

Mirror Mount External Cavity Diode Lasers

Peter Heuer and Jeffery Kleykamp

Summer 2013

1 External Cavity Diode Laser Design




Figure 1: A closeup of the mirror mount laser's external cavity

An external cavity diode laser (ECDL) uses optical feedback to force a commercial laser diode to emit a narrowed and tunable range of frequencies. The beam originates at one side of the cavity in a commercial laser diode (figure 1). This beam is then incident on a diffraction grating at the other end of the cavity. The first order diffraction beam is directed back into the laser diode, forming the cavity. However, the diffracted beam is spread out with an angular dependence on frequency, allowing only a small range of frequencies to be fed back into the diode. Since this feedback causes the stimulated emission in the gain medium of the diode, the laser lases only over this narrow range of frequencies. The laser frequency can be tuned by rotating the diffraction grating to change the range of frequencies entering the diode.

Finer adjustments to the lasing frequency can be made by changing the length of the laser cavity. The diffraction grating is mounted on a piezoelectric actuator with a range of around a μm . Changing the cavity length changes which frequency modes the cavity can support, shifting the output frequency of the laser. Scanning the cavity length by applying a sawtooth wave to the actuator allows the laser frequency to be swept back and forth rapidly.

This technique is used to stabilize the output frequency of the laser over time, a process known as frequency locking. A sample of the laser's output beam is used to generate an absorption spectra. A circuit monitors this signal as a measure of the laser's frequency and responds to drifts by changing the voltage of the piezoelectric actuator (and thus the cavity length) to compensate. This process helps hold the laser frequency constant against temperature fluctuations and vibrations.

The mirror mount ECDL design described here is based on of a design published by A.S. Arnold, J.S. Wilson, and M.G. Boshier (1998). Using a mirror mount as a body provides a very

mechanically and temperature stable laser with a very short cavity length. Shortening the length of the cavity increases the free spectral range (FSR), the effective range of frequencies over which the laser can scan without jumping to a new cavity mode. These characteristics make the laser better behaved and thus easier to work with.

This design is also easier to maintain than other external cavities. The laser diode is held in a tube mounted in a precisely fitted hole that is accessible from the back of the mount. This allows the diode to be replaced without misaligning the cavity.

The prototype laser described in this document has successfully met these design goals. The laser is very resistant to changes in temperature and vibrations, minimizing frequency drift. For greater stability, the laser can remain frequency locked for days at a time. The short cavity length allows continuous scanning over at least 5 GHz.

The greatest weakness of this design appears to be its large thermal mass. Since the temperature of the entire mount is controlled, the system has a great deal of thermal inertia and responds slowly to rapid changes in temperature. Cooling the mount also creates substantial waste heat, which makes the thermal design of the laser's base and enclosure important.

2 Machining




Figure 2: The fully assembled mirror mount laser prototype

Select a mirror mount with a thick back plate and a flat front plate. The back plate must have a solid area large enough to accommodate a hole for the diode collimation tube. For the prototype lasers described in this manual, we used a Lees mirror mount. Another suitable choice would be a Newport U200-P with a UPA-PA2 adapter to hold the laser diode (Arnold et al. 1997) ¹.

The diode collimation tube was mounted in a hole in the back plate of the mount (figure 3a). This hole should be drilled as close to the edge of the mount as possible so that the output beam of the cavity will have ample clearance. The diode collimation tube used in this prototype (Thorlabs LT230P-B) has a diameter that does not match a standard drill size. To machine the hole, a smaller hole was drilled in the mount and then widened with a boring tool on a mill. This

¹This paper suggests a U100-P mount with a UPA-PA2 adapter, which is a 1" rather than a 2" mount and would be a tight fit with the approximately .85" diameter Thorlabs collimation tubes.




Figure 3: Closeups of machined parts

technique allowed the hole to be sized within several thousandths of the collimation tube's actual diameter. The inside of the hole was then sanded with emery cloth until the tube fit perfectly inside. To maximize thermal contact between the tube and the mount, the outer surface of the tube was coated with thermal grease before being inserted. The tube was secured inside the mount with a #10-32 set screw in a hole drilled in the top of the mount.

A small hole to hold a thermistor was drilled as close as possible to the bottom of the hole for the collimation tube (figure 3a). This hole should be deep enough that the thermistor is in the middle of the mount.

The bulk of the front plate of the mount was cut away to make room for the grating (figure 3b). Be careful not to damage any of the mounting equipment on the front plate while cutting.

The grating mount (figure 3c) should be machined as close as possible to the appropriate Littrow angle. This was done by clamping a piece of aluminum to a sine plate set to approximately 27.9° and then milling down the surface. The precise dimensions of the mount are not critical. For this prototype, the grating mounts were approximately $2'' \times \frac{5}{8}'' \times \frac{1}{2}''$.

Throughout the machining process it is advisable to cover the ends of the threaded bushings on the mirror mount with tape to prevent particles of metal getting in and damaging the threads.

3 Assembly




Figure 4: Closeups of assembled parts

After the mirror mount had been reassembled, the rest of the components were epoxied into place. First the piezoelectric actuator was attached to the grating mount. Be sure to use a small bead of epoxy that will not press out and cover the sides of the actuator, inhibiting its movement. Then glue the grating onto the piezo. The grating should be oriented so that the arrow indicating its blaze direction points in the direction of the output beam (figure 4a).

Finally the grating mount was epoxied to the front plate of the mirror mount. The alignment

of this joint is critical, as it determines the position and angle at which the laser beam will hit the grating. It is easiest to align the arm by laying the mirror mount on it's back and looking up through the collimation tube hole while the epoxy dries. The grating should be positioned so that the beam will strike its center but the mounting arm is perpendicular to the side of the mirror mount.

A 20 k Ω thermistor was cemented into the mounting hole using a thermally conductive epoxy (figure 4b). This epoxy was prepared by mixing very finely ground silver particles into a batch of regular five minute epoxy. Ideally the epoxy should fill the entire hole around the thermistor to maximize thermal contact with the mount.

The entire laser was attached to an aluminum base with a screw that threads into the bottom of the mirror mount. On the prototype laser, 30W thermoelectric coolers (TECs) (figure 4c) were placed in between the laser and the baseplate. Two TECs were used because no TECs were available with a hole large enough to accommodate the screw that attaches the baseplate to the mount. A different arrangement of TECs could be used as long as they mechanically support the laser and are sufficiently powerful to control its temperature.

In order to cool the mount effectively, the TECs must have a way to dispose of heat. When sealed inside the aluminum enclosure the only way for heat to exit the system is through the aluminum baseplate. During the final assembly of the laser, the bottom of the mirror mount, the TECs, and both sides of the laser base plate should be coated with heat sink paste to allow heat to flow freely.

Finally, pieces of sorbothane were placed underneath the corners of the enclosure to provide vibration isolation.

4 Collimating and Aligning the Laser Diode

In order to collimate the beam, place the laser diode in a mount to hold it steady. Turn on the diode and direct the beam along a free path of several meters. Use an IR card to observe the diameter of the beam along this path, and adjust the lens inside the collimation tube until the beam size is uniform along the path.

The collimation tube can now be inserted into the laser. Since the diffraction phenomenon depends on polarization, the output power of the laser is dependent on the angular position of the collimation tube. To align the tube properly, rotate the tube until the power in the output beam of the laser cavity is maximized.

In order to maximize the mode-hop free tuning range of the laser, the length of the external cavity should be approximately .6 times the length from the pivot point to the center of the grating (Saliba et al., 2009) for a Littrow angle of 27.9°. The latter distance is fixed, while the cavity length can be controlled by moving the collimation tube in or out. For our prototype laser, the front of the collimation tube was set flush with the front of its mounting hole.

5 Aligning the Feedback Beam

The goal of aligning the feedback beam is to aim the 1st order diffracted beam from the diffraction grating back into the laser diode. First, roughly align the beam by holding an IR card inside the cavity. It may be easiest to purposefully misalign the horizontal direction so that you can more easily align the vertical direction. Then bring the two beams as close together as possible until they appear to be on top of one another.

As the feedback beam approaches the diode, you may see a second reflected beam in the output of the laser. This second spot is the reflection of the feedback beam off of the metal casing of the diode. Overlapping this beam with the output beam of the laser will bring the laser very close to the proper alignment.

The alignment can now be fine-tuned by 'flashing' the laser. Turn the laser current down to just below the diode's threshold current and make small adjustments to the alignment of the

grating. When the feedback beam enters the diode, it will decrease the threshold current, causing the diode to suddenly become much brighter. When close to feeding back, the output power will become extremely sensitive to vertical adjustments.

If flashing the laser is difficult, a power meter or photodiode can be used to make the process easier. With the feedback beam misaligned, turn down the diode current until the output power is approximately $200\ \mu\text{W}$. This current will be below the threshold current. Then adjust the alignment of the grating to improve the power out without changing the diode current. Be careful: shifting the grating will change the direction of the output beam possibly shifting it off of the power meter, which will look like a drop in beam power. When well aligned, the output power should reach close to 2mW .

Once the laser has been flashed, reduce the current slightly and make small changes to the grating to re-optimize the output power. Once you have reduced the current by several mA from the original threshold current, the feedback beam should be well aligned.

6 Getting the Laser On-Line




Figure 5: A line in a rubidium cell as seen through an IR viewer

Once the laser is feeding back, the laser must be tuned to the appropriate range of frequencies by adjusting angle of the grating in the horizontal plane. In order to adjust the laser to match an atomic transition frequency, the laser is shone through a cell filled with a vapor of the atom in question. When the laser's frequency passes over the transition frequency, a line will be visible (either by eye or through an IR viewer) through the cell (figure 5). When this occurs the laser is said to be 'on-line'.

The frequency of the laser beam is influenced by the temperature of the diode, the angle of the grating, and the length of the cavity. As the temperature of the diode is lowered the diode contracts, shortening the natural wavelength (raising the frequency) of the diode. The diodes used in the prototype laser were designed for 785nm , so they were cooled in order to reach their target of 780nm . The temperature of the laser was measured with a $20\text{k}\Omega$ thermistor, where higher resistance corresponds to lower temperatures. The ideal operating temperature varies from diode to diode, but we had success in the range $13\text{k}\Omega$ to $17\text{k}\Omega$ (where room temperature is approximately $11\text{k}\Omega$). Beware that running a diode extremely hot or cold may reduce its lifetime or cause problems such as condensation.



Once the laser is set to a reasonable temperature, apply a sawtooth wave to the piezoelectric actuator to rapidly scan the cavity length. This will effectively broaden the range of frequencies emitted by the laser, making it easier to stumble across a line. We used a frequency of 50Hz with

an amplitude of several volts. Use the horizontal adjustment knob on the mirror mount to rotate the grating back and forth, watching the cell for a line to appear. Be careful not to rotate the knob too far, as doing so will reduce the feedback you previously aligned. We found that a single rotation of the adjustment knob in either direction was an acceptable range.

If no line appears, adjust the diode current (in steps of 5-10mA) to a different value and try again. If no current and grating combination works, adjust the temperature (in steps of 1k Ω) and repeat the process.

When all of these techniques fail, it may be necessary to measure the wavelength of the laser with a wavemeter. The laser can then be easily adjusted towards the desired wavelength.

7 Locking the Laser



(a) The control panel of the locking circuit used to test the prototype laser.

(b) A scan through several rubidium absorption peaks.

Figure 6

The following directions are given assuming the signal has been displayed on an oscilloscope in X-Y mode with the PZT voltage on the X channel and the spectroscopy signal on the Y channel.

The first step of locking the laser is to identify the slope that you wish to lock on to. The larger and steeper the slope is, the easier and more stable the lock will be. If the feature you need to lock to is not large enough to lock easily and/or stably, different spectroscopy methods such as DAVLL (dichroic atomic vapor laser locking) can be used to generate a better signal.

Once the target slope has been identified, turn the ramp gain down while shifting the DC offset as necessary to narrow the scanning range to include just the slope you want to lock to. Slowly turn up the feedback knob, noting which direction the signal tries to "run" to. Turn the feedback back to zero, then change the set point. The direction you need to change the set point knobs depending on the direction the signal runs may depend on the lock box. Repeat this process until the signal begins to run the other way. When this happens, slowly continue the process in the other direction. Your goal is for the signal to not run either direction when the feedback is increased. If your lock box has a fine set point adjustment knob, you can use it in the same manner once the regular setpoint knob has become too coarse.

Once the signal has stopped running, turn the ramp gain all the way down (so that the laser is no longer scanning) and turn the feedback all the way up. The oscilloscope should now show only a bright dot. The laser is now locked.

If the signal refuses to settle on the slope you have chosen (or seems driven away from the slope) try flipping the invert switch on the lockbox. This switch controls whether the lockbox locks to an increasing or a decreasing slope.

The easiest way to make sure your laser is locked is to clap your hands, drop a small weight, or otherwise create some vibration that affects the laser. If the laser is locked you should see a quick flash of circles around the dot on the oscilloscope screen. To test the strength of the lock, create increasingly strong disturbances until the laser falls off lock. When this happens the circles will flash quickly then disappear and the dot on the oscilloscope screen will be unresponsive to further disturbances.

If you are careful, you can adjust the frequency of the laser while it is locked by slowly tuning the setpoint. Be aware that, should you tune the setpoint off the slope you originally locked to, the laser will fall off of lock.

8 A Brief Encyclopedia of Rubidium Absorption Lines




Figure 7: This published spectra (from the Georgia Tech Sat. Spec. lab manual) is helpful for identifying transitions

The easiest way to distinguish different rubidium absorption lines is to compare the shapes of the spectra observed to a published spectra like the one shown in figure 7. While each line has a distinct appearance, the spectra can look drastically different depending on your laser, oscilloscope, and spectroscopy settings. Once you have found a line, try to optimize the spectra by eliminating mode hops and scaling the oscilloscope properly in order to avoid misidentification.




Figure 8: Rb87 $F=2$

The Rb87 $F=2$ line is the cooling transition for most rubidium magneto-optical traps. The laser should be locked to the left side of the small peak at the deepest part of the dip.




Figure 9: Rb85 $F=3$

The Rb85 $F=3$ transition is most easily identified by its lopsided dip, as well as the small dent at the bottom of the right-most dip. When viewed in a cell, this transition is almost always the brightest line you can see.




Figure 10: Rb85 $F=2$

The Rb85 $F=2$ transition is easily confused with the Rb85 $F=3$ peak. While both have two dips side-by-side, the $F=2$ transition's peaks are of equal length while those of the $F=3$ are lopsided.




Figure 11: Rb87 $F=1$ (Note that this image has been enlarged: this transition is much smaller than the other peaks).

The Rb87 $F=1$ transition is easiest to identify by the process of elimination. Scan through the spectra until a (relatively weak) line appears in the cell where no obvious peak is visible in the absorption spectra. Then make small adjustments to try and bring the line into view. Its most characteristic feature is its size, along with the jagged appearance of the hyperfine transitions.