assuming $L_{int} \approx$ 1 mm and $n = 3.49$ for aluminum-gallium-indium-phosphide (AlGaInP), what do you expect for $\delta\lambda_{internalcavity}$ and $\delta\nu_{internalcavity}$? Express your answers in nm and GHz. These values give the spacing between the cavity modes in Figure 6.

(b) Moving on to the Grating diffraction curve, the width of the peak can be estimated from the resolving power of the external grating, which is given by

$$R \equiv \frac{\lambda}{\delta\lambda_{grating}} \approx N \tag{5}$$

where N is the number of grating lines sampled by the light. Assume that the grating in Figure 5 has 2400 lines/mm, the width of the laser beam is about 2 mm, and it strikes the grating at about a 45-degree angle. From this, calculate the width $\delta\lambda_{grating}$ and $\delta\nu_{grating}$ of the grating peak in Figure 6. Express your answer in nm and GHz.

(c) Finally, for the External cavity curve, assume $L_{int} \approx 1$ cm and $n = 1$ (air) to calculate $\delta\lambda_{externalcavity}$ and $\delta\nu_{externalcavity}$.

If you find this all somewhat confusing, that's because the physics of actual devices is messy. When billion-dollar industries are being developed, technology evolves with amazing speed, often by trial-and-error, and the underlying material science may not be well understood by anyone. Nevertheless, most of this lab focuses on the physics sketched in Figure 6, including observing it in the lab. Although this is a qualitative picture of diode laser behavior, it turns out to be remarkably handy when using these devices in experimental physics. As usual, it begins to make more sense when you start working with these devices in the lab....

Laboratory Exercises

To begin, locate the *Diode Laser Spectroscopy* hardware on the large optical table, consisting of a 12x12-inch aluminum breadboard with some optics already mounted. Also locate the laser temperature controller (Thorlabs TED200C) and the laser current controller (Thorlabs LDC205C) on the shelf above. You should see a pair of cables connecting the two controllers to the laser mount (Thorlabs LDM56) on the breadboard. As the names suggest, these instruments are designed to control the temperature of the diode laser and drive it with a known amount of current. Notice also that many optical components and connections have already been assembled on the breadboard; you can leave those alone for now as we examine the hardware in detail.

The LDM56 fastens the diode laser can (see Figure 4) to a rigid mount, as illustrated in Figure 8, and it places an aspherical collimating lens in front of the laser. The lens takes the strongly diverging laser output and produces a collimated laser beam that is neither diverging nor converging. (There is no such thing as a perfectly collimated laser beam, but the lens does a pretty good job.) We are using an HL6501MG diode laser that emits red light at a wavelength of about 658 nm.

Figure 8. The LDM56 holds the diode laser can (lower right) and a collimating lens (upper right). The bulk of this unit is a large heat sink used for controlling the temperature of the diode laser.

Next turn on the TED200C and make sure the Sensor is set to "Th 20kΩ", which is a 20kΩ thermistor. If you set the Display to Tact, this shows the value of the thermistor in kΩ. (Alas, this instrument does not convert kΩ to temperature for you. You have to use the conversion table, which you can find on the shelf.) Next select the Tset display to show the temperature set point (again in kΩ). Make sure that the instrument is set to a target temperature of 11 kΩ (about 23C). Press the TEC-ON button and you can see that the TED200C will regulate the laser temperature so that Tact goes to Tset. At this point all we want is a stable laser temperature near room temperature… the exact value of the temperature is not important.

Next turn on the LDC205C and adjust the display to show the I-LIM setting, which should be about 75 mA. Leave this setting alone, as it limits the laser current so you cannot burn it out by accidentally driving it with too much current. Then set the display to show the I-LD value, which is the laser-diode injection current in milliamps. Turn the current knob to full off (full CCW… gently), and then press the Enable button to send current to the laser. Adjust the current knob to 50mA and you should see a bright red laser beam, about what you get with a laser pointer.

When using the HL6501MG laser, you needn't worry much about laser safety, and no laser goggles are needed. It is difficult to do harm to yourself or anyone else with a few milliwatts of visible laser light, to the point that even children are allowed to play with low-power laser pointers. Of course, you still should never look directly into any laser beam … for the HL6501MG, that would be roughly equivalent to staring directly at the Sun. However, because several billion people could stare at the Sun but don't, having the Sun overhead is not much of an eye safety hazard. And neither is the HL6501MG.

Exercise 3. For your first laboratory task, measure the optical output power as a function of the injection current to observe the laser's threshold behavior. Leave the paper cup you see on Grating 2, and you should see a faint beam coming out of the small prism. Open the slit in front of the prism (please be gentle with the micrometer!) to let most of the beam through. Then measure the power in this beam using one of the New Focus model 2031 battery-powered photodetectors (same as you used in the first part of the Optics Track). Measure the laser output as a function of input current (up to 70mA) and plot your results. You should see a clear threshold in the laser power at around 40mA.

Recall what this means from the discussion above. Below threshold, the laser acts much like an LED; the recombining electron-hole pairs emit photons mainly via spontaneous emission. Above threshold, there is enough light in the laser cavity to cause stimulated emission, and this quickly becomes the dominant emission mechanism. There is no classical analog for this threshold behavior, making the laser a truly macroscopic quantum device.

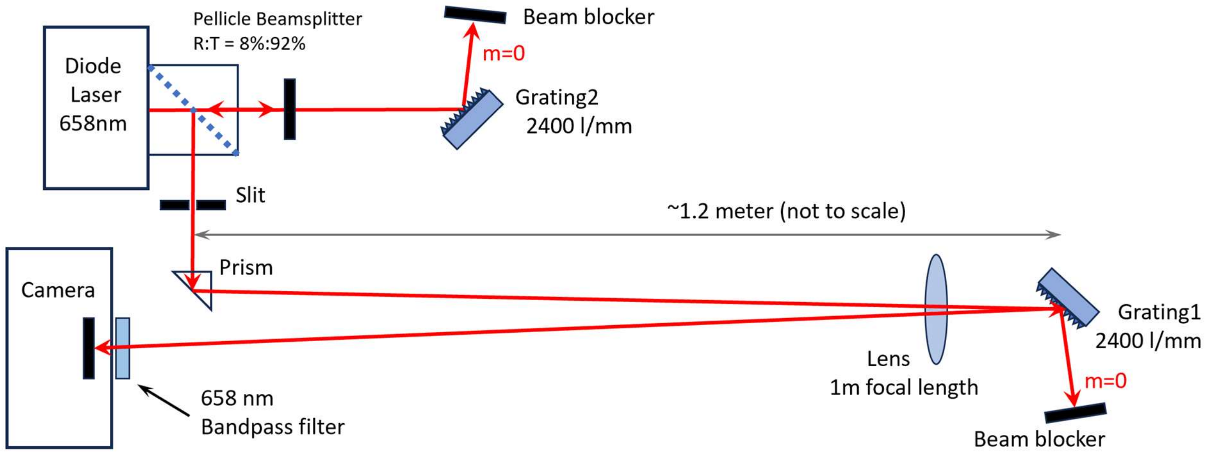

Figure 9. The optical layout for examining the diode laser spectrum.

Your next steps involve setting up some optics to observe the diode-laser behaviors described in the previous section. Figure 9 shows the overall optics layout, and many parts of this setup are already assembled for you (just to save time), as shown in Figure 10. Several of these optical elements are a bit tricky to align, and these are held down using black screws that require a smaller wrench than what you have on the optical table. Best to leave these optical elements unchanged unless you know what you are doing.

To begin setting up the optics, make sure that Grating2 is still covered with its opaque paper cup. Grating2 is part of the External Cavity described above, and we want to look at the bare laser without this grating first. Be especially careful that you **do not touch any grating surfaces**; there is no way to clean the fine grooves.

As can be seen in Figure 10, the camera lens has been replaced by a short tube, and a 658nm bandpass filter is attached to the end of the tube. If this filter is not already attached, you should be able to find it in a labeled box on or near the breadboard. Attach it as shown, and please screw it in *gently*; it needs to come off again later. This filter blocks most of the room lights while letting the 658nm laser light pass through with little loss. Without this filter, the room lights would interfere with the laser signal.

After placing the camera, set it to the correct height (4” above the breadboard), turn it on and set it to shutter priority mode (S on the mode dial). Then you can turn the unlabeled knob on top of the camera to adjust the shutter time (try it). The value of the shutter time appears on the lower-right

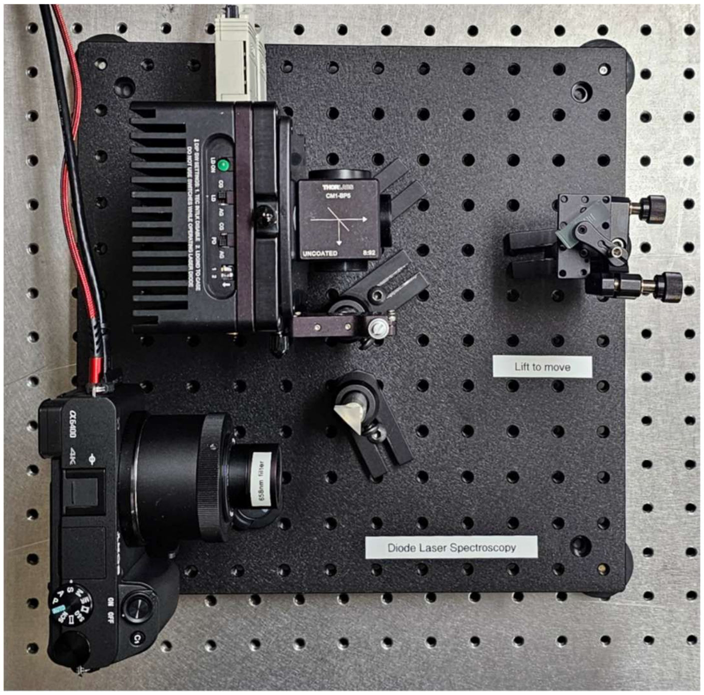

Figure 10. The Diode Laser Spectroscopy hardware in the lab.

corner of the TV, and it works best to set this to no longer than 1/100th of a second, so you needn't worry about camera shake. Also, best to use the remote shutter release. You already know about the ISO setting on this camera, so set the ISO to 500 to begin. Then adjust the shutter time when you want to change the image brightness. As usual, auto-focus and auto-brightness do not work well in the lab, so both are disabled.

To make room for the lens and Grating1 in Figure 9, you may need to lift & place the LIDAR breadboard back on the optical table (where it fits under the shelf) without disconnecting any of its cables. Then place Grating1 about 1.2 meters from the camera, as shown in Figure 9, but without the lens. Adjust the grating so that the m=1 back-reflected light produces a diffuse line of red light on the camera sensor, as shown at the top of Figure 11. To help align the grating, turn the laser current up to 60mA and try sending the prism output beam through a small hole in a white card placed fairly close to the grating. You should then see the back-reflected beam on the other side of the card, and you can send it back through the hole. Once you locate the back-reflected beam this way, it is usually straightforward to get the beam into the camera. Note that the direct (m=0) reflection off the grating in Figure 9 is at about 100 degrees, while the m=1 retroreflected beam can be quite faint.

When you have the alignment right, the m=1 diffracted light should enter the camera, while the much stronger m=0 beam (basically a mirror reflection off the grating surface) goes to a beam blocker. This beam blocker is mostly there to avoid the potential annoyance of unwanted laser beams for people seated nearby. It is generally good practice to keep your laser beams confined to the tabletop as much as possible. Set the laser current to 45mA, reduce the width of the slit, and your laser streak should look something like the top image in Figure 11. Note that nothing is in focus yet; this diffuse streak of light is just an alignment guide.

Once you have a nice diffuse red streak on the TV, add the lens to focus the diffracted beam and create a high-resolution spectrum of the laser output. If you place the lens properly (positioning both

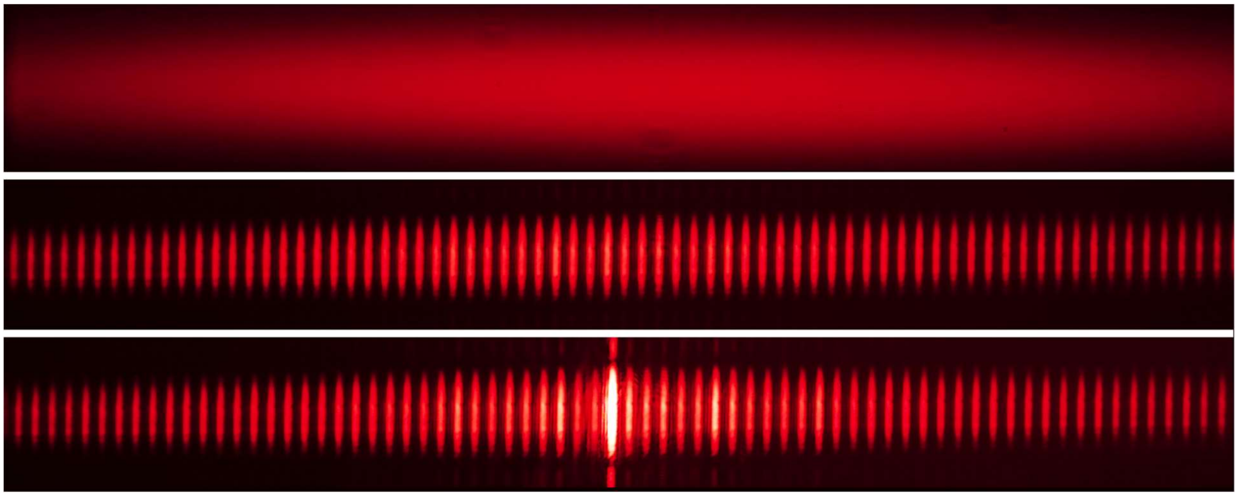

Figure 11. (Top) This image shows the laser spectrum with lens removed, giving just a diffuse streak of red light. (Middle) The lens now focuses the slit onto the camera, revealing a set of bars in the spectrum. The overall broad streak comes from the semiconductor bandgap, while each bar corresponds to an individual cavity mode. The laser current is set below threshold, yielding a uniform array of mode intensities. (Bottom) With the injection current now above threshold, one cavity mode dominates the laser emission. This mode is even more dominant than it looks here because the bright bar saturates the camera sensor.

up/down and left/right), then the new line of laser light should hit the camera sensor at about the same place as it did without the lens. If the lens is properly centered on the beam, then there will be minimal beam deflection. If the lens is not centered on the beam, then there will be a lot of beam deflection.

Finally, move the lens back/forth along the beam path to focus the vertical lines, and adjust the slit width so the lines can be resolved, as illustrated in the middle of Figure 11. If you see some odd blotches of light come and go on the screen, these may come from reflections off the lens surfaces. Tilt the lens a small amount to make these go away. Note that a clean image like that in Figure 11 may not come easily. Getting good results in experimental physics usually involves some tweaking. Adjust the slit width, lens position, and maybe even a bit of lens tilt for best results, and you may need to iterate all these adjustments. Of course, ask your TA for help if needed.

Now let us pause to ponder for a moment, because the middle image in Figure 11 contains a lot of interesting physics to unpack. First of all, the dispersion produced by Grating1 produces an optical spectrum of the laser on the camera. Thus, what you are looking at is laser power as a function of wavelength. The horizontal scale is not precisely linear, but it is roughly linear because $\delta\lambda/\lambda \ll 1$.

Analyzing this optical spectrum, the overall streak of light shows the laser acting like an LED, as depicted by the top curve in Figure 6. With the current set below threshold, there is very little laser action taking place, and the overall width of the streak represents a spread of about 2 nm centered at about 658 nm.

Imprinted on this broad streak are numerous vertical lines, and these correspond to the cavity modes depicted by the second curve in Figure 6. Because the optically active semiconductor material is enclosed by an optical cavity, the cavity modes are excited by optical emission inside the cavity. So far, this is just classical physics; the situation is similar to what happens when you generate acoustic noise inside an organ pipe – the resonant modes of the pipe are strongly excited, producing a series of enhanced tones. In the laser cavity, the LED-like emission similarly excites the resonant optical modes of the internal semiconductor cavity, which are equally spaced in frequency.

This excitation of resonant modes is a ubiquitous phenomenon in physical systems. When a strong wind blows on a playground swing, it excites the resonant motion of the swing, even though the wind gusts are certainly not tuned to that frequency. The same thing is happening inside the diode laser. The emission from electron-hole recombination is not tuned to the cavity modes, but these high-Q resonant modes are preferentially excited. Injecting noise into a resonator is a subtle concept, not usually discussed in undergraduate courses. But it shows up frequently in many physical systems.

If you turn the laser power slightly above threshold, you should see a spectrum that look something like the bottom image in Figure 11. Here a single mode dominates the spectrum, indicating the onset of laser behavior. As you turn the laser power up/down, you will probably see the laser running single-mode (or close to it) at certain currents, but multi-mode at other currents, and mode hops may occur as well. Usually the laser spectrum is somewhat erratic and unpredictable at this point.

If you turn the laser current up and down and look closely, you will also see the positions of all the modes shift slightly with the current. Place an “arrow sticker” (located near the TV) on the TV screen to serve as a marker, and then you can see the individual cavity modes moving to the right as you turn the current up. (Please do **not** write on the TV screen with any kind of marker!) This coordinated frequency shift mainly comes from an increase in the laser temperature as you turn the current up. The temperature controller only regulates the temperature of the laser can, so the actual semiconductor temperature will increase as you send more current into the device. Temperature changes affect both the cavity length and the index of refraction of the laser medium, thus changing the frequencies of the cavity modes.

The main take-away message from these spectra is that the overall behavior of the laser spectrum is essentially what you expect from the top two curves in Figure 6. The broad spectral envelope comes from the intrinsic semiconductor properties, and the bars reveal the resonant modes of the laser’s internal optical cavity. Also, the concentration of optical energy in a single mode arises from stimulated emission and the mode-selection instability it produces, as described above. Qualitatively, at least, what you observe makes some amount of sense.

Exercise 4. Produce three images corresponding those shown in Figure 11 and add them to your e-notebook.

Exercise 5. Calculate the internal-cavity mode spacing from your images and express your results as a frequency width $\delta\nu$ in GHz and the corresponding $\delta\lambda$ in nm. Provide an uncertainty estimate with your measurement. To do this analysis, you will need several pieces of information, plus you will need to make some simplifying approximations:

1) The wavelength of the laser light is about 658 nm, and the spectrum covers a narrow range with $\delta\lambda \ll \lambda$.
2) For your calculation, use an approximation of the optical layout that looks like that shown in Figure 12. Assuming perfect back-reflected light from the grating simplifies the calculation without adding much error to the resulting $\delta\nu$. This geometry is called the Littrow configuration.

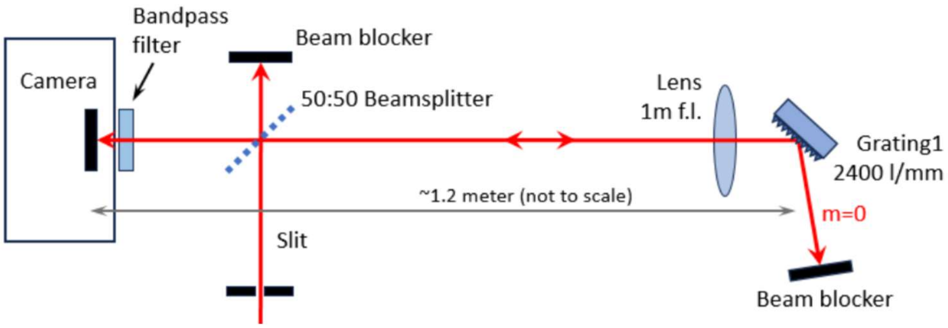

Figure 12. When analyzing your spectrum, use this approximation of the optical layout in Figure 9. Making the input and output grating angles are nearly identical simplifies the calculation with little loss of accuracy.

3) Calculate the grating angle (relative to the normal vector perpendicular to the grating surface) in this case. Start with the grating equation: $d(sin\,\theta_{in} + sin\,\theta_{out}) = m\lambda$, where d is the spacing between the grooves of the grating, θ_{in} and θ_{out} are the input and output angles, and $m = 1$ for the case shown in Figure 9 and Figure 12. Call this angle $\theta_{Littrow}$.
4) Next you want to know the dispersion of the laser light by the grating. In Figure 12 you can see that $\theta_{in} = \theta_{Littrow}$, and this angle does not depend on λ… this is just the angle of the laser hitting the grating. Given this fact, differentiate the grating equation to get $d\theta/d\lambda$, the dispersion of the outgoing laser light from the grating.
5) Next use the fact that the camera sensor is 25.1 mm wide and note that $dx = Ld\theta$ on the camera sensor, where L is the focal length of the lens. (You might think that L should be the distance from the camera sensor to the grating, but it's not. Why? If you draw a thin-lens ray diagram for this optical setup, you will see why.)

Exercise 6. What is the approximate width of the semiconductor-medium curve in Figure 6, in nm? (It is equal to the width of the overall streaks of light in Figure 11.)

Exercise 7. Assuming a typical value of $n = 3.5$ for the index of refraction of the AlGaInP semiconductor in the HL6501MG, use your measured $\delta\nu$ to give you the length of the internal optical cavity.

Exercise 8. As a final exercise in this set, observe how the optical spectrum changes with laser temperature. Set the injection current to 50mA and record a spectrum for temperatures ranging from 17C to 26C in steps of 1C. Extract strips from each image and stack them into a single composite image made from all ten of your images. For example, Figure 13 illustrates a composite made from just two spectra. Comparing your composite image with Figure 7, you can see this laser's mode-hopping behavior as a function of temperature. Note that the bright bars in Figure 13 are highly *saturated* in the image. These modes typically have a *lot* more optical power than the fainter modes.
If you spend more time on this and study the spectra carefully, you can measure that the cavity modes shift slowly with temperature, with an individual mode moving at a rate of about 0.05nm/C (one mode width in 1.4C). Meanwhile the semiconductor medium peak shifts at about 0.15nm/C. As

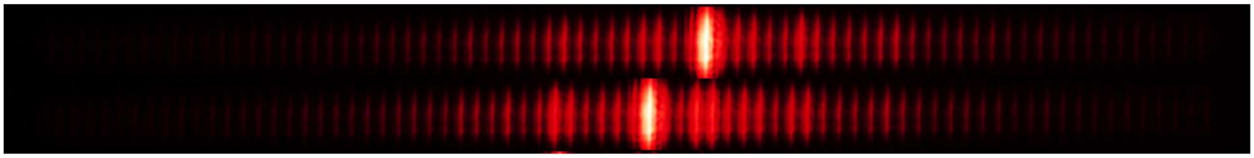

Figure 13. A series of two laser spectra taken at different temperatures. As in Figure 11, the bright modes are actually more dominant than they appear, because of saturation of the camera sensor. The dominant mode "hops" somewhat erratically with temperature, making it difficult to use bare diode lasers in spectroscopy applications where one wants a continuously tunable monochromatic light source.

described in connection with Figure 6, this difference is what causes the mode-hopping behavior you see in your composite image.

Imagine now that you would like to have a laser that excites a specific atomic transition, as we do in the Rb spectroscopy lab in the Atomic Track. The set of spectra you just took illustrates how this might not be possible using a bare diode laser. The mode-hopping behavior is quite unpredictable, and the laser might not operate at the wavelength you need. Adding an external grating provides a viable solution to this problem, as you will now demonstrate.

Making a frequency-tunable laser

Your next step is to uncover Grating2 to realize the full optical set up shown in Figure 9. The trickiest part now is to align the grating angle so that the beam retroreflects back into the laser. Because the laser channel is just a few microns in cross-section, this is a small target to hit, making for a challenging alignment task. As you did with Grating1, you might try inserting a white card with a small hole into the beam and look for the retroreflected beam on the other side of the card. This will not be sufficient to hit the laser channel, but at least it will get you headed in the right direction.

When you think you are getting close, set the laser current to about 45mA, right around threshold, and view the m=0 beam coming off Grating2 with a white card. If you slowly adjust the grating up/down knob, you should see the beam become brighter when the retroreflected beam is going back into the semiconductor laser. The brightening can be quite subtle, so it may take some effort to find. Try looking at different laser currents, but always near threshold. Once you see the brightening, watch the laser spectrum on the TV when you adjust the grating up/down knob. You will probably see a qualitative change in the spectrum when the retroreflection is correct, and you can use this to fine-tune the up/down angle further.

After optimizing the up/down angle, you should then be able to tune the laser frequency using only the grating left/right angle, with the laser running single-mode most of the time. This will be immediately obvious when you have it. At this point, congratulations … you now have a frequency-tunable, grating-stabilized, single-mode semiconductor laser. You have taken the erratic semiconductor behavior of a bare diode laser (random mode hops, running multi-mode) and turned it into a precision spectroscopy tool.

Exercise 9. With the external grating now selecting any cavity mode on demand, place an arrow sticker at some random place on the spectrum (not on an existing mode), and imagine that this is the precise laser wavelength you must have. You should be able to adjust the grating left/right angle and

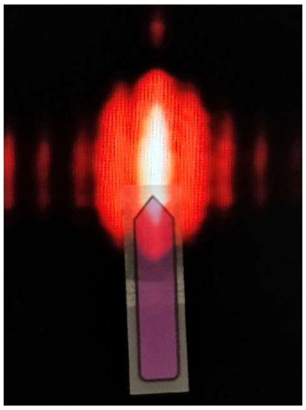

Figure 14. By adjusting the external cavity left/right angle and the laser current, you should be able to coax the laser to operate at any predetermined wavelength, as long as it is within the reach of the semiconductor medium.

the laser current to produce single-mode operation exactly at your randomly chosen wavelength (as close to the tip of the arrow as you can see on the screen, anyway). Document this by taking a close-up photo that shows your arrow sticker and its accompanying laser mode, as shown in Figure 14. (If you happened to pick an easy spot, move your arrow sticker to make the task more challenging.) As you can see, adding the external grating now allows you to tune your laser to any desired wavelength within the semiconductor range. Nice!

Once again, the main take-away message is that the laser behavior generally follows expectations from the qualitative model shown in Figure 6. If we had greater spectral resolution, you could go on to examine the external cavity modes shown in the last curve in Figure 6. But that is left for some future date, should you continue working in laser spectroscopy or some related field. Because ECDLs are so inexpensive and versatile, they have found many applications and are commercially available at many wavelengths. So you might well encounter this technology in your future endeavors.

Part 2. LIDAR

As illustrated in Figure 15, the basic physics of LIght Detection And Ranging (LIDAR) is simple: send out a pulse of laser light and observed the pulse of scattered light that returns. Then use the known speed of light (30 cm/nsec) to measure the distance to whatever scattered the laser photons. This basic time-of-flight concept has been applied to a remarkably diverse array of remote-sensing tasks, including: 1) Improving photography by measuring subject distances (LIDAR is built into many smartphones); 2) Testing General Relativity by measuring the distance to the Moon (specifically, the distance to a retroreflecting mirror placed on the surface by the Apollo astronauts) to an absolute accuracy of a few millimeters as a function of time; 3) Obstacle detection and avoidance for autonomous vehicles, including agricultural robots; 4) Earth surface topography; 5) Atmospheric monitoring; 6) Remote sensing on planetary spacecraft … and a great many other applications.

One way to think about LIDAR is that it expands imaging to the third dimension. Cameras produce 2D images with little information about the object distance. One can get some depth

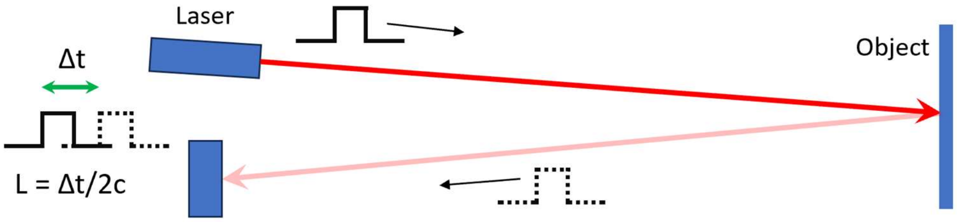

Figure 15. The basic operation of time-of-flight LIDAR. A laser sends out a pulse of light and the (much weaker) scattered pulse is detected. The distance to the object can then be calculated from the pulse travel time.

perception from stereoscopic techniques (analyzing images taken from different camera positions) or from focusing (analyzing images taken with different camera focus), but these methods are inaccurate unless the object is quite close to the camera. LIDAR is capable of measuring even large distances with extremely high accuracy, as long as the return pulses are bright enough to be detected. Imaging LIDAR is often done by sweeping the laser direction (both horizontally and vertically), but this method is being replaced by the use of high-speed array detectors in many high-end applications.

LIDAR technology relies on high-power, high-speed lasers (especially semiconductor diode lasers), along with sensitive, high-speed detectors (including detector arrays) and fast integrated electronic processing. Obtaining centimeter-scale accuracy means operating at GHz frequencies, and new imaging detectors are appearing with megapixel arrays and close to single-photon sensitivity for each pixel. Even phased-array detectors (observing both the intensity and phase of the incoming light) are being demonstrated in research labs, promising another quantum leap in functionality. As the underlying technology advances and becomes cheaper, LIDAR is rapidly finding new applications in diverse areas.

In Ph77, we will not be exploring these cutting-edge (a.k.a. expensive, time-consuming) imaging LIDAR technologies but will stick with a single pulsed laser and a single high-speed photodetector. Our focus is on a demonstration where you will use a LIDAR setup to measure the distance from the Ph77 lab in East Bridge to the nearest exterior wall of West Bridge. The overarching goal of this exercise is to give you some experience working with the hardware, and to examine how one makes a model of the LIDAR signal and (more importantly) the signal-to-noise ratio (SNR).

Exercise 10. Your first task in any experimental venture is to estimate the signal and determine the feasibility of the measurement. Building hardware is expensive and laborious, so you want to know what you are getting into. Before buying a lot of equipment, you need to convince yourself (and the funding agencies) that you have a reasonable plan with a well-defined goal. This kind of model building is a process, so this exercise involves a chain of steps.

1) Your goal is to measure the distance from the Ph77 optical table to the nearest wall of West Bridge, so what is that distance? This is easily estimated using Google Maps (right-click on a map to measure a distance), so do that and report your results. For now, just measure the distance from the west wall of East Bridge to the east wall of West Bridge, as these both appear on the map (Figure 16). Be sure to include an uncertainty estimate in your measurement.

2) What is the light travel time between these two walls (in nsec) ... again with an uncertainty estimate?

3) LIDAR begins with a laser pulse out. What is the peak laser power in this pulse? For this, consult the laser specifications in Appendix I below. (Note that eye safety is mainly determined by the time-averaged power, which is quite low with these rapid pulses.

4) Next this pulse hits the wall of West Bridge. What fraction of this light scatters off the wall? This is difficult to estimate because the light scatters into a large solid angle. Nevertheless, this is an important number in your model, so don't make a wild guess. Look up the albedo of various surfaces online and look into the properties of diffuse scattering; then make an educated guess. Express your answer as a reflection coefficient per steradian. Because you do not know the condition of the wall surface, this is probably the least accurate estimate in this exercise.

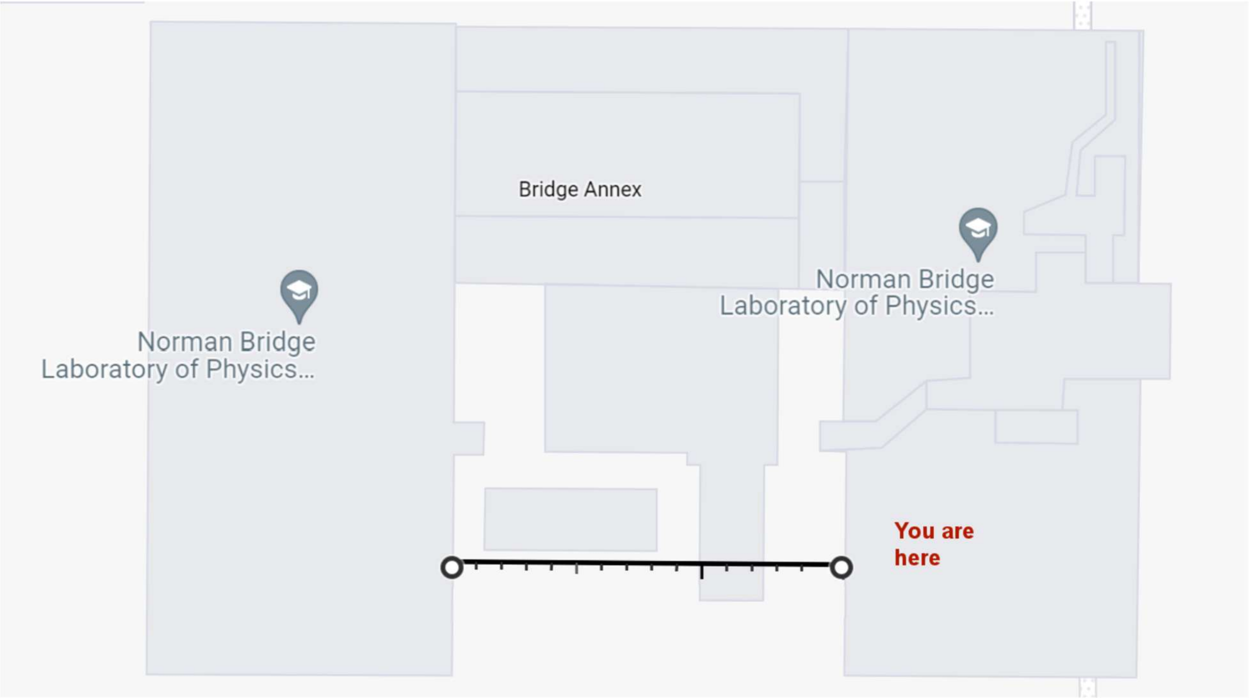

Figure 16. Your goal in this laboratory exercise is to use LIDAR to measure the distance from the Ph77 lab in East Bridge to the nearest exterior wall of West Bridge.

5) From this, what is the peak laser power entering the main lens of the detector system (in nW)? Hint: You will be using a 500mm f/8 telephoto lens for this measurement. Calculate the lens aperture from this information.

6) What fraction of this power lands on the photodetector (PD)? To make this estimate: a) Calculate the size of the laser spot on the wall of West Bridge, using the laser spot size and divergence given in the specifications; b) Use the thin-lens approximation to calculate the size of the image of the laser spot on the focal plane of the main lens; c) compare this with the detector size given in the specifications.

7) With the detector amplification set to maximum, what is the peak output voltage from the laser pulse? How does this compare with the "integrated noise" listed on the detector spec sheet?

8) Because the detector noise is greater than the signal, you would not see the signal from a single trace on the oscilloscope. How many traces would you need to average to obtain a SNR (signal-to-noise ratio) of 5? [The good news here is that this experiment operates at MHz frequencies, so averaging a few thousand traces takes just a fraction of a second.]

These kinds of SNR calculations are a staple in any field of experimental physics because this is how you decide what experiments are feasible. In large-scale experimental projects, where a lot of time and money is being spent on hardware and data analysis, people often spend years developing complex computational models along these lines. When you build things, planning is a big part of the process.

Laboratory Exercises

If you moved the LIDAR apparatus to the back of the optical table, first move it to the front again so it is easily serviceable. Again, some of the optical and electronic elements have been locked in place for you (to save time), so it is best to leave these alone. For best results, adjust the leveling screws on the breadboard so they are as short as possible, as this is a good starting point.

Figure 17 shows the overall layout of the LIDAR apparatus. To complete this setup you should remove the 658nm filter from the camara (and place it inside its box), and install the camera in the post holder provided (but note that the beamsplitter cube is probably blocked with a plastic cover, so remove this before installing the camera). Adjust the camera placement so the tubes line up with about a 1mm gap between them, as shown in the image on the right. Note the gap between the two tubes in the image. This gap lets unwanted light in, so cover this gap with the small "sleeve" tube that should be on the beamsplitter side. If the sleeve is missing, ask.

Next turn on the camera, remove the lens cap on the 500mm lens, and switch to aperture priority. If you rotate the breadboard a bit, you should be able to focus on the tree leaves outside the window. This is a telephoto lens, so you will not be able to focus nearby. You might have to tilt the assembly up using the "feet" on the right side of the breadboard, and of course the blinds will have to be up for you to see the leaves on the tree outside. Also, you might find a thin lens (labeled *Close Up +1 67mm*) on the front of the telephoto lens; if so, remove this extra lens for now and place it in its storage box on the breadboard.

Once you see the tree, adjust the breadboard feet so the camera points toward the mirror sitting on the window sill. If you set the lens to its minimum focus distance (10m), you should be able to see an out-of-focus hand on the monitor if you wave your hand around near the mirror.

Next turn on the 640nm laser (Thorlabs NPL64B) by turning the key, and you should see a flashing purple light. There is a sliding cover on the front of the laser, so make sure that it is in the open (up)

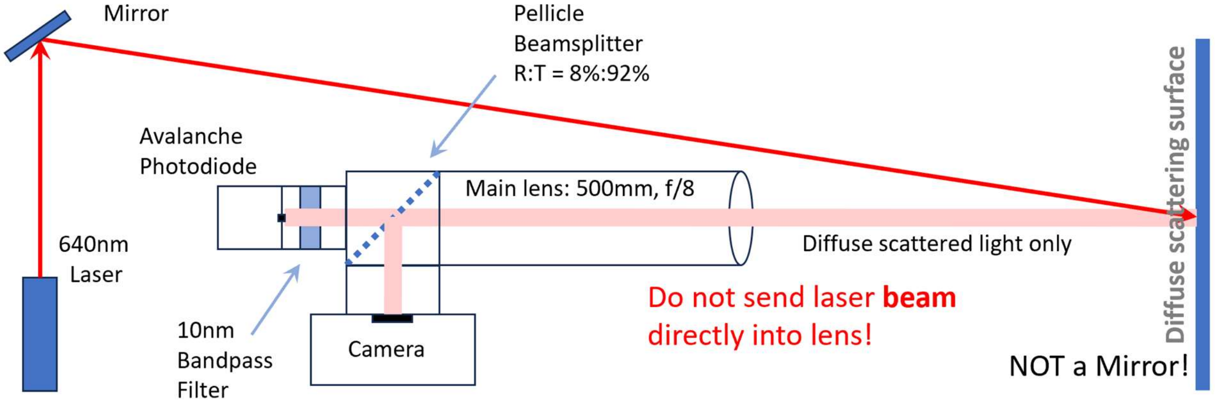

Figure 17. The basic layout of the LIDAR apparatus. Laser pulses are produced by the 640nm diode laser, and these pulses strike a far-away scattering surface. Some light then scatters off the surface, enters the lens, and is detected by a high-sensitivity Avalanche Photodiode. The 10nm bandpass filter allows 640nm light to reach the detector while rejecting broadband ambient light.

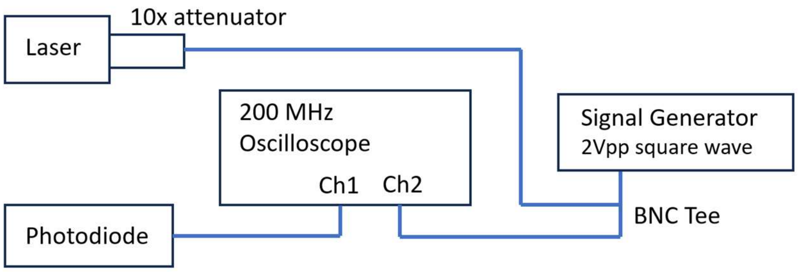

Figure 18. The connection diagram for the LIDAR apparatus. High-speed electronics is needed for LIDAR, so we use a Siglent SDG2042X signal generator that goes up to 40 MHz and a Keysight DSO1202A oscilloscope that goes to 200 MHz.

position. The laser needs an electronic trigger signal to operate, so there will not be any laser light yet. To supply this trigger, connect the thin SMA cable to the signal generator (Siglent SDG2042X) following the connection diagram in Figure 18. Send a 2Vpp square-wave signal at 5 MHz to the laser and oscilloscope as shown, and trigger off ch2. You should see the signal on the ‘scope and some red laser light coming out of the laser. The laser has been configured (on its back panel) to produce pulses with maximum duration when an external trigger is applied, and you can leave these settings alone.

Next turn on the avalanche photodiode (PD) by turning on the power supply block (labeled; a green LED lights up) and setting the 0/1 switch to 1 on the photodiode housing. Now you should see two green LEDs – one on the PD power supply and another on the PD housing. Avalanche photodiodes are much more sensitive than normal photodiodes (and you can read about how they work online if you are interested), and this is needed here because LIDAR works by detecting a small amount of scattered light. You will want to leave the PD set at maximum gain. **Never send laser light directly into the photodiode**, as this can damage it. Scattered light only.

Now that the laser and photodiode are functioning, place the +1 diopter close-up lens on the front of the 500mm lens (again, labeled *Close Up +1 67mm*.) The “diopter” term refers to the focal length of the lens… one diopter means $f = 1$ meter, two diopters means $f = 0.5$ meter, etc. (This notation is most commonly used with eyeglasses.) With the +1 diopter lens in place, setting the telephoto lens distance to infinity means the focal distance will actually be 1 meter. (Think about it. We covered this kind of thing at the beginning of the Optics Track. Make a thin-lens diagram if you really want to understand it.)

Next place a white paper card about 25 cm in front of the 500mm lens, using this for the scattering surface in Figure 17. Steer the laser beam onto this card and you should see a splattering of red light on the TV. One you see this, you should be able to adjust the ‘scope settings to see some scattered-light pulses on the ‘scope, as illustrated in Figure 19. If you do not see this signal, try some of these:

1) Is the camera looking toward the laser spot on the white card? The lens will not focus this close, but you can wave your fingers around to see where you should place the card. The scattered red light should nearly fill the monitor, but it will not be super bright because the focus is way off.
2) Is the ch2 square-wave signal clearly visible on the ‘scope? If not, this is probably a triggering problem.
3) The photodiode signal will be weak, so increase the ch1 to 1mV/division (but not 500µV/division).

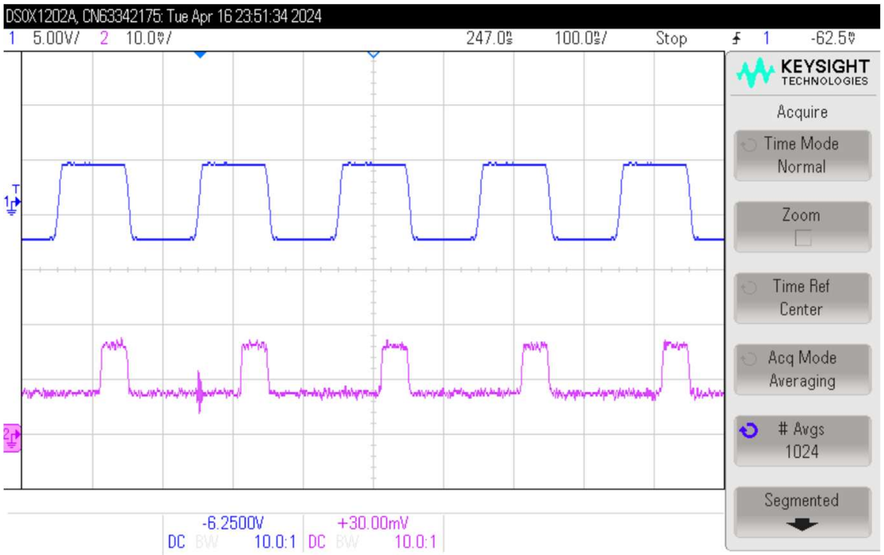
DSOX1202A, CN63342175: Tue Apr 16 23:51:34 2024

1 5.00V/ 2 10.0mV/ 247.0ns 100.0ns/ Stop 1 -62.5mV

Time Mode Normal

Zoom ☐

Time Ref Center

Acq Mode Averaging

Avgs 1024

-6.2500V DC BW 10.0:1

+30.00mV DC BW 10.0:1

Figure 19. A 'scope screenshot showing the 5MHz input signal (top) and the resulting scattered-light pulses coming from the photodiode (bottom). (There is also a slight glitch at the trigger point… probably a 'scope artifact.)

4) Average 128 traces to reduce the noise. With a 5 MHz signal you can average quite a few traces in a short amount of time.
5) Make sure lens 500mm lens aperture is full open (O ↔ P should be set to O)

Exercise 11. Once you see some signal, optimize it by moving the card around, changing the laser pointing, and switch to averaging 1024 traces. You should be able to produce some clean, flat-topped pulses like those shown in Figure 19. Save a nice screenshot using: Run/Stop button, insert a USB drive, hit Save-Recall, Format = PNG (top menu), Save to USB… and add it to your e-notebook.

Exercise 12. Note that the 5MHz square wave (ch2) transitions are not clean and vertical but are sloped. Measure the rise time of the ch2 transitions and add another screenshot to your e-notebook showing the measurement using the 'scope cursors. The signal generator lists a sampling rate of 1.2 GSamples/sec (as printed on the front of the signal generator). So how many sample points make up one rise time?

Exercise 13. When you have this signal on the screen, try changing the 'scope gain from 1mV/div to 500µV/div. You will see the signal get larger on the screen, but the pulses will now be rounded, as shown in Figure 20. (You need some clean, flat-topped pulses at 1mV/div to see this effect well.) The reason for this rounding is that the 500µV/div setting is mostly a software enhancement; it smooths the signal using temporal averaging to reduce the noise, which is not the same as doing a clean electronic gain change. Reproduce Figure 20 for yourself by saving one reference trace (using the Analyze/Features menu on the 'scope) and add a screenshot to your e-notebook. Because you will be wanting sharp pulses to measuring time delays, you should **not** use the 500µV/div setting for the remainder of this lab.

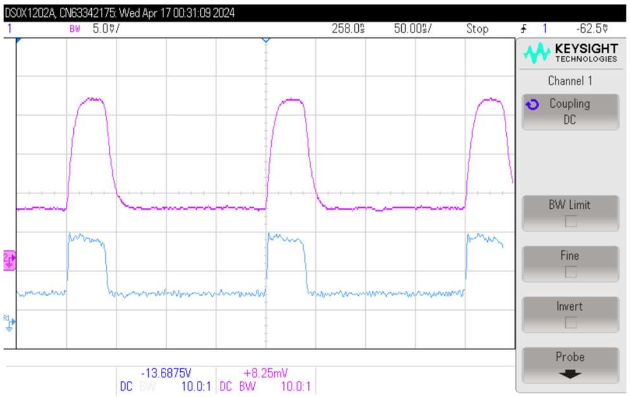

Figure 20. A set of clean, flat-topped pulses at 10mV/div (lower trace) along with the same pulses at 5mV/div (using the same trigger point). The 'scope applies some temporal averaging at 5mV/div that substantially distorts the intrinsic pulse shape.

Going back to Figure 19, note that there is a phase delay between the ch1 and ch2 traces, with the scattered-light pulse lagging the laser trigger (change the signal generator from 5 MHz to 1 MHz to be certain which signal is lagging which). How much delay do you expect from the light travel time, and how much delay do you measure on the 'scope? You can conclude most of this delay is from the electronics, with only a small portion coming from the LIDAR time-of-flight delay.

Exercise 14. Deploy the swing-out bar provided on the optical table and align the bar angle and breadboard so the camera lens points straight down the length of the bar. Move the white card so it sits at the edge of the optical table, adjust the optics to get a photodiode signal, and measure the pulse delay on the 'scope (from the half-max of the ch2 pulse to the half-max of the ch1 pulse). Zoom in on the delay part of the signal on the 'scope and you should be able to measure the delay time to an accuracy of a few tenths of a nanosecond. Record this delay time, giving you essentially just the electronic delay in the system.

Without changing anything else, move the card to the end of the swing-out bar, tweak the laser mirror slightly, and again measure the delay time to high accuracy. Use a tape measure to measure the distance between the two white-card placements and compare your measurements with what you expect for the change in light-travel time. If your LIDAR answer does not fit expectations, redo both measurements. Unless the speed of light has changed since this writing, your results should agree with theory, up to measurement uncertainties.

Exercise 15. With the +1 diopter close-up lens on the 500mm lens, set the lens focal distance to infinity and move the card so it is about 1 meter from the lens. Now you should be able to move the card along the bar until the scattered light comes into focus on the monitor. You should see a bright spot, often a double-lobed spot because the diode laser does not produce a clean, round Gaussian beam. Move the spot around on the monitor to maximize the signal, now yielding much larger PD pulses, typically about 200 mV on the 'scope. Take another screenshot and add it to your e-notebook.

This last exercise is also useful for connecting the photodiode signal to the TV image. Referring back to Figure 17, you can see that the photodiode sensor and the camera sensor sit at "conjugate" image planes. That means they are essentially the same image plane, but in different locations. And these planes are connected via their placements relative to the beamsplitter.

Maximize the PD signal (by pointing the laser beam on the card while you watch the 'scope signal) and then use an arrow sticker (next to the TV) to mark the position of the laser spot on the TV screen. By virtue of the two conjugate image planes, this pointer also marks the position of the PD sensor. If you move the laser spot away from the PD position by more than a couple of beam widths, the signal quickly goes to zero. This means that the laser spot size is comparable to size of the photodiode active area (with a 200 μm diameter, as listed in the spec sheet). And the main takeaway from this discussion is that we can now use the TV image to easily focus the scattered laser spot onto the tiny photodetector.

Your next step is to measure a longer distance while still staying inside the Ph77 lab (one step at a time). First return the swing-out bar to its closed position (to get it out of people's way) and remove the +1 diopter close-up lens. Move the breadboard so the lens points toward the mirror sitting on the windowsill on the west wall of the lab (if the mirror is missing, ask) and point the laser at this mirror as well. **Do not point the mirror so the laser goes directly into the lens**. Instead, use the mirror to send the laser beam to a spot just above the door to the room. Make sure the signal generator is set to 5 MHz, so the laser beam is easily visible.

Align everything so you see the laser spot on the TV screen. The optical path is then laser → mirror → wall above the door… then the scattered light goes from the wall above the door → mirror → camera. Only scattered light should go back into the lens… no direct reflections! Use the TV image to position the laser spot on the photodetector, at the conjugate position of the photodiode. This should focus the scattered light onto the photodiode.

If all is well, then you should see a small PD signal on the 'scope. If not, make sure the 'scope is set to 1 mV/div with averaging, because the signal is not going to be very strong. If you are having no luck, go back to the card on the swing-out arm, because that gives you a much larger signal.

When you do locate the signal, try moving the breadboard a bit. You will see that the laser spot does not move on the TV screen, although it does move on the ceiling. (Think about it.) Just for consistency, make sure the laser spot hits the *east wall of the lab above the door.*

Exercise 16. Do another LIDAR distance measurement and verify the light-travel distance (using an 8-meter tape measure in the lab). Compare the pulse delay with what you got when the card was at the edge of the table, as you did before. Add a relevant screenshot to your e-notebook. Again, assuming the speed of light hasn't changed recently, your distance and light-travel-time measurements should agree. Include measurement uncertainties on both.

Okay, now you are ready for the final challenge of measuring the distance to the east wall of West Bridge. Finding the signal is substantially more difficult this time, but that is why you worked up to it in steps. This measurement should **not** be attempted in the morning because the laser spot can be exceedingly hard to find when the morning sun is shining on the wall. The procedure is about the same as what you have just done, except be warned that the TV and PD signals will both be quite small.

To get set up, open the window wide (just the right side of window; leave left side closed and bolted) and point the camera out the window. You should be able to focus on the opposite wall or the trees between the two buildings. Raise one side of breadboard using the leveling screws so the treetops appear at the bottom of the TV screen. The branches are useful as a reference, so it helps to have them on the bottom of the field of view.

Next point the laser so the beam spot appears on the TV (this is the hard part). You should be able to see the laser spot on the wall with your naked eyes, but we also have binoculars in the lab to make this easier. Once you have the red beam on the TV, place it at the known PD focus spot and move it around to optimize the PD signal. If all is well, you should see a PD signal of around 1-10 mV on the ‘scope.

Exercise 17. Measure the pulse delay and compare with your estimate from Exercise 10, again with measurement uncertainties. (To avoid measurement ambiguities, you might want to switch to a lower pulse frequency. Why?)

Exercise 18. (Optional) If you are really feeling inspired, go outside and measure the distance directly using the 30-meter tape measure available in the lab (with a partner, this only takes 10 minutes). Obviously, your direct measurement should agree with the LIDAR measurement... but does it? ... and to what accuracy?

Appendix I: LIDAR Component Specifications

This table lists several important specifications for the NPL64B Nanosecond Pulsed Laser Diode System running at 640nm. It was taken from the Thorlabs online catalog.

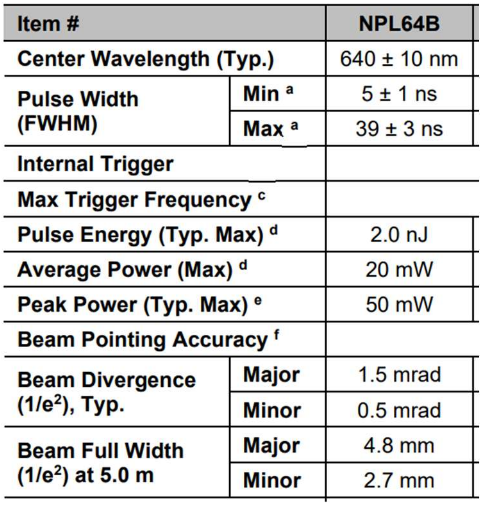

Item #		NPL64B
Center Wavelength (Typ.)		640 ± 10 nm
Pulse Width (FWHM)	Min [a]	5 ± 1 ns
	Max [a]	39 ± 3 ns
Internal Trigger		
Max Trigger Frequency [c]		
Pulse Energy (Typ. Max) [d]		2.0 nJ
Average Power (Max) [d]		20 mW
Peak Power (Typ. Max) [e]		50 mW
Beam Pointing Accuracy [f]		
Beam Divergence ($1/e^2$), Typ.	Major	1.5 mrad
	Minor	0.5 mrad
Beam Full Width ($1/e^2$) at 5.0 m	Major	4.8 mm
	Minor	2.7 mm

And here is the spec sheet for our APD430A2 Avalanche Detector (central column in the table)

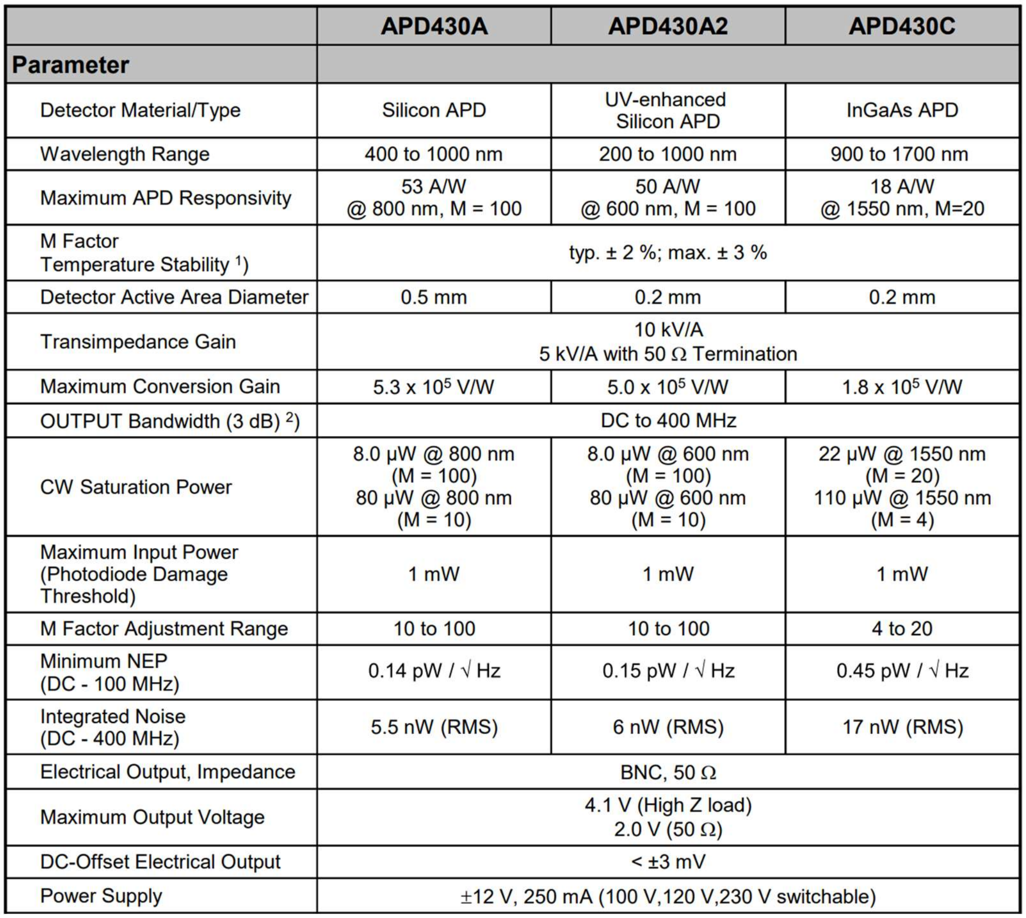

	APD430A	APD430A2	APD430C
Parameter			
Detector Material/Type	Silicon APD	UV-enhanced Silicon APD	InGaAs APD
Wavelength Range	400 to 1000 nm	200 to 1000 nm	900 to 1700 nm
Maximum APD Responsivity	53 A/W @ 800 nm, M = 100	50 A/W @ 600 nm, M = 100	18 A/W @ 1550 nm, M=20
M Factor Temperature Stability [1])	typ. ± 2 %; max. ± 3 %		
Detector Active Area Diameter	0.5 mm	0.2 mm	0.2 mm
Transimpedance Gain	10 kV/A 5 kV/A with 50 Ω Termination		
Maximum Conversion Gain	5.3×10^5 V/W	5.0×10^5 V/W	1.8×10^5 V/W
OUTPUT Bandwidth (3 dB) [2])	DC to 400 MHz		
CW Saturation Power	8.0 µW @ 800 nm (M = 100) 80 µW @ 800 nm (M = 10)	8.0 µW @ 600 nm (M = 100) 80 µW @ 600 nm (M = 10)	22 µW @ 1550 nm (M = 20) 110 µW @ 1550 nm (M = 4)
Maximum Input Power (Photodiode Damage Threshold)	1 mW	1 mW	1 mW
M Factor Adjustment Range	10 to 100	10 to 100	4 to 20
Minimum NEP (DC - 100 MHz)	0.14 pW / √ Hz	0.15 pW / √ Hz	0.45 pW / √ Hz
Integrated Noise (DC - 400 MHz)	5.5 nW (RMS)	6 nW (RMS)	17 nW (RMS)
Electrical Output, Impedance	BNC, 50 Ω		
Maximum Output Voltage	4.1 V (High Z load) 2.0 V (50 Ω)		
DC-Offset Electrical Output	< ±3 mV		
Power Supply	±12 V, 250 mA (100 V,120 V,230 V switchable)		