OPTIX Module 3 – Intermediate
Laser Diodes and Diode Lasers

Michaela Kleinert

1 Objectives:

In this module you will learn about

- laser diodes and how changing the temperature of the diode and the current through it affects its wavelength;
- what mode hops are, and how you can recognize and avoid them;
- the difference between a free-running laser diode and an extended-cavity diode laser.

Note that you have to take this module before you can take MODULE 3 - ADVANCED to investigate the spectrum of rubidium since that module relies heavily on the basics we'll cover in this module.

Use this manual as you work through the module to keep track of your notes and thoughts. In addition, you will have to add a few printouts or add additional sheets of paper containing data tables, sketches, or additional notes. Please note that this is *not* your lab report and that you are expected to complete a full lab report written in *LaTeX* after you have completed this module. Keep in mind that this should be a publication style report, which means that it should place a big emphasis on your data and data analysis, and not so much on all the nitty-gritty details of how to assemble the apparatus. Keep this in mind as you work through this module. Your instructor can provide you with more information and will send you a template file that you can use for your report.

2 Tests and assessment:

In preparation for this module, read through the whole manual and answer the questions that are marked with a *. You should also watch the VIDEOs that are posted on our website (www.willamette.edu/cla/physics/info/NSF-OPTIX). They are meant to accompany this manual and will show you some critical steps of the module. When you come to lab, be prepared to discuss your answers to these questions with your classmates and your instructor.

In order to assess the success of this module, you will take a short assessment test before you start ("pre-assessment"), and another one after you have successfully completed this module ("post-assessment"). At this point you will also have the opportunity to provide us with feedback about the module that we will use to improve it for the next student generation. Thank you for your support!

3 Equipment:

You can find most of the equipment for this module in the box labeled 'Module 3' that is located in one of the cabinets in the OPTIX lab. Standard equipment that is used for multiple modules will be located in the cabinets in the OPTIX lab. Please feel free to also ask your instructor for help. You will need:

- a Thorlabs laser diode driver with current driver and temperature controller (LTC100-B);
- laser diodes centered at 780 nm (L780P010);
- a collimation lens for laser diode;
- a diffraction grating attached to a PZT and a mirror mount;
- silver coated mirrors and mirror mounts;
- the *Mightex* IR1 spectrometer.

4 Required background knowledge and things you need to do before starting this module:

In this module you will learn how to work with laser diodes. Review the *Laser Safety Material* before coming to lab; there will be a brief test about laser safety before you will be allowed to start this module. Unlike the HeNe laser that you used in Module 1, in this module you will use a laser that emits light at 780 nm, which is in the near-IR. Remember that lasers that emit light in the invisible range are more dangerous to you than visible lasers of the same power because they circumvent your blinking reflex. And while 780 nm is just at the edge of the visible range – in fact you can still see it when you place a piece of paper into the laser beam –, make sure to **always wear your laser goggles unless your instructor tells you that it is safe to take them off**. [Your instructor will show you the beam without goggles, but after that demonstration, please wear your goggles at all times.] Note that the goggles you need to use for this lab are different from the ones you used in Module 1 (you can easily distinguish them by their color: The goggles in this module are yellow-green, while the ones in Module 1 were blue.) The correct goggles should be located in the blue box labeled 'Module 3'. Besides general laser safety you don't need to know anything about diode lasers; you will learn all about them in this module.

Make sure to read through the *whole* manual before coming to the OPTIX lab, and mark everything that you find difficult to understand. In addition, **work through all the boxes marked with a yellow star like this: *. Be prepared to discuss your completed work with the group and your instructor before starting the lab. This counts as your pre-lab.** During your lab time, you can work through parts of the manual with your lab group and instructor, and you will get a lot of hands-on experience.

If you damage anything, please tell your instructor *immediately*. While we of course try to avoid damaging the equipment since it is rather pricey, especially in this lab, accidents can happen. Please don't try to cover up any damage; that only makes life harder for us as we try to figure out why something is no longer working. Tell us exactly what happened, and we can either fix it or replace the equipment quickly.

*** Lastly, write down everything related to lasers and laser diodes you remember from the Modern Physics or from another previous classes (any relations, sketches, key words that pop into your mind). If you can connect them in a meaningful way, even better! And now - have fun in the lab!**

5 Introduction to lasers:

In its very basic form, a laser consists of three components: The medium, the pump, and the resonant cavity.

The **medium** is the material in which lasing occurs. It consists of atoms or molecules in either gaseous (for example in a Helium-Neon laser), liquid (for example in a dye laser), or solid form (for example in a ruby or Nd:YAG laser). The medium of a laser diode is a doped semiconductor. You'll learn more about laser diodes later on, but for now you don't need to worry about these details. It is sufficient to think of the medium as a bunch of atoms that are typically in their ground state but can be excited somehow, and once they are excited, they can emit photons in a stimulated way, leading to the emission of laser light.

The job of the **pump** is to excite the atoms. Depending on the laser, the pump can be a flash lamp as in a Nd:YAG laser, or a potential difference as in laser diodes and Argon-Ion lasers. It can even be another laser. Really, pretty much anything that can excite atoms can be used as a pump. After a certain - typically very short - amount of time on the order of 10^{-8} to 10^{-4} s, the atoms that have been excited by the pump drop back into their electronic ground state. Remember that this emission can be either spontaneous or stimulated. In the box below, draw a simple two-level atom (that's an atom that has only one ground and one excited state) and sketch the three possible atom-light interactions listed in the box.

* Absorption: Spontaneous Emission: Stimulated Emission:

Spontaneously emitted photons from two different atoms have no fixed phase relation to one another; they are typically not emitted in the same direction, don't have the same wavelength, and are generally completely independent. In an analogy, you can think of them as kids on the play ground. They'll run all over the place! Photons emitted through *stimulated* emission on the other hand have a fixed phase relation and are emitted in the same direction as the incident photon that stimulated the emission, leading to *coherent light*. They are more like a marching band, completely in sync. This is what we need for a laser.

The last component of a laser, the **cavity**, increases the number of photons that were emitted through stimulated emission and keeps them in a well-defined direction. In the simplest case you can think of it as two plane mirrors that face one another. If the emitted photons have just the right wavelength (which also means just the right energy since the energy $E = hf$ and the wavelength λ are related through the speed of light $c = \lambda f = \lambda E/h$, they form standing waves between the two mirrors. This situation is very similar to standing waves on a string.

*** To refresh your memory, sketch the first three standing waves on a string of length L.**

Let's assume that we place the medium inside the cavity, and that the pump has excited a good fraction of the atoms in the medium. Let's further assume that one of those atoms releases a photon spontaneously and drops back into the ground state. Let's follow the path of this spontaneously emitted photon! As it moves through the medium it can stimulate emission in other atoms. For the sake of the argument, let's say that it "succeeds" once while traveling through the medium, so that there are now two photons exiting the medium. If the initial photon was moving perpendicular to the axis of the cavity, i.e. parallel to the mirrors, it (and its follower) will just continue on its path and leave the medium. If, however, the photon happens to travel parallel to the axis of the cavity, i.e. perpendicular to the mirrors, it (and its follower) now hits the mirror under a right angle, is reflected, and travels back in the exact opposite direction from where it came. As the two photons now travel through the medium in the opposite direction, they again have a chance to stimulate emission of more photons. If each of these photons again succeeds once, there will now be four photons exiting the medium on the opposite side, where they once again bounce off a mirror, and the process repeats over and over again.

*** Estimate how many times the initial photon and its followers need to bounce off the mirrors to generate a million photons inside the cavity.**

Calcuate the energy of a million photons assuming that they have a wavelength of 633 nm and compare it to the energy that a 5 mW laser pointer of the same wavelength emits every second.

The cavity thus creates an *avalanche* of coherent photons all traveling in the same direction, perpendicular to the mirrors. These photons bounce back and forth repeatedly and form a standing wave between the mirrors if they have the right wavelength. Thus, more and more energy builds up inside the laser cavity. We can extract part of that energy by making one of the cavity mirrors a little bit leaky, so that during each round trip a few photons leave the cavity and become the laser beam that is used in the experiment. Typically, about 3-5% of the light that hits this leaky *output coupler* is transmitted. So only a tiny fraction of the actual laser power circulating within the resonant cavity is used!

Here's a two-part trick question for you:

*** Assume that you have a very powerful laser that *outputs* 100 W. Would you place your hand in the laser beam that is coming out of the laser? If yes, why? If no, why not?**

*** Based on what you just learned about the output coupler and the cavity, would you place your hand *in between* the cavity mirrors, i.e. *inside* the laser cavity? If yes, why? If no, why not? Think carefully about this; this is the trick part!**

As often, life is a bit more complicated than this simple 'two-level atom in a cavity consisting of two plane mirrors' descriptions makes you believe, and the complication has to do with the so-called *population inversion*. Without any pump present, most of the atoms are in the ground state because it takes energy to get to the excited state. A few atoms can be in the excited state if there is enough thermal energy and if the two atomic levels are close together. However, for most atomic systems, this thermal excited state population is negligible. We call this situation in which more atoms are in the ground than in the excited state a *regular population*. As we increase the pump power, the fraction of atoms in the excited state increases because more atoms absorb energy from the pump and thus push their electrons to the excited electronic state. But there is a cap, and it has to do with the *rates* at which absorption and emission happen. Quantum Mechanics teaches us that the rates for (stimulated) absorption and stimulated emission are exactly the same! Both depend on the strength of the pump (which makes sense since both require a photon to work), but the way they depend on it is exactly the same for both processes. In simpler terms: Imagine what happens when a photon of correct energy "meets" two atoms, one in the ground and one in the excited state. The probability that it interacts with atom one and gets absorbed is exactly the same as the probability that it interacts with atom two, which would lead to stimulated emission. So let's think about what happens when we shine an intense pump beam into an ensemble of two-level atoms that are all initially in their ground state. For a short time after we turn on the pump beam, absorption dominates since all atoms start out in their ground state. But as more and more atoms enter the exited state, it is not clear whether the absorption or the stimulated emission "wins". In fact, the best we can achieve is a stalemate! That's what we mean when we say the rates for absorption and stimulated emission are exactly the same.

*** As an example, draw this for an ensemble of 10 atoms. Show the distribution (i) before the pump is turned on, (ii) immediately after it has been turned on, and (iii) after the pump has been been on for a while (and assuming that it is quite intense).**

It thus makes sense, and one can easily show this mathematically, that therefore lasing is not possible in this simple two-level system, because the best we can achieve is a stalemate: The amount of energy that is absorbed is equal to the amount of energy that is emitted through stimulated emission, so there is no net gain in the number of photons, and no laser beam that can be used in an experiment is emitted. If you are not fully convinced have a look at this PHET simulation: `phet.colorado.edu/en/simulation/lasers`. The 'lamp control' button allows you to set the intensity of the lamp that excites the atoms. You can also set the wavelength to be in resonance with the transition. Go to the 'multiple atoms' tab and set the intensity to very high. Make sure you display a two-level atom, and that your lamp is in resonance with the transition. Then observe the populations in the ground and excited states. The following three screenshots were taken just a few seconds apart. While at any given moment there may be more photons in the excited than in the ground state (see figure on the left), there may also be times when there are more atoms in the ground state (middle figure). *On average*, the two populations will be very similar (figure on the left).

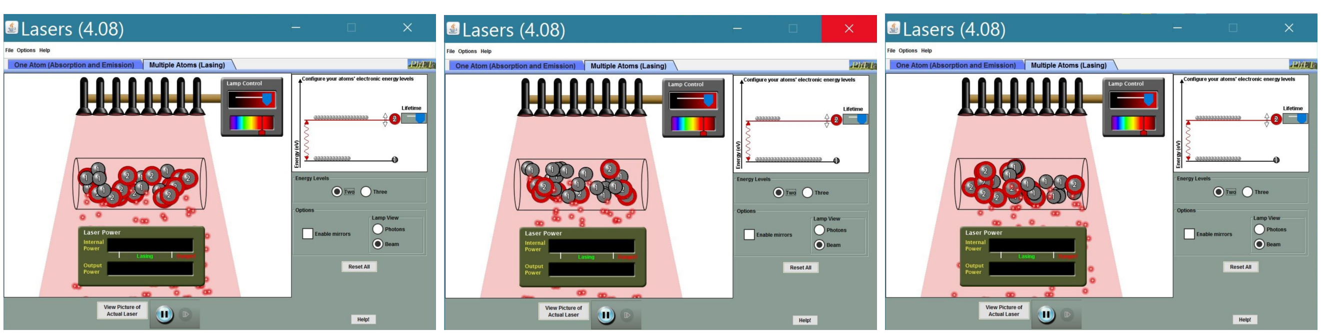

In order to achieve a *net gain*, which means that more photons are produced through stimulated emission than are absorbed, we need a *population inversion*, which means that *on average* more atoms are in the excited than in the ground state. As we just reasoned, this is not possible in a two level system, and that's why lasers involve at least three, many even four atomic energy states. You can adjust the settings in the PHET simulation to turn your two-level atom into a three-level atom. When you pick the wavelength and intensity of your pump lamp just right, you can get a working laser!

*** Take a screenshot of your working laser and attach it to this manual. Briefly describe how you had to adjust the settings to make it work.**

The figures below, taken from the Wikipedia page on population inversion, show you generic examples of three- and four-level lasers, and you can learn more about them in our upper-level Optics course or in any optics or laser textbook. If you look for instance at the three-level system (the left figure), you can see that the pump connects the ground state (Level 1) with an excited state (Level 3) that decays rapidly (on the order of 10^{-8} s) into a lower-lying excited state (Level 2). Level 2 has a very long lifetime, on the order of 10^{-6} s or longer, which means that the electrons stay in this excited state much longer than they stayed in Level 3. So while the pump can only achieve a 50/50 splitting of the population between Levels 1 and 3, Level 2 accumulates more and more electrons, and eventually ends up having more electrons than the ground state at any given time.

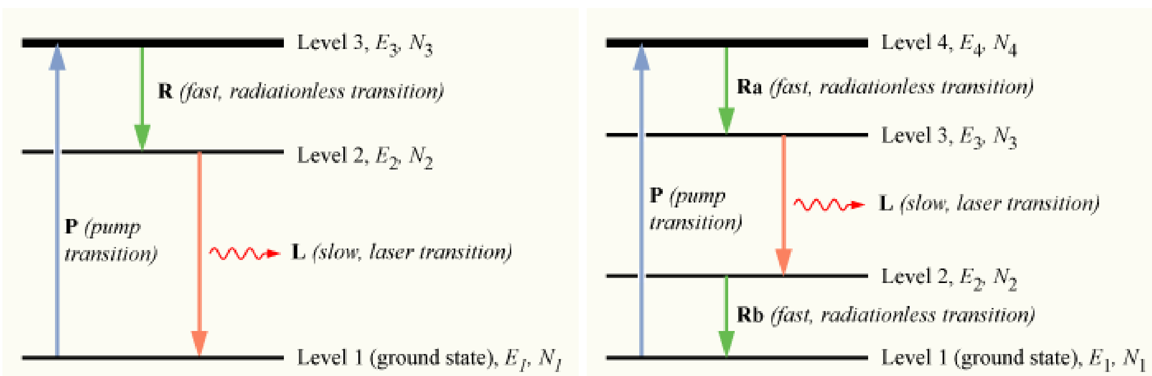

But more electrons in Level 2 mean that there will be more stimulated decays from Level 2 into Level 1 than absorption from Level 1 into Level 2, and that means we can extract energy in form of laser light.

*** The four-level system works in a very similar way. Describe it in the box below.**

6 Introduction to laser diodes:

In this module you will work with a special type of laser, a *laser diode*. The figure below, taken from Thorlabs' website, shows you what laser diodes look like. They often come in two sizes (i.e. two diameters): 5.6 mm and 9 mm (although there are also other ways to build and package a laser diode).

Laser diodes are semiconductor devices that are often fantastic choices because of their relatively low cost (at least at the wavelength you will be working with), their availability over a wide wavelength and power range, their ruggedness and reliability, and their ease of use. If you want to learn in more detail how they work, search the internet or pretty much any optics or laser textbook. For example, "Laser Physics" by Milonni and Eberly is a good starting point.

Very briefly, a semiconductor material is, as its name suggests, a "half-conducting" material, halfway between an insulator and a conductor. Unlike individual atoms that have discrete energy levels in which their electrons are located, bulk materials like metals, semiconductors, or insulators, have *bands* in which their electrons reside. These bands happen naturally when you bring lots of atoms close together, because the atoms "feel" the other atoms in their vicinity. Quantum mechanically, that leads to an interaction term in the Hamiltonian of the system, which in turn leads to a shift in the allowed energies. Since the environment is a little bit different for each of the individual atoms, their energy levels are all shifted by a slightly different amount. And because there are so many atoms in a bulk material, this leads to a broad band of allowed energy levels over which the electrons are distributed. Schematically, this effect is shown in the figure below which was taken from `www.ibiblio.org/kuphaldt/electricCircuits/Semi/SEMI_2.html`:

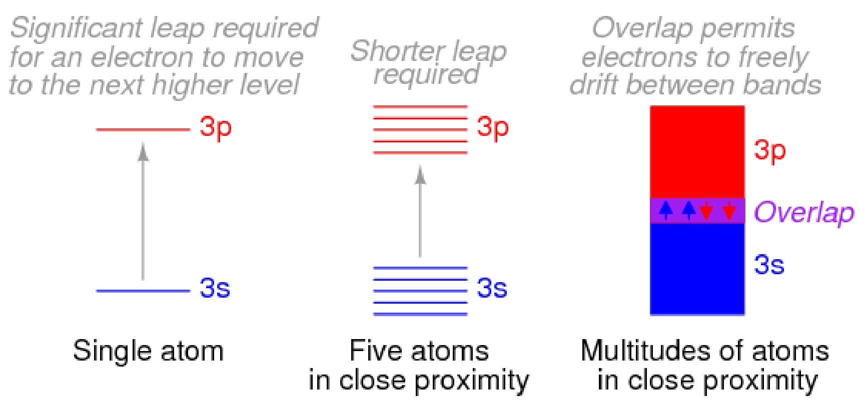

The highest-energy band that is completely filled is called the *valence band*, and the band above it is called the *conduction band*. In insulators, the valence band is completely filled and no electrons are in the conduction band. This explains why insulators can't conduct a current: The electrons in the valence band are so tightly packed that they can't really move freely - think of a very crowded room, you wouldn't be able to move through it easily either -, and there are no electrons in the conduction band. And no moving electrons means no electrical current. In metals on the other hand, there are a few electrons in the conduction band. While their fellow electrons in the valence band still can't move, the conduction band electrons see all this free space and can move around very easily. That's why currents can flow easily in metals. Semiconductors are somewhere in between: While they have no electrons in the conduction band, the *band gap* (i.e. the energy difference between the valence and conduction band) is very small, so that only a small amount of energy is necessary to move a valence band electron into the conduction band. Thus, semiconductors become more and more conductive as you increase the temperature of the semiconductor (remember that increasing the temperature means you increase the thermal energy of the electrons, and that increase in energy is enough to lift them from the valence into the conduction band).

* **Draw a sketch of the valence and conduction bands for insulators, metals, and semiconductors in the box below. Indicate where the electrons are located. To keep things simple, assume that each band can hold up to six electrons.**

Doped semiconductors are even more useful. Doping a semiconductor means that you add a small percentage of atoms from a different column of the periodic table. For example, a typical semiconductor is silicon, an element in Group 14, the carbon group. If a few atoms of Group 13 are added to the silicon, it becomes a *p-doped* semiconductor, while it is called *n-doped* if electrons from Group 15 are added. Group 15 elements have one more electron than silicon, which means they donate an extra electron to the bulk material. Since the valence band is completely filled, these extra electrons move close to the conduction band (they don't have the exact same energy as the silicon conduction band, so they end up just below the lower edge of the conduction band), which means that these electrons need a lot *less* energy to hop into the conduction band where they can conduct electrical current. Similarly, an element from Group 13 has one less electron than silicon and thus provides an extra free space, or *hole*, in the valence band. With these extra empty seats, electrons can now move around in the valence band and thus also contribute to a current. The figures below, taken from `hyperphysics.phy-astr.gsu.edu/hbase/solids/dsem.html#c3`, show you the effect of doping.

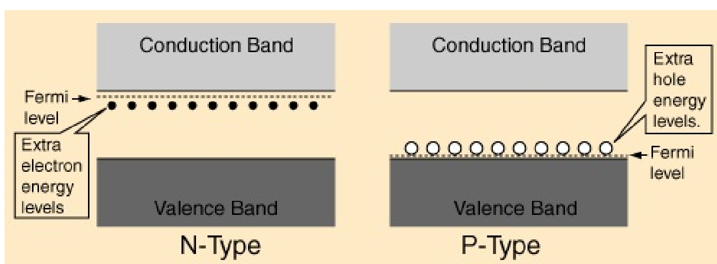

By connecting a p- and an n-doped semiconductor and applying a potential difference across the p-n junction, the valence and conduction bands are shifted and distorted such that, at some point within this p-n junction, more electrons are in the conduction band than in the valence band as shown in the figure below, taken from `www.ni.com/white-paper/14878/en/`. This again is population inversion, leading to more stimulated emission than absorption, and thus to the emission of a laser beam, as we discussed above.

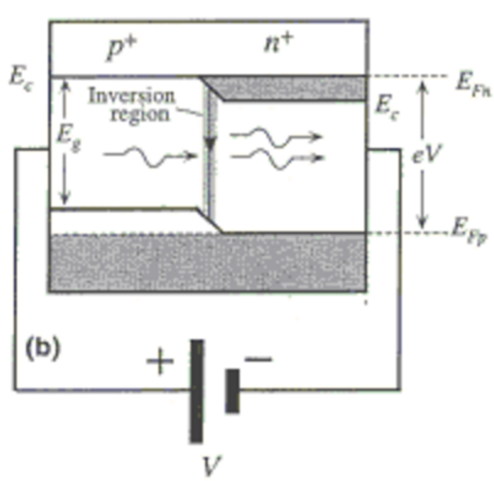

7 Turning laser diodes into diode lasers:

Now that you know what a laser is and how laser diodes work, let's use them! The blue container box labeled 'Module 3' contains a mount for a laser diode (Thorlabs TCLDM9 or LDM21), a temperature controller (Thorlabs TED200C), and a current driver (Thorlabs LDC205C). Take them out of the box and place them on the optics table. You should also find the manuals for all three pieces of equipment. Take a moment to familiarize yourself with the equipment. In particular, focus on what each of the buttons does, and determine which ones you will need in order to operate the equipment correctly. **Note that the following instructions in this manual are meant to *supplement* the equipment manuals; they are not meant to replace the manuals completely and are not detailed enough to provide you with enough information to run the equipment correctly.** Please also watch the `VIDEO` called `HOW TO SET UP THE DIODE LASER SYSTEM` on our website.

Record your notes here.

You should also find a few laser diodes (L780P010) in their protective bags in the box labeled 'Module 3'. The spec sheet for these laser diodes is copied below (feel free to go to Thorlabs' website for a larger version).

Laser Diode

ATTENTION
OBSERVE PRECAUTIONS FOR HANDLING ELECTROSTATIC SENSITIVE DEVICES

L780P010

Description

Thorlabs Ø5.6 mm, TO-18 can package discrete laser diode is a compact light source suited to many application. Our lasers are fully compatible with our entire line of Laser Diode and TEC Controllers as well as our selection of Laser Diode Mounts and Collimation Solutions.

Specifications

Specification	Symbol	Max
Optical Output Power, mW	P_o	12
LD Reverse Voltage, V	$V_{R(LD)}$	2
PD Reverse Voltage, V	$V_{R(PD)}$	30
Operation Case Temperature, ºC	T_{op}	-10 to 60
Storage Temperature, ºC	T_{STG}	-40 to 85

(P_o = 10 mW, T_c = 25 °C)

Specification	Symbol	Min	Typ	Max
Threshold Current, mA	I_{th}		14	20
Operation Current, mA	I_{op}		24	40
Operating Voltage, V	V_{op}		1.8	2.5
Lasing Wavelength, nm	λ_p	770	780	790
Beam Divergence	$\theta_{//}$	6°	8°	12°
Beam Divergence	$\theta_{\perp}$	25°	30°	32°
Slope Efficiency, mW/mA	η	0.5	0.65	0.9
Monitor Current, mA	I_m	0.3	0.7	1.5
Astigmatism, µm	A_s		10	15

Drawings

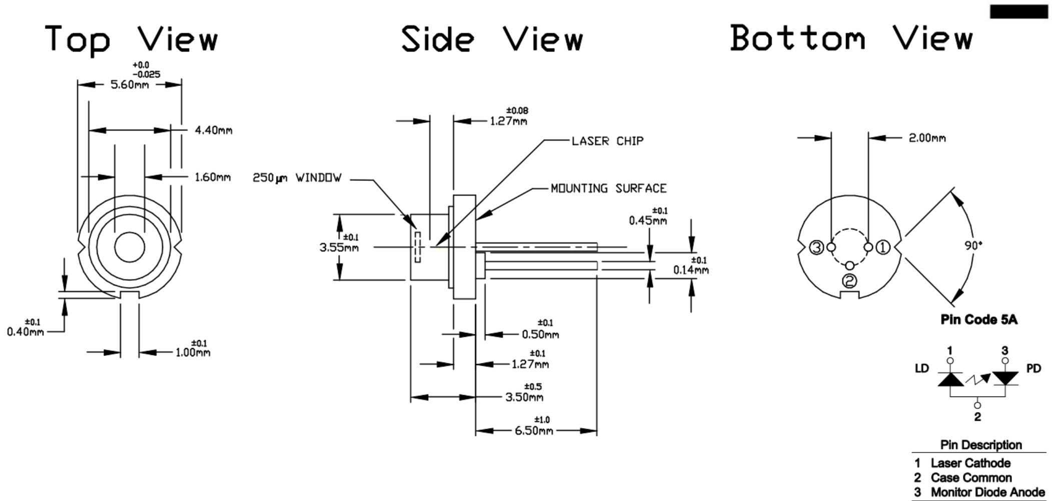

Take one of these laser diodes out of the box (**but leave it in its protective bag!**) and place it on the table. Confirm that the table is properly grounded (that means that there is a cable running from (and safely connected to) the table top to the ground part of an outlet). If it isn't, ground it now or ask your instructor for help. Laser diodes are very sensitive to static charge. This means you should never handle them without properly grounding all the equipment they will come in contact with. That includes yourself! The plastic bag in which the laser diode arrives is designed to protect the laser diode from static charge, but as soon as you take the laser diode out of the bag you must be properly grounded. To do this, take the light blue anti-static wrist strap shown in the figure below, attach it tightly to your wrist, and connect the other end firmly to the table. Note that only students who are properly grounded may handle the laser diode until it is safely mounted in the driver!

Your first task is to insert the laser diode correctly into the mount and turn it on. The final setup is shown in the figure below. Here are the step-by-step instructions on how you do it.

The black mount has a four-pin plug in the center to hold laser diodes with 5.6 mm (like the one you are using) or 9 mm housings. Flip the mount over so that it rests on its back with the part that will hold the laser diode facing up. That'll make inserting the diode a lot easier. The diagram on the front plate of the mount shows you that the top and the bottom pins are connected to ground while the left one (when looking directly at the mount) is for the photo diode (PD) and the right one (when looking at the mount) is for the laser diode (LD). You want to make sure that you insert the diode correctly, i.e. you must find out which of the three pins of the diode is ground, which is the LD, and which the PD. The ground is easy: It's always the middle pin, the one that is visibly connected to the rest of the laser diode housing ("can"). The can also typically has a rectangular notch on the side where the ground pin is located. Look at the laser diode's spec sheet and find the diagram that shows you the bottom view with the three pins labeled.

Copy it here.

You want to line up the three pins such that they match the diagram on the mount. Discuss with your group members how you need to insert the diode (remember that the spec sheet shows you the *bottom* view of the diode!).

Describe where the notch is located (up or down) when the diode is inserted correctly and you are looking at it.

Now carefully open the protective bag and let the diode gently glide onto the grounded table (be careful, don't let it fall into one of the holes!). You want to hold it gently from the side, touching the metal can, and without touching the glass surface (wearing gloves is a *great* idea!). Confirm that the three pins are sticking out straight, maybe even outward just a tad bit; that will make inserting it into the mount much easier. Then carefully insert it in the orientation you just determined and press down on the metal ring of the can until the diode is flush with the mount. You can also use the foreceps to push it down (just be careful that you don't scratch the glass surface!). If you can't fit the pins into the mount, bend them outward just a little bit and try again. When the diode is pushed all the way against the mount, insert the black retainer oval (shown below mounted inside the laser diode mount) and tighten the two plastic screws to safely hold the diode in place (careful not to over-tighten them, they might break).

Connect the two drivers to the mount using the two 9-pin cables and stack them on top of each other on the table next to the mount, just as shown in the figure on page 11. Flip the mount back up and secure it to the table using a post and post holder with base. Before you can turn it on you have to check that the mount is correctly calibrated for this particular diode. Depending on the brand of laser diode you are using, the laser diode *anode* or the laser diode *cathode* may be connected to ground. The Thorlabs mount has two sliding switches located underneath the black square front cover that allow you to run laser diodes with their cathodes on ground (CG) or their anodes on ground (AG), and similarly for the two PD combinations.

Go back to the spec sheet of the diode and copy the pin code diagram into this box. It shows you which part of the laser and photo diode is connected to ground, the cathode or the anode.

If, for example, the laser cathode is connected to ground, you want to toggle the corresponding switch on the mount to CG (for 'cathode ground'). If the laser anode is connected to ground, you want to switch it to AG (for 'anode ground').

Record the settings of the two toggle switches here.

If you accidentally set the switches incorrectly, the current driver will beep, a red LED will turn on, and no current will flow. This is a smart driver that protects your laser diode. Note, however, that if you use a home-built or a cheaper driver it may not warn you and instead destroy the laser diode if you accidentally connected it incorrectly and enable the current through it.

Now we are ready to test the diode. First, **turn on the temperature controller**. The purpose of this controller is to keep the laser diode at a certain pre-set temperature. It does that using a so-called PID controller (more in a little bit). You can change the set temperature by pushing the buttons near the middle of the controller until the green LED next to TSET lights up. Then rotate the knob to change the set temperature. Careful, though: The temperature is shown as the *resistance* of a thermistor located inside the mount. You can find the resistance-temperature conversion chart in the mount's manual. You want to keep the temperature between about 15°C and 25°C. That avoids accidental condensation of water on the laser diode, which could damage it, or excessive heating, which could shorten its lifetime.

What resistance range does that temperature range correspond to? Does a higher resistance correspond to a higher or a lower temperature?

Set the temperature to about 20°C. Push the button in the middle of the controller once more to switch to TACT, the actual temperature of the diode. Since the controller is not yet enabled, TACT will be different from the set temperature. Convert the reading of the actual resistance to a temperature using the chart in the diode mount's manual.

What does this temperature correspond to, and does it surprise you to get this reading?

Now press the 'ON' button and the actual temperature should change and approach the set temperature. Note that it is normal for the actual temperature to oscillate around the set temperature a few times before settling down. **So wait a few minutes until the temperature has stabilized.** If it doesn't stabilize but continues oscillating, the *PID* values of the temperature controller are probably set incorrectly. In that case, you can adjust them by turning the small potentiometers labeled *P*, *I*, and D with a screwdriver (or, in a pinch, your fingernail). To start, you want to turn them all the way counterclockwise; that sets them to their lowest possible values. Wait until the temperature has stabilized a bit, then increase P just a little bit by turning the potentiometer clockwise. P stands for proportional, and the value of the P-potentiometer adjusts how quickly the controller reacts to a change in temperature. Basically, when the temperature deviates from the set-point (let's say it drifts up a little), then the P-part of the controller sends out a voltage signal that compensates for this drift (in this example it would be a small negative voltage). Since this is an instantaneous effect, the resulting changes can be quite abrupt (WHAM! WHAM! WHAM!). That's where the I-part comes in (I stands for integral). The I-part allows the controller to wait a little by integrating the drift for a certain amount of time, finding the average fluctuation of the temperature from the set-point over that time, and then sending out a signal that corrects for this *average* fluctuation, leading to a much gentler change (nudge, nudge, nudge). In a nutshell, P makes the controller fast, I smooths out sharp fluctuation spikes. Both together, if set correctly, allow for a very stable control of the temperature. Note that for diode lasers you typically don't need to use the D-part (D stands for differential), so leave the potentiometer rotated all the way counter-clockwise. Also note that these diodes typically also need only small values for P and I. The actual temperature reading will be a little bit different from the set temperature. That's normal, and as long as the actual temperature is stable we don't care too much about this discrepancy. In the following, just make sure to record the *actual* temperature and not the set point.

As soon as the actual temperature has stabilized you are ready to **turn on the current controller**. Make sure that the current knob is rotated all the way counterclockwise, i.e. no current will flow when you enable the controller. Then turn on the controller by pushing the power button and enable the current output by pressing the top right button. A short beep followed by a *green* LED next to the 'LASER ON' button should indicate that the diode and mount are set up correctly. If you hear a beep and a *red* LED turns on, chances are you made a mistake in the position of the two switches on the mount or in the orientation of the diode in the mount. Check them again. If the problem persists, check in with your instructor.

Starting at zero current, increase the current slowly (5 mA per second). With just a few mA of current through the laser diode, you should be able to see a faint glow on a detector card that is held no more than an inch away from the laser. This glow should increase in brightness drastically once the current through the laser diode is above about 20 mA. If you see this, everything is working as intended! Set the current to about 30 mA. Note that from now on, you may enable and disable the laser diode simply by pressing the LASER ON button, i.e. you don't have to turn down the current with the knob first.

Note that the small potentiometer located next to the ILIM LED can be used to set the limiting current, i.e. the maximal current the current driver will provide independent of the position of the knob. To get there, press the middle button until the LED next to ILIM lights up. Then use a very small screw driver to gently rotate the potentiometer. Check what the maximum allowed current for this diode is (it's in the spec sheet) and set ILIM to slightly below that current. This will protect the diode and prevent you from accidentally setting the current too high.

Below is a brief summary of everything you have done so far:

How to correctly mount a laser diode:

- Properly ground yourself and all equipment before touching the laser diode.
- Never touch the front facet (the glass surface) of a laser diode with your fingers. Wear gloves!
- Identify which laser diode pins correspond to the LD, the PD, and the ground. Then insert the diode with the correct orientation, making sure that it sits snugly against the back of the mount.
- Check the diode spec sheet to see whether the anode or cathode is connected to ground and adjust the toggle switches on the mount accordingly (AG or CG for both the LD and the PD).
- Connect the temperature and current driver.
- Turn on the temperature controller first. It should turn on with just a single beep and you should see a green LED. Set the temperature to somewhere between 12 kΩ and 14 kΩ. Press the ON button and make sure that the actual temperature reaches a stable value. If not, adjust the PI values.
- Turn on the current driver. Rotate the current knob all the way counterclockwise. Enable the current output. It should turn on with a single beep and a green LED should light up. Slowly increase the current while monitoring the output of the laser with a detector card close to the diode (divergence!). You should see a faint glow above a few mA, and the laser should be clearly visible above about 20 mA. After you have confirmed that the laser diode is mounted correctly following these initial steps you can turn the current on and off simply by pressing the 'enable' button without rotating the knob down to zero first. However, never press the power on/off button without first pressing the enable button; that could permanently damage the laser diode!
- Make sure to set ILIM correctly using the potentiometer and a small screw driver.

You should now have a laser diode that is safely mounted inside a Thorlabs mount and connected correctly to a current and temperature controller. Its temperature should be very stable, and the current through it should be around 30 mA.

However, you probably noticed that the beam emerging from the laser diode diverges quite rapidly. That's because the region in which the laser emission is created is so small (the size of the semiconductor material is on the order of just a few micrometers!). You've already seen this phenomenon before when you studied the single slit experiment: The smaller the slit width, the broader the interference pattern on the screen, i.e. the more divergent the beam after it passed through the slit. In order to turn the laser into a useful tool that works over several meters you must *collimate* the laser beam. Note that students often call this "columnating" (which is not even a word!) or "focusing" (which is incorrect - a focused laser beam is one that comes to a focus (gets very small) and then diverges from that focus. Exactly what we don't want!). As you remember from MODULE 1 - INTERMEDIATE, a single converging lens is all you need to collimate a diverging beam.

Thorlabs' mounts are well-designed because they allow you to easily attach a black disk that screws into the front plate of the mount and holds a small lens in its center. Even when rotating that disk, the lens stays, at least roughly, centered. Remember from MODULE 1 - INTERMEDIATE that this is important to avoid astigmatism. Both the disk and the lens are in the box labeled 'Module 3'. Find them. They should look like this:

In the tool drawer you should find another spanner wrench similar to the one that you used in Module 1, but this one has two metal prongs sticking out that fit nicely into the two holes of the black disk. This tool allows you to gently rotate the disk with the lens while at the same time monitoring the beam. A third, much smaller (and silver) spanner wrench can be used to mount the lens inside the black disk.

Before you go ahead and mount the lens, let's think about *how exactly* we want to **mount the lens in the black disk** since there are four different ways how you can orient the lens relative to the black disk. Have a look at the box in which the lens came. It should tell you that the lens' focal length is pretty short, on the order of 5 mm. That means that we need to position it almost directly behind the laser diode.

Discuss with your group members how you want to mount the lens in the disk and why. A sketch might also be helpful. Then mount it safely in the disk using the small silver spanner wrench.

Turn the laser off by pressing the ON button on the current driver. Screw the black disk (with the lens safely held in place) into the mount in the orientation you and your group members agreed on. Keep in mind that the focal length of the lens is pretty short, so you want to screw the black disk in quite a bit. Place a detector card a few inches away from the laser and enable the laser current (make sure it is set to about 30 mA). Rotate the disk with the lens using the spanner wrench and observe the beam profile on the detector card.

Describe what you see. In particular, focus on the position of the laser beam on the card. Does it change? If so, how? And if so, why? Also focus on the "roundness" of the beam and compare it to the beam profile of the HeNe laser you worked with in `MODULE 1 - INTERMEDIATE`.

To collimate the laser diode properly, rotate the disk until you minimize the spot on the card (i.e. until the focus of the laser beam is on the card). Then move the card farther back and rotate the lens again until you get a focus on the card. Repeat this several times. This pushes the focus farther and farther away from the laser diode, until it is essentially at infinity (or as close as we can get to infinity in the limited space we have in the lab!). Remember that a focus at infinity corresponds to a collimated laser beam. You thus want to make sure that the distance over which you check the beam's collimation is larger than the distance the beam will travel over the optical table. Note that you probably have to move the card away slowly while adjusting the position of the lens, instead of hoping to find the beam somewhere on the wall. As always: It is easy to optimize a signal you can see; it is very difficult to optimize a signal you have lost.

Comment on the shape of the beam *in the focus*, and whether the positions of the horizontal vertical focus are the same. As mentioned in `MODULE 1 - INTERMEDIATE`, laser diodes are typically highly astigmatic. You should be able to see that by moving the detector card through the focus of the laser beam.

Once the laser beam is reasonably well collimated, observe the beam profile on the detector card as you slowly move it from the laser to the far spot near the edge of the room.

Describe it in the box below and compare it to the HeNe beam profile.

Here's the summary of what you've just done. You may also watch the `VIDEO` called `HOW TO COLLIMATE THE DIODE LASER` as a refresher.

How to collimate a laser diode:

- Note: Collimated is not the same as columnated or focused!
- Mount the lens in the black disk and screw it into the mount. Note: The lens has a very short focal length.
- Monitor the beam with a detector card at a certain distance away from the laser. Rotate the lens until the spot on the card is minimized.
- Move the card farther away from the laser and again minimize the spot on the card by rotating the lens.
- Repeat until the card is as far away as possible (typically on the other end of the room). By minimizing the spot on the card at that distance you are essentially "focusing the beam at infinity", which is the same as saying you are collimating it.
- You can use an additional cylindrical lens if the beam is strongly astigmatic or a prism pair if the beam is highly elongated.

You should now have a nicely collimated laser beam that diverges only minimally across the full length of the optical table. So let's characterize it! For the following exercises, create a new Jupyter notebook that you can share with your group members and your instructor. If you've never worked with Jupyter before let your instructor know and they can help you get started.

Your first goal is to **measure the output power of your laser as a function of current** through the laser diode. Discuss with your group members what the experimental setup should look like and which experimental parameters you want to vary, in which order, and by how much.

Record the result of your discussion below. Feel free to check with your instructor to see if you forgot anything. Record your data in a properly labeled table.

You just discovered the so-called **threshold current**, the current at which the output of the laser diode suddenly increases significantly. Below this threshold current the laser diode is only emitting random photons (this is called **amplified spontaneous emission or ASE**). Above the threshold, the laser diode is actually lasing.

Sketch your data in the space below (you will generate a nice graph for your report, but you don't need it at this point) and indicate the threshold current in your sketch. You can find it by extending the linear increase at larger currents back and finding the crossing with the x-axis.

Observe the laser beam on a detector card as you slowly adjust the current from below to above the threshold current. It should be very clear when you hit the threshold current.

Record the value of the threshold current here and compare it to your estimate from the graph. Hopefully they agree! We will revisit this threshold current in a little bit.

Your next goal is to **measure the wavelength of the laser diode as a function of temperature and current** using the Mightex spectrometer. If you need a refresher, watch the `VIDEO` called `HOW TO USE THE MIGHTEX SPECTROMETER` that you can find on our website. Remember that temperature changes always take a long time to take effect. Discuss with your group members what the experimental setup should look like and which experimental parameters you want to vary, in which order, and by how much (note: To get good data, please vary the current in increments of no more than 2 mA!).

Record the result of your discussion below. Feel free to check with your instructor to see if you forgot anything. Record your data in a properly labeled table and feel free to add additional pages if you need more space. Please record the wavelengths of *all* peaks that you observe for a given current!

Let's analyze your data. Generate a nice plot of the wavelength as a function of current *for one temperature* in your Jupyter notebook. In a second plot, pick a *fixed current* and show how the wavelength changes with temperature. Add a print out to this manual.

Try to come up with some general statements along the lines of "as the temperature increases, the wavelength..." and "as the current increases, the wavelength ...".

We'll now try to understand where these dependencies come from. Let's first look at the **temperature dependence:** Remember that in diode lasers the lasing happens between the bent conduction and valence bands. As you increase the temperature of the semiconductor, the material expands a little. That means its *spatial* dimensions increase. How does that translate into a wavelength? In solid state physics you learn that the spatial and the frequency domain are inversely related. But you know this already from waves: The argument of a sinusoidal wave has the term $kx = 2\pi\frac{x}{\lambda}$. So if the spatial dimension (x) gets larger, then λ must also get larger.

Thus, as the temperature increases, the wavelength

From your data, estimate the rate of change (in nm per Kelvin) *for a fixed current.* Note that you are looking at the short and slowly increasing parts of your data set, not at the abrupt changes. Those are *mode hops* and we will analyze them later!

Explaining the change with **current** is a bit more tricky since there are two effects at work: Increasing the current again increases the temperature (slightly), so we would expect the same dependence as above. In addition to that effect, however, the density of electrons in the junction between the p- and n-doped semiconductors also changes, which affects the index of refraction. As the current increases, the density of electrons increases, which also increases the index of refraction. So, in order to maintain a standing wave inside the cavity, the wavelength has to increase to compensate for the increasing refractive index, since $\lambda_n = \lambda_0/n$.

Thus, as the current increases, the wavelength

From your data, estimate the rate of change (in nm per mA) *for a fixed temperature.* Again, make sure to analyze the short and slowly increasing parts of your data set, not the abrupt changes.

You should have noticed two things that are worth discussing in more detail: (1) You probably saw that for certain currents you get a lot of peaks, while for others you get just a few. Lasing at multiple wavelengths at the same time is called *multimode behavior*, and it is very common in free-running laser diodes. (2) If you look carefully at your data, you should see an overall linear trend as a function of current, but there should also be sharp kinks between linear sections. These kinks or jumps are called *mode hops* and they, too, are very common in diode lasers. Let's first look at mode hops and then analyze the multimode behavior some more.

Mode hops arise because laser diodes are essentially two resonance cavities stacked inside one another. To understand this, let's draw a very simplified version of a laser diode. The semiconductor material that is responsible for the lasing is a highly polished rectangle. One of the polished sides is lined up with the back of the laser diode can, while the other one is inside it, parallel to the front window. The two polished sides act as the laser cavity, because as you learned in Module 1, glass reflects about 5% or so of the incoming light. So, the lasing wavelength has to form a standing wave within this cavity. But the laser diode can also has a glass window, which acts as *another* cavity mirror because it, too, reflects some of the incident light. We thus have a small cavity (the semiconductor) inside a larger cavity (formed by the glass and the back of the semiconductor).

Draw this here.

Now, remember that in order to get stable amplified stimulated emission and thus lasing, you need to form standing waves inside the cavity. But because we have *two* cavities here, the situation is a bit more complicated: Only waves that form standing waves in *both* cavities simultaneously can be amplified. **Sketch such a special standing wave into your drawing in the previous box.**

Then use this information to explain mode hops (and in particular the fact that there are certain wavelengths that you cannot reach as you increase the current for a fixed temperature (or the temperature for a fixed current).

Lastly, let's analyze the **multimode behavior.**

Count the number of modes for each of your current values and measure the distance between two consecutive modes. Record that data in a properly labeled table in the box below. Then convert this distance into a frequency distance (careful here: $\Delta f = \frac{c}{\lambda^2}\Delta\lambda$! Show the derivation of this expression as well!). What can you say about the spacing of the modes?

8 Extended cavity diode laser

While changing the current through or the temperature of a laser diode to change its wavelength is a perfectly viable method, it can be a little bit too blunt of a tool for certain applications. The spectral range, i.e. the width of the emitted light, of a free running laser diode covers several nanometers as you have just seen, and mode hops and multimode lasing is very common. All of these makes a free-running laser diode too broad (in frequency) and too unreliable to precisely "talk" to the very narrow atomic transitions. For example, rubidium has four different electronic energy states in the first excited $P_{3/2}$ state that are only a few tens to hundreds of megahertz apart each, so if you want to precisely address just one of these states, then the laser linewidth must be smaller than a few ten or hundred megahertz. This is similar to painting with a very broad brush that can smear out various details versus painting with a sharpened pencil that can precisely draw a thin line from one point to another. You can learn more about these atomic energy states in the advanced version of this module.

One way to narrow the linewidth of a laser diode is to provide *optical feedback* by sending a small amount of light of a certain frequency back into the laser diode. The active medium of the laser diode thus "sees" more photons at that particular frequency, and - through stimulated emission - is more likely to emit light at that particular frequency, leading to an overall narrowing of the frequency spread of the emitted light. Let's look at this statement in a bit more detail.

Sending a small amount of light back into the laser diode can be achieve by using a diffraction grating (remember what you learned about them in `MODULE 1 - INTERMEDIATE`). The idea is this: A diffraction grating not only reflects light, but it also diffracts it into multiple orders, and within each order, it separates the light by wavelength, see the figure below as a reminder.

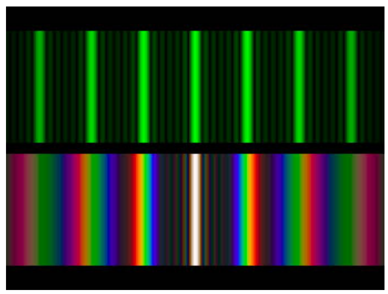

The top row shows you a green laser beam while the bottom row shows you white light after each has been diffracted by a grating. The zeroth order beam is in the center of the figure (this is the reflected but not diffracted beam), with the first three diffracted orders on either side of it. Again notice that the higher diffracted orders are more spread out in wavelength (the rainbow gets wider). The white light (bottom of the figure) is separated into its wavelength components, forming the characteristic rainbow spectrum. In the case of the green laser beam on the other hand, all orders appear green since the laser is highly monochromatic. If you looked very, very carefully, though, you would see that even the green light is separated by wavelength, but since the overall rage of wavelengths is so much smaller in the laser due to its monochromaticity, the effect of the splitting is much less noticeable; the full split spectrum appears green to our eyes.

So what happens if we adjust the angle of the grating such that one of these diffracted orders is sent back into the laser diode? Inside the laser, atoms are being pumped into an excited state; they are "ready to lase" as soon as they receive a photon that they can follow through stimulated emission. And sending some light from the grating back into the laser diode provides just that: A photon that can stimulate emission. By carefully adjusting the angle of the grating we can influence the wavelength of the photons that are sent back. And because there are now more photons at that particular wavelength, the stimulated emission predominantly happens at that wavelength. Thus, by changing the angle of the grating we can actually change the frequency of the laser in a very precise and controlled way. As an added bonus, this also narrows the profile of the laser diode to just a few megahertz; the broad laser emission "collapses" to emission at just the wavelength that is sent back from the grating. Additionally, multimode lasing is also strongly suppressed. There will still be mode hops, because as we discussed above, that is something that is inherent to laser diodes.

This type of laser consisting of a laser diode and a diffraction grating that provides optical feedback is called an *Extended Cavity Diode Laser (ECDL)*. There are two different ways to insert the diffraction grating: In *Littrow* or in *Littman-Metcalf* configuration.

*** Look up both of them and sketch them in the box below.**

*** Comment on similarities and difference between these two configurations. Think about ways when you would use one over the other.**

Your first task is to **assemble an ECDL in Littrow configuration**. You should find a mirror mount with a diffraction grating attached to it in or next to the box labeled 'Module 3'. Note that the grating is attached to a long green rectangular piece of ceramic; that's a *piezo-electric transducer* ('PZT'). We will use it in the advanced version of this module, so don't worry about it here.

To start, set the laser current to about 30 mA and turn off most of the lights in the room. Place the diffraction grating into the laser beam at an angle of about 45 degrees and about a cm or so away from the laser mount. Using a detector card, find the zeroth order diffracted beam (that's the beam that follows the regular Law of Reflection), as well as the first order diffracted beam (this one should be weaker than the zeroth order and leaves the grating at an "unexpected" angle). Rotate the mirror mount to roughly align this first order diffracted beam back into the laser. Make sure that this really is the first order diffracted beam and not the zeroth order beam as that could destroy the laser diode immediately! Move the detector card into the laser beam with the pink side facing toward the grating. This will block the laser beam, which means that the diffracted beam will disappear. The original laser beam will shine *through* the card and appear as a *blurry* spot. By moving the detector card in and out of the laser beam, you can "blink" between seeing the diffracted beam and seeing the original beam, which will help you find the angle that directs the first diffracted order directly back into the laser diode. Note that the brightness of the spots on the detector card may fade over time if the card is held in the same place for a while. This is due to *bleaching* of the pink material on the card; the intensity of the beam itself is not changing.

When the two beams appear fairly well overlapped, lock the diffraction grating down by using a table screw (or two dogs from opposite sides) and tightening the thumb screw that holds the post in the post holder thoroughly. Check that your alignment hasn't changed and that the diffracted beam is still fairly close to the laser beam. **Double check that all screws are thoroughly tightened; this is important** as bumping the grating mount even a little bit can knock off the alignment significantly. Check that the two screws on the mount that move the grating up/down and left/right are roughly in their middle position; we will use them in just a moment to fine-tune the alignment of the beam. It is pretty unlikely that your alignment of the diffracted beam back into the laser diode at this point is perfect, because it really needs to be aligned *very* precisely. Thankfully, we don't have to rely on this rough alignment method. Remember that the emitted wavelength is affected when the diffracted beam is directly sent back into the laser diode, because the active medium "sees" more photons of the very specific frequency that is sent back into the diode. By observing the spectrum with the Mightex IR1 spectrometer, we can easily see this effect. Just as in the previous section, insert the cosine corrector of the fiber cable into the beam of the laser and observe the spectrum on the computer. Make sure that you are not saturating the CCD camera in the spectrometer (the maximum reading at the highest peak of the spectrum should be no more than about 30,000. If it is too high, move the fiber a little bit to the side so that less light hits it.).

Discuss with your group members what exactly you are looking for and how the spectrum will change when the diffracted beam is sent perfectly back into the laser diode. Summarize the results of this discussion in the box below, and feel free to check in with your instructor to confirm that you're on the right track.

In order to move the beam systematically, you should use a good strategy. Here's what typically works well: To begin, use the screws to move the diffracted beam such that it is located slightly higher and slightly to the left of the laser beam (you can check this with the detector card). Because the diffracted beam is slightly off to the side and slightly too high, we know that the laser doesn't see optical feedback yet, so the spectrum should show you the 'free-running' or 'natural' wavelength of this particular diode. It should be close to 780 nm, but could be off by a few nm. It will probably also show multi-mode behavior or mode-hops. Because we know exactly where the diffracted beam is located with respect to the actual laser beam, we can now move it in a very controlled way using the two screws on the grating mount as shown schematically in the figure below. First, move the up-down screw slowly and steadily until the beam is now a bit too low (and still too far to the left since you haven't touched the left-right screw yet). You should check with the detector card to get a feeling for how much you have to rotate this screw. Then use the left-right screw to move the diffracted beam just a little bit to the right. Again use the up-down screw, but this time move the beam back up until you reach approximately the height you started at. Use the left-right screw to move the beam a bit more to the right, then use the up-down screw again to move it down, and so on. The two bolded red lines in the figure show you the positions of the diffracted beam at which you should see some effect of the optical feedback.

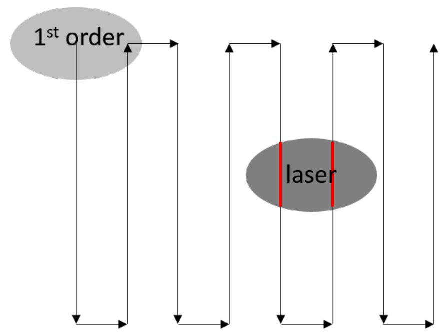

Doing this procedure just right takes some practice (and sometimes some luck too!), and there are multiple things that can go wrong. Here are a few hints that can help you find optical feedback:

1. It is very easy to move the screws too much or too quickly and skip over the perfect alignment, especially when you've never done something like this before. Thus, make sure that you move the left-right screw only a *tiny* amount each time, maybe a sixteenth of a turn or less. The alignment is especially sensitive to the up-down motion; that's why you use that screw to scan continuously over a fixed left-right position. Make sure to move that screw slowly and steadily.

2. It is very easy to lose track of what you just did and what you need to do next. It can help to announce loudly to your group members which screw you are moving and in which direction. Have them keep track of this as you tweak the alignment.

3. In order to make this process as reproducible as possible, you want to ensure that you rotate the screw about the same amount each time, especially the up-down screw, so that you end up pretty much at the same height at which you started, and you don't walk the beam off too much. It helps to place your hand on the table and hold it in a fixed position, and then rotate the screws with your fingers only. That of course won't give you a lot of rotation, but that's okay, because the amount of rotation is surprisingly reproducible that way. To rotate the screw more simply count how many times you've touched the screw, rotated, let go of the screw, touched it again, rotated, let go again, etc. That gives you an idea of how far the beam has moved. So, for example, if you start with the diffracted beam in the upper left corner of the original beam, you can count to - say - ten rotations or so as you steadily move the beam down. Then move the beam a tiny amount to the right, and again count to ten rotations as you move the beam back up. It's not a foolproof method, but in most cases it's actually pretty good!

4. Alternatively, you can make use of the white dot on the thumb screw on the mirror mount. Record its orientation, then rotate the screw until you know that your diffracted beam has definitely covered the beam diameter and record this new orientation. Then rotate between these two angles. This method works especially well when you have two people, one to watch the spectrum and the other to rotate the screw and monitor the position of the white dot.

Once you have found the optical feedback, adjust the up-down and the left-right screws and set them roughly to the middle of range over which you maintain optical feedback. Before we move on, let's confirm that you indeed have achieved good optical feedback. Since optical feedback sends more photons back into the laser diode, it actually lowers the threshold current (you'll confirm this experimentally in a moment). So, reduce the current until you are just barely above the threshold current. Then rotate the up-down screw a little. If the laser stops lasing your alignment was indeed perfect. Move the up-down screw back and leave it in the middle of the range over which you see laser emission.

How far can you rotate the screws before you completely lose feedback? Is it similar in both axes or different? Record your observations and any additional comments here.

Below is a summary of these last few steps:

How to build an extended cavity diode laser:

- Position the diffraction grating in front of the laser diode and roughly align the first order back into the diode. You can check the alignment by blinking between the diffracted beam on a detector card and the original beam passing through the back of the same card (it'll appear fainter and blurrier). Tighten the grating mount.
- Use the mirror mount screws to position the diffracted beam slightly to the side of the original beam (e.g. above and to the left).
- Observe the spectrum of the zeroth order diffracted beam with the Mightex spectrometer.
- Use the up-down and left-right screws to scan over the beam emerging from the laser diode in a grid-like pattern (down-right-up-right-down-right etc.) until you see the effect of the optical feedback on the spectrum.
- Adjust both screws until they are in the middle of the up-down and left-right scanning range.
- Confirm that you have good optical feedback by lowering the current to just above the lasing threshold and moving the up-down screw. The laser should stop lasing, and start again when you move the up-down screw back to where it was.

Now let's characterize this new laser. Record the output power as a function of the laser current.

Sketch this new data set and compare it to the one you took without the grating. What is similar, what is different? Explain the difference by remembering that the laser "sees" more photons when the diffracted beam is sent directly back into the laser diode.

Next, observe the **spectrum at a fixed temperature and a current of about 30 mA**. Gently rotate the up-down screws and observe the spectrum closely.

Does it change? If so, how?

Move the up-down screw back to the middle range and then gently rotate the left-right screw. Again, observe the spectrum.

Does it change this time? If so, how?

Explain your observations! It helps to look at the orientation of the grating and think about the direction in which the laser beam is split into its wavelength components.

Adjust the up/down screw so that you're in the middle of the range.

Then rotate the left-right screw and record the wavelength range over which you only see a single peak in the spectrum that follows nicely as you rotate the left-right screw. Over this range you have perfect optical feedback and thus perfect control over the emission of the laser diode.

Lastly, record the *wavelength as a function of the current through the laser diode for one fixed temperature* and compare it to the data set you took earlier without the grating. *Use a temperature that you used previously to make this comparison more meaningful.* Also use the same current steps (about 2 mA) when you take your data. Follow the instructions starting on page 17.

Chances are you saw some mode hops. This makes sense, because by adding the grating we add yet *another* cavity to the whole system (therefore the name 'extended *cavity* diode laser'!).

Sketch this, and explain how this additional cavity affects the standing waves that are formed. Thus, do you expect the ECDL to have more or less mode hops than the free running laser diode? Does this agree with your data on the previous page?

Your laser is now (almost!) ready to interact with rubidium atoms (you will do that in the advanced version of this module.

Let's play with it!

If you have time, align the diffraction grating in the Littman-Metcalf configuration. Again, find feedback (but this time use the additional *mirror* to move the diffracted beam, not the grating!). Repeat the previous exercises and record your observations and data in this box (and maybe some additional piece of paper).

Even if you don't have time to do the experiment, answer these questions: What are the advantages and disadvantages of each of these two configurations? When would you choose one over the other?

And that's it! You made it successfully through the first module and are now qualified to perform the more advanced modules. Please leave us any comments, suggestions, or concerns in the box below, so that we can optimize this module for future student generations. Thanks!