Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

- Atomic Track -

Rubidium Spectroscopy using a Tunable Diode Laser

Introduction

The Rubidium Spectroscopy Lab provides an introduction to the field of laser spectroscopy, in which frequency-tunable laser light is used to probe the electronic structure of atoms, molecules, and a wide variety of condensed-matter systems. Our focus will be on the spectroscopy of neutral rubidium atoms, as this is a reasonably simple experimental system for gaining some experience working with laser hardware in the physics lab.

A typical schedule for completion is:

Weeks 1 & 2: Rb absorption spectroscopy. Complete measurements of temperature-dependent and intensity-dependent atomic absorption and their associated modeling. Submit your Assignment 1 e-notebook at the end of Week 2.

Weeks 3 & 4: Saturated absorption spectroscopy. Complete observations of saturated absorption and measure the associated Rb hyperfine splittings. Submit your e-notebook at the end of Week 4.

The hardware for this lab is in Room 209 (Figure 1), which is somewhat isolated from the rest of the Ph77 lab for safety reasons. You will be using a near-infrared laser operating at 780 nm, which is invisible to the naked eye. For this reason, you must wear safety goggles whenever the laser is on.

Rubidium atomic structure

Rubidium has two naturally occurring isotopes, ^{85}Rb (72 percent abundance), with nuclear spin quantum number I = 5/2, and ^{87}Rb (28 percent abundance) with I = 3/2. Although ^{87}Rb is not

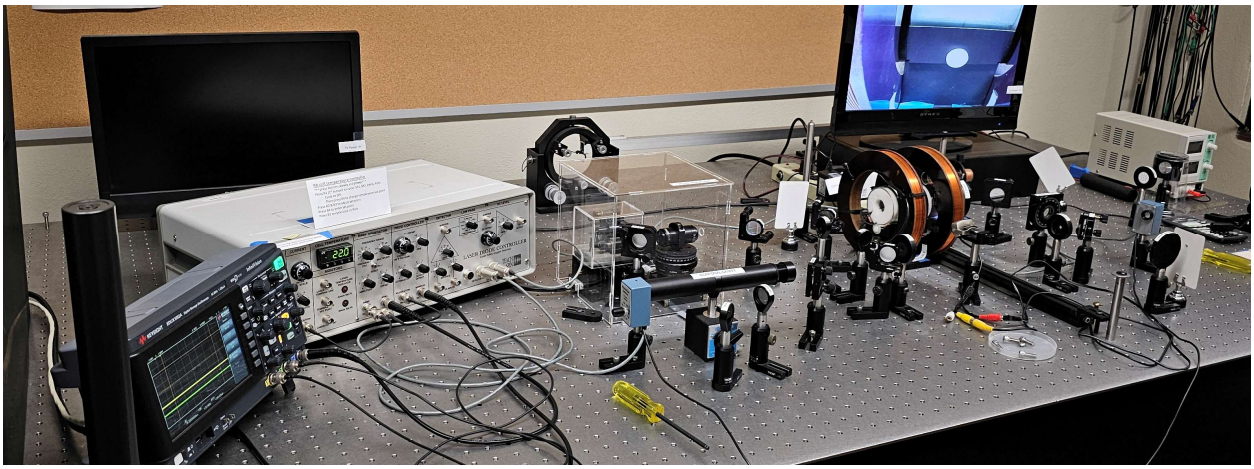

Figure 1. The Rubidium Spectroscopy setup, located in Room 209 East Bridge.

stable against radioactive decay, it has a half-life of 4.8×10^{10} years, which is more than three times the age of the universe. The ground-state electronic configuration of rubidium consists of closed shells plus a single 5S valence electron, giving a hydrogen-like electronic level structure. For the first excited state, the 5S electron is moved up to 5P.

The different energy levels are labeled by the usual Russell-Saunders notation $^{2S+1}L_J$, where S represents the total spin angular momentum (S = 1/2 for a single electron), L specifies the total orbital angular momentum (designated S, P, D, etc. for L = 0, 1, 2 respectively … alas, this S is different from the total-spin S), and J refers to the total angular momentum. (You learned about the quantum physics underlying spin-orbit coupling in Ph125, so we will not discuss the matter further here. Consider this handout as a good opportunity to refresh your memory and take a deeper dive into a specific case study.)

For the ground state of rubidium S = 1/2 (since only the single valence electron contributes) and L = 0, giving J = 1/2 and the ground state $^2S_{1/2}$. For the first excited levels we have S = 1/2 and L = 1, giving J = 1/2 or J = 3/2. Thus there are two excited states labeled $^2P_{1/2}$ and $^2P_{3/2}$, and these energy levels are split by spin-orbit coupling. Note that Rb, like all the alkali metals, has a level structure that is essentially the same as the Hydrogen structure you studied in Ph125, again reflecting the fact that the electronic structure at low energies is dominated by that single valence electron.

Adding in the nuclear spin yields hyperfine splittings to all these electronic states, and Figures 2 and 3 show the lower S and P energy levels for ^{85}Rb and ^{87}Rb including the hyperfine splittings. Magnetic dipole interactions are often the dominant hfs term in the atomic Hamiltonian, giving an energy splitting

$$\Delta E \approx \frac{A}{2}[F(F+1) - J(J+1) - I(I+1)] \tag{1}$$

where F=I+J is the total angular momentum quantum number including nuclear spin, and A is the "hyperfine structure constant." In addition, there are significant hfs energy-level shifts that arise

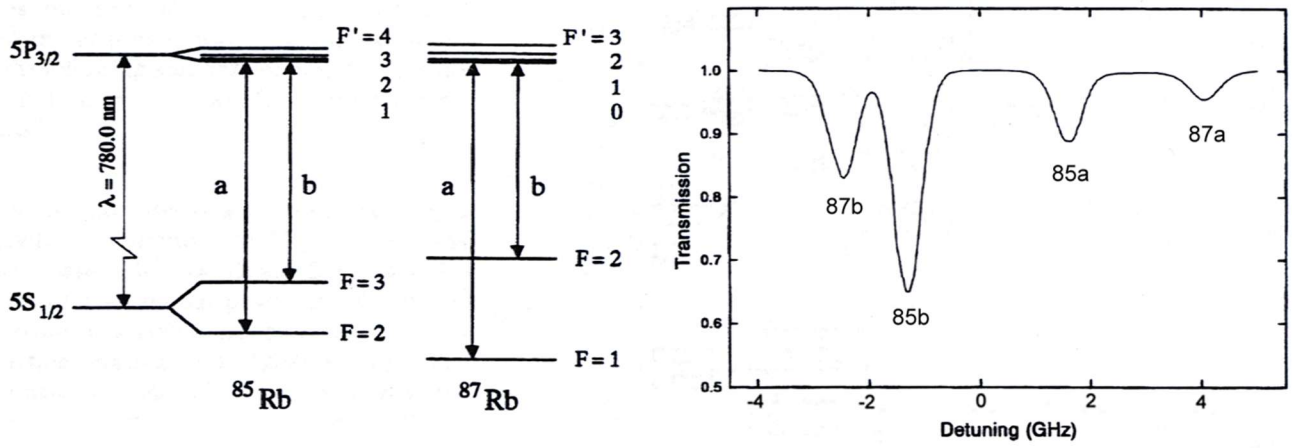

Figure 2. (Left) Level diagrams for the D2 lines of the two stable rubidium isotopes near 780 nm. (Right) Typical absorption spectrum for a rubidium vapor cell, with the different lines shown.

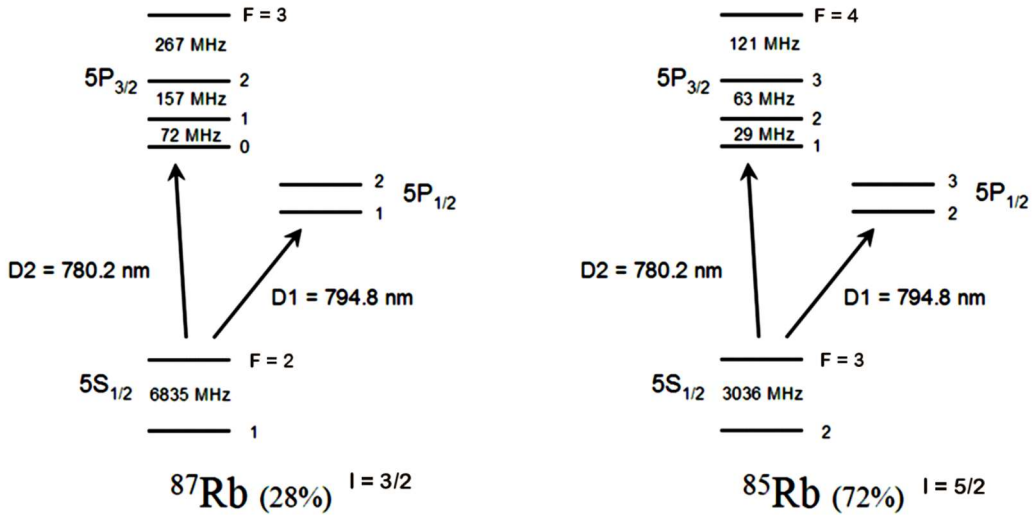

Figure 3. A more detailed rubidium level diagrams, showing the hyperfine splittings of the ground and excited states.

from interactions with the nuclear quadrupole moment. In general, multi-electron atomic energy levels are difficult to calculate because the various interactions yield a complex many-body problem that can be quite challenging. Because of these theoretical limitations, experimental measurements from laser spectroscopy are vital for making progress toward useful applications in remote sensing, chemical analysis, and other areas of science and technology.

Because our goal in this lab is not to teach you all the many intricacies of atomic energy levels, we will end our brief discussion of the topic here. Instead we will turn our attention to how one might use laser spectroscopy to measure the numbers shown in Figure 3. Gaining experience with some of these tools may be beneficial if you end up going into any area of AMO (atomic/molecular/optical) physics, which is a diverse and exciting branch of modern physics research. In addition, tunable lasers are commonly used in many laboratories for many purposes, so you may find it worthwhile to gain some experience with laser optics and related opto-electronic devices.

Exercise 1. The absorption dips in Figure 2 are mainly *Doppler broadened* by the thermal motion of Rb atoms in a vapor cell, expressed as Gaussian absorption profiles $\sim \exp[-(\Delta f/\sigma)^2]$ where $\Delta f = f_{laser} - f_{resonance}$. Assuming a cell temperature of 40 C, calculate the Doppler broadening parameter σ in MHz. [Note that this is a good exercise for ChatGPT (or your favorite AI tool); the physics is fairly simple and it involves some tedious calculations you don't need to do yourself. If you go that route, however, be sure to document your work in your e-notebook, as you would a Mathematica notebook or some Python code. In all these cases, including some screenshots is a good way to document what you have done, so you can remember it better and communicate it to others. Also, be sure to use Figure 2 to (roughly) check your results. AI tools often give you incorrect answers if you are not paying attention.]

Laboratory Exercises – Assignment 1

Your first laboratory task is to observe the Doppler-broadened Rubidium absorption spectrum illustrated in Figure 2. To begin, *put on your laser safety goggles* and start assembling the optical layout shown in Figure 4. You will likely need some help from your TA at first, because getting some one-on-one help is the fastest way to learn how to use this laser system (and there are too many operational details to include in this handout). With your safety goggles on (and only when you have them on!) turn the Laser Power switch to ON (at the left side of the laser controller; note that the power switch on the back of the chassis is always on) and set the Current dial to 5.0. This should produce a laser beam with about 10 mW of optical power, although the 780nm laser light is not visible to the naked eye (and certainly not through your goggles). The easiest way to "see" the laser beam is by using an IR Detection Card, and a few of these can be found on the optical table. In a nutshell, the active area of the IR card absorbs UV photons from the overhead lights, putting some dye molecules into a metastable excited state. When the 780nm laser photons tickle these excited molecules, they decay back down and emit visible light in the process, yielding an orange spot you can see even through your laser goggles. These cards are more expensive than they look, so please handle them with care.

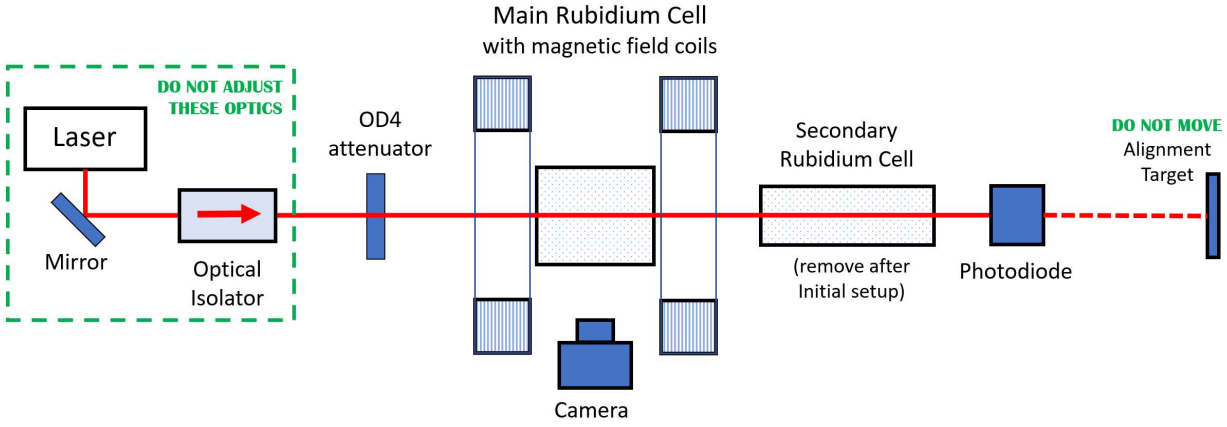

Figure 4. The optical layout for making the first observation of the rubidium absorption spectrum.

As you first use the IR card, note that the orange spot fades if you keep the laser beam on it, so it does not work to place a stationary IR card in the laser beam (try it). This fading happens because the continuous IR light "bleaches" the dye molecules, sending them all to the ground state. And it takes time for the UV light to re-excite them. Thus, the best way to use an IR card is to briefly wave it around in the beam as you follow its path. Also, as you can see, it is possible to burn spots on these cards, so try not to do that.

In former times, the overhead lights used fluorescent bulbs, which emitted plenty of UV light to excite the IR cards. Modern lights are LEDs, however, which emit very little UV. If you find that your IR card is acting a bit "tired", place it under the UV light on the optical table for a recharge. If you keep one IR card under the UV light while using the other, you will always have one ready to go.

If you want to view the IR light continuously without waving a card around, you can use the TV camera that is also on the optical table. And a photodiode detector will give you a more quantitative measure of the beam brightness, as you will soon see.

Note that the **laser** and **optical isolator** in Figure 4 are difficult to align properly, so please do not move or adjust these elements. Check that all is well by removing the attenuator, secondary rubidium cell, and the photodiode, and then confirm that the laser beam is centered (or nearly so) on the alignment target. Note that the laser beam in Figure 4 runs along a row of holes on the optical table, at a fixed height of four inches above the table surface. This is not by accident but is part of good optical alignment technique (more on that later). Note also that the main rubidium vapor cell is a cumbersome item, with its magnetic field coils and electrical cables, so this device should be pushed up against the guide rail already attached to the optical table. You can slide the cell assembly along the guide rail as needed, and it is so heavy that it need not be bolted down. The cell temperature should be near 23C at this point (as seen on the digital readout on the laser controller).

Once you have verified that the laser beam strikes the alignment target, place the photodiode (PD) using an IR card to see that the beam strikes the PD sensor as best you can. The OD4 attenuator should not be in the beam path for this step, as this optical element reduces the beam intensity by a factor of 10^{-4}. Send the PD output cable to ch1 on your oscilloscope, turn the photodiode gain (knob on the back of the PD box) to its minimum position (CCW) and view the photodiode output on the 'scope. The PD output signal is probably not very interesting at this point, but you should see the PD signal go to zero and back if you block and unblock the laser beam.

The PD output voltage is negative when light is striking the sensor (not the best design choice), but you can correct for this using the oscilloscope's *Invert* feature on ch1 (this might already be set on the 'scope). While observing the PD signal on the 'scope, adjust the position of the photodiode on the optical table (up/down and right/left) to maximize the signal voltage, which tells you that the laser beam is striking the center of the photodiode sensor. When you have the photodiode optimally placed, clamp the PD down and place the secondary Rb cell in the beam path (this does not need to be clamped down). Hold onto the fragile Rb cell using the handle or the base; do not grab onto the glass cell or the mounts touching the glass.

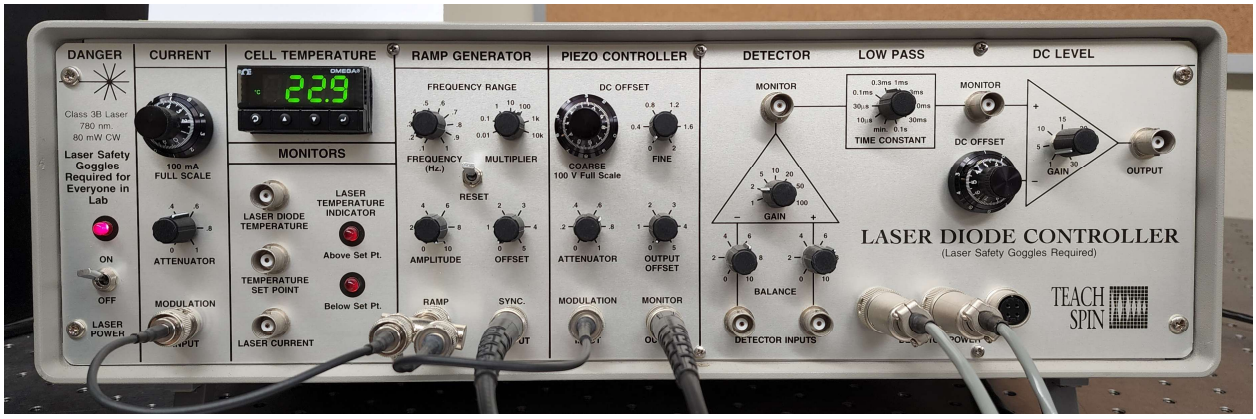

Figure 5. The laser controller in its typical configuration.

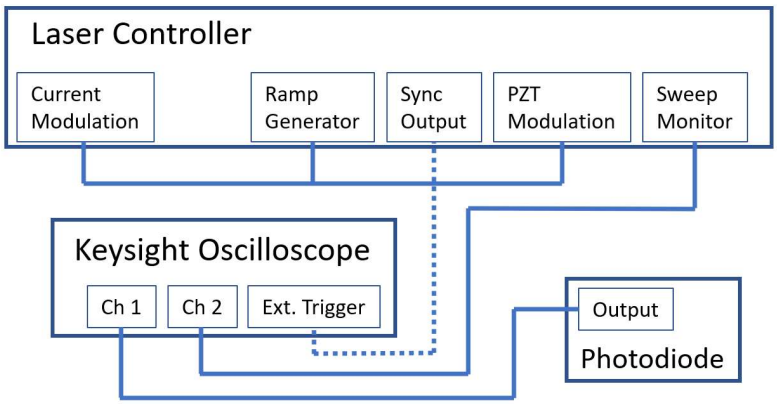

Figure 6. The electronics setup used to scan the laser frequency and observed the light transmitted through the Rb cavity. The connections shown are all BNC cables.

Set up the laser frequency scanning

Your next step is to set up the laser controller so that it continuously scans the laser frequency back and forth. Figure 5 shows the controller front panel and Figure 6 shows the electronics layout. Several of these cables and connections may already be in place. Assuming you are brand new at this (although perhaps not if you already did the Optics Track), let's walk through the settings to see what the different controls are for.

1) The *Ramp Generator* produces a triangle-wave voltage signal for scanning the laser frequency, and the Ramp output goes to both the *Current Modulation* and *Piezo Modulation* inputs. As the name suggests, the Current Modulation input modulates the diode-laser injection current, thus changing the laser brightness. The piezo (also called a PZT, short for piezoelectric transducer) moves the external grating right next to the diode laser tube (you can see this small diffraction grating inside the small acrylic laser box; the piezo only moves it a few microns). The physics behind all this is described in the Diode Lasers handout (part of the Optics Track), but for now just remember that the main purpose of this setup is to change the laser frequency by an amount that is proportional to the applied Piezo Modulation voltage. With a triangle-wave voltage, the laser frequency then scans back and forth linearly in time. This voltage also modulates the laser brightness, which is a necessary but somewhat undesirable additional effect.
2) For a reasonable starting point, set the *Current Attenuator* knob to 0.6, the *Ramp Generator Amplitude* knob to 6, the *Piezo Attenuator* knob to 1, the *Ramp Generator Offset* to zero, and set the *Piezo Controller Output Offset* to zero. Some of these are "set and forget" items, and we will not change them going forward.
3) You will often want to monitor the laser-scanning signal using the *Piezo Monitor* output, so send this signal into ch2 of your oscilloscope. Make sure ch2 on the 'scope is not set to Invert. Send the *Sync* signal into the external trigger port on the oscilloscope and adjust the 'scope so you see a stable triangle-wave signal. You will likely have to set the 'scope to Normal trigger mode, because the triangle-wave signal frequency is typically too slow for Auto mode.

4) Once you have the triangle-wave on the oscilloscope, use the Measure feature to measure the signal frequency, and then set the Ramp Generator Frequency so the 'scope measures between 12-13 Hz. Then get the 'scope set up so you reproduce the bottom trace in Figure 7, with a single half-cycle of the triangle-wave visible, with a negative slope. A half-cycle is all you need, because the other half-cycle just shows the laser scanning in the opposite direction, providing no additional useful information. The negative slope is just a convenient convention. Note the zero-voltage flags (left side of screen, with the small "ground" symbols) and the trigger point (solid triangle on the top left of the screen). Use these to help you adjust the 'scope settings. If your memory is a bit foggy when it comes to using the oscilloscope, ask your TA. This sweep trace has become our standard laser scanning setup, and you can probably leave it essentially unchanged from this point forward. The important knobs you need to remember at this point are:
5) *Current Set* (100mA Full Scale)– used to change the laser current. This requires occasional tweaking to get the laser to operate well, as we will see below.
6) *Piezo DC Offset*– Turn this knob and watch what happens to the Piezo Monitor trace on the 'scope. As the name suggests, this knob just changes the DC offset of the piezo signal. Note that the Piezo Monitor voltage should always be positive throughout its sweep.
7) *Ramp Generator Amplitude*– Again, turn this knob and see what happens. You will use this knob to adjust the range of the laser frequency scan. Turn this knob up until the Piezo Monitor signal on the 'scope has a peak-to-peak amplitude of about 4 volts.

For both the Piezo and Ramp knobs, changes in laser frequency (if all goes well) will be proportional to the Piezo Monitor voltage you see on the 'scope. Remember, however, that the laser frequency will always remain close to $f_0 \approx 3.85 \times 10^{14}$ Hz (*aka* 385 THz, corresponding to a wavelength of 780nm), because a full frequency sweep will be no more than about 10^{10} Hz.

Find the Rb absorption features

At this point the laser frequency will be scanning, but not over the desired frequency range. So power up the TV camera and TV monitor so you are looking into the side hole in the rubidium cell, as shown in Figure 4. Put the TV camera about 2cm from the side of the Rb cell, so you get a good view inside the cell. Turn up the laser current slowly (without exceeding 8.0 on the dial) while you watch the TV monitor. At some point you should see a flickering line of fluorescence inside the cell. You may need to turn the dial up to 6-7 and then maybe back down to maybe 5.5, but your results may vary. There is typically a lot of hysteresis in the laser behavior, so you will likely have to turn the current up and down a few times before getting stable fluorescence. Ask your TA for help as needed. What the camera sees here is the Rb atoms absorbing laser light and emitting it in all directions as excited atoms decay back down to the ground state. And this scattered light at 780 nm is easily detected by the silicon camera sensor.

The presence of fluorescence means that the Rb atoms are absorbing laser light, and this absorption can also be seen in the photodiode signal. Put the OD4 attenuator into the beam path at this point, and turn up the PD gain so you again have a signal on the oscilloscope that comes

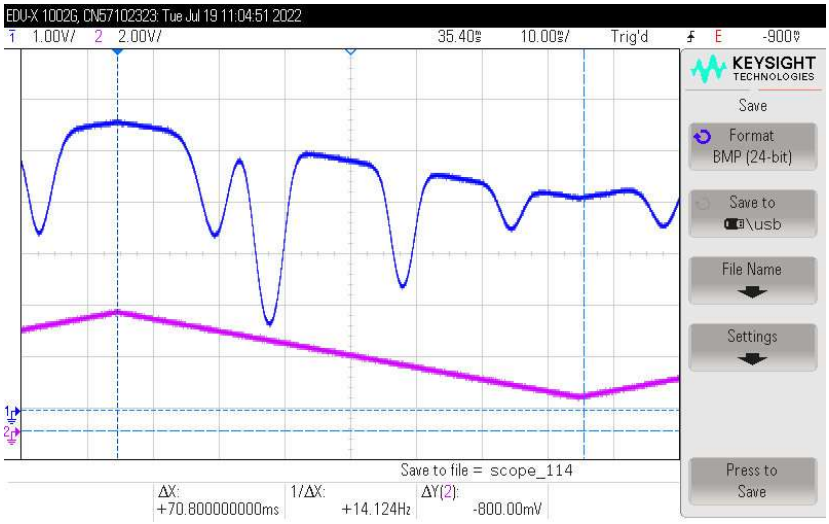

Figure 7. A screenshot (negative image) showing the four Rb absorption features, also shown in Figure 2.The bottom trace shows the voltage sent to the PZT to scan the laser frequency (note this voltage is always greater than zero volts). The top trace shows the photodiode output. Note the scale settings shown (ch1=1V/division; ch2=2V/div, time=10msec/div) and the zero-voltage flags for both channels. Note also that the trigger starts on the left side of the screen. The overall slopes of the curves should both be negative; if not, check that ch2 is not set to Invert.

and goes as you block and unblock the beam. It is usually best to keep the 'scope at about 500mV/division and adjust the PD gain accordingly. If all goes well, you should see some absorption dips on the photodiode signal, as illustrated in Figure 7.

Observe how your signal changes on the 'scope as you adjust some of the parameters described above. Turn up the *Ramp Generator Amplitude* to get a larger laser sweep, and adjust the *Piezo DC Offset* to move the absorption dips back and forth. You will probably also have to tweak the laser *Current Set* so all four absorption features are nicely displayed. Play around with these three knobs (Current Set, Piezo DC Offset, Ramp Generator Amplitude) to see how these affect the PD signal. Ask your TA for help if needed. As a useful bit of practice, try to reproduce the screen in Figure 7, so all four Rb absorption dips are clearly visible. You will want to reproduce this screen many times in this lab, so take some time now to learn how to do it. We like this screen because it nicely shows all the absorption features as you see them in Figure 2.

Once you have this screen, it is usually best to change the laser sweep by adjusting the Piezo DC and Ramp Amplitude knobs only, and *not* by changing the horizontal scale on the oscilloscope or changing the Piezo Attenuation. If you leave the ramp frequency alone and leave the 'scope horizontal scale alone as well, then you will always have a single half-cycle of the triangle-wave visible on the 'scope, with a negative slope. Just remember that the Piezo DC Offset and Ramp Generator Amplitude are usually all you need to change the laser sweep.

One of the trickiest parts of this setup is the laser current, so slowly adjust this up and down (by a small amount) to see what happens. The laser requires a certain "sweet spot" current to operate properly; otherwise it will *mode hop* or just run at some frequency far from the Rb resonance. Diode lasers are finicky that way, and the optimal current may require tweaking

throughout the day. So remember to adjust this knob from time to time, especially if your spectra look odd.

We also recommend that you show your oscilloscope screen to your TA at this point for confirmation. You may need to adjust other laser parameters to obtain a nice clean signal, and there are too many possibilities to write down here. It takes some experience to know when the signal looks wrong, and what settings might be suboptimal. Also, sometimes a spectrum looks okay but still has some subtle problems that should be addressed before moving forward.

Exercise 2. When you are satisfied with your absorption signal, save a screenshot (a .png file) similar to Figure 7 to the USB drive on the oscilloscope (use the Save/Recall button on the 'scope to set this up). Add this image to your e-notebook, along with a description of your lab setup. You are welcome to reproduce figures from this handout (e.g. Figure 4) if you wish, but this is not necessary. Include as much detail in your e-notebook as you think is required to describe your work. There is no fixed rubric itemizing exactly what you need to do for this and the other lab tasks, so treat them like small research projects. This is the advanced physics lab, so describe your work, your analysis, and your results as you see fit.

Exercise 3. With the same absorption signal on the 'scope, save the numbers to a .csv file as well. Using your own computer, save this data and plot the absorption spectrum. Use the known Rb hyperfine splittings in Figure 3 to calibrate the frequency axis, so you can plot the absorption signal as a function of the Δf, the change in laser frequency. If all went well, you should get a plot that looks something like Figure 8.

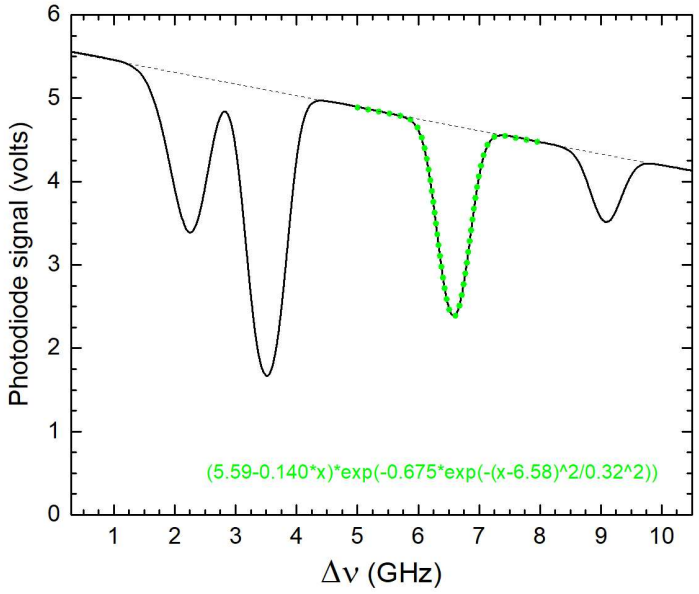

Figure 8. This plot shows the laser absorption as a function of laser frequency with the beam passing through the main and secondary Rb cells at 24C. The known hyperfine splittings were used to convert the oscilloscope time axis to laser frequency change in GHz.

For data analysis and plotting you might use Mathematica, Python, or some other software of your choosing (but not Excel, please; that is a terrible choice for scientific data.) Many of the plots in this document were made using a package called Microcal Origin (for Windows), and a free 6-month trial version is available if you are interested. The Mathematica Notes document on Canvas shows how you can do the analysis using Mathematica, and ChatGPT is a useful resource whenever writing code. And yes, it generally takes some effort to obtain quality data plots, but this is a useful skill to master. Presentation quality makes a difference whatever your career choice, because you have to communicate your ideas and results to others. And a good place to start is with nice-looking data plots; this is an example of doing the easy parts well.

While you still have the Rb absorption spectrum on the oscilloscope, try turning the photodiode (PD) gain knob to its minimum value and then turn up the 'scope gain to look at the resulting signal. With so little light hitting the PD sensor, and the PD gain so low, the PD output signal has a low voltage with quite a bit of noise. When you subsequently amplify this with the oscilloscope, both the signal and noise get amplified together, giving a poor signal-to-noise ratio (SNR). The oscilloscope's Average feature can clean up the noise (try it), but this is the wrong solution. Better is to keep the 'scope at about 500mV/division and turn up the PD gain accordingly – doing this improves the SNR without the need for averaging. (Don't forget to set the 'scope back to 500mV/div before moving on to collect more data!)

One last trivial thing to observe: remove the secondary Rb cell and note that the absorption dips become less pronounced. No surprise here, as both glass cells have the same Rb vapor density inside, because the vapor pressure in both cases is in equilibrium with solid Rb at room temperature. More on this later.

Modeling Rb absorption data

Your next step is to try some detailed modeling of the absorption data in your .csv file, as illustrated in Figure 8. In deconstructing this plot, the first thing to consider is the slope in the no-absorption signal (away from the absorption dips), modeled by the thin dotted line in the plot. This happens because the laser current and piezo voltage are scanned simultaneously using the same triangle-wave input. Doing this is necessary to record a full sweep of all four absorption dips without any laser mode hops, which is a common feature of tunable diode lasers (studied further in the Optics Track). This double modulation is something of a technical point with our laser system, but the end result is that scanning the laser injection current produces a change in laser intensity during the frequency scan. In Figure 8 we simply drew a straight line through the data to match the no-absorption data.

The atomic absorption is usually described in terms of the optical depth of the absorption cell, defined by the relation

$$P_{out} = P_{in} e^{-\tau} \tag{2}$$

where P_{in} is the input laser power (say in mW) going into the cell and P_{out} is the power exiting the cell. If you think about the physics a bit, you can convince yourself that τ is proportion to the

column density of gaseous Rb atoms in the cell. More specifically, τ is proportional to the column density of atoms that can absorb laser light, which means that τ depends on frequency because of the Doppler shift, as described above. Thus we can write

$$\tau(\nu) = \tau_0 \exp\left[-\frac{(f-f_0)^2}{\sigma^2}\right] \tag{3}$$

where f is the laser frequency, f_0 is the Rb transition frequency, σ is the Doppler width, and τ_0 is the optical depth at $f = f_0$. This gives the photodiode signal

$$P_{out} = P_{in} \exp\left\{-\tau_0 \exp\left[-\frac{(f-f_0)^2}{\sigma^2}\right]\right\} \tag{4}$$

which, amusingly, has an exponent within an exponent.

This model strictly applies to 2-level atoms, while Rb has multiple hyperfine states that overlap within the Doppler profile (see Figures 2 and 3). As a result, the actual absorption profile will be somewhat distorted from a perfect Gaussian form, and σ will be a bit higher than calculated from the thermal velocity distribution. Nevertheless, as seen with the green dotted line in Figure 8, a fit to an absorption feature using this model reproduces the data quite well.

A few notes about fitting. The green dotted line in Figure 8 was not obtained from a numerical fitting algorithm but was simply drawn through the data. A curve with the functional form in Equation (4) (plus the linear trend) was plotted along with the data in a single graph, and then the parameters were adjusted by hand until the curve passed through the data adequately. We call this a "chi-by-eye" fit, as one just adjusts the curve by eye until it looks good. Doing a chi-by-eye fit is often a good first step in data analysis, as the process gives you valuable information and physical intuition about how the fit depends on the various parameters. Also, your brain is pretty good at seeing what looks right and what doesn't, whereas nonlinear fitting algorithms can give wildly inaccurate results if you run them blind (*i.e.*, without plotting the fit with the data to do a chi-by-eye confirmation). For this lab, chi-by-eye fits are mostly all you need to get good results, but you are welcome to apply more advanced techniques to see how the results differ. Of course, if your aim was to publish this result in a scientific journal, you would want to put in the extra time and effort to obtain a suitable numerical fit.

Exercise 4. Perform this same analysis on your own data and overlay a fit with one or more absorption dips.

Comparing with Figures 2 and 3, you can see that the two deepest absorption dips come from ^{85}Rb, which is more abundant than ^{87}Rb. Beyond that, the different transitions have different strengths from the underlying quantum mechanics (complicated!), and the hyperfine splittings of the two ground states are clearly visible. The excited-state hyperfine splittings are smaller than the Doppler widths of the absorption dips, so these transitions do not appear separately in the data. There is quite a lot you can see by analyzing just a single absorption spectrum.

Temperature-dependent absorption

Your next step is to look at how the Rb absorption spectrum changes with the temperature of the Rb vapor cell, and you can then compare your results with theoretical expectations. The overall setup is the same as shown in Figure 4, and the cell temperature should start out at room temperature (about 23C). To reduce the background light on the PD sensor, you might want to turn out the overhead lights and turn on the floor lamp instead. Not only does this reduce the total room light, but the floor lamp runs on DC current, producing just a small voltage offset in the PD signal, while the overhead lights tend to flicker at 120Hz.

You might also notice that there is a narrow-band optical filter that sits at the front of the PD (it looks like a mirror). This filter blocks all light except from a narrow band around 780nm, which greatly reduces the scattered-light problem. Importantly, however, you should *not* be tempted to darken the room completely to avoid this minor background signal. Working in near-total darkness tends to be slow and unproductive, and it is easy to break things when you are stumbling around.

It is a good idea to double-check that everything is working as expected at this point, especially if you are coming back to lab after a break. With the secondary Rb cell and the OD4 attenuator in place, get the system set up to reproduce the signal shown in Figure 7. However, this time increase the sweep range a bit so the shallow 87a dip is not too close to the end of the sweep. Having the 87a feature more toward the center of the sweep will make it more stable, but do include the other three absorption dips in your sweep as well (if possible). Tweak the laser current a bit to make sure your spectrum is stable with respect to small current and/or temperature drifts.

Exercise 5. When all looks good, **remove the secondary Rb cell** and save a .csv file with your absorption spectrum. You might try averaging 4-8 sweeps on the oscilloscope for these measurements, as this gives you a cleaner signal at essentially no cost. Also record the cell temperature when you save the spectrum, which should still be about 23C. Of course, with only the main Rb cell in place, the absorption features will be substantially weaker than shown in Figure 7.

As you did before, do a fit to the 87a dip and plot your results for your e-notebook. Make sure this dip is stable (with regard to changing the laser injection current), and make sure the dip is not so close to the end of the sweep that it becomes distorted. It's okay if the 87b dip is distorted (or even not included in your sweep), because you will only be analyzing the 87a dip for now.

Exercise 6. Measure temperature-dependent absorption. If the system appears to be operating stably, begin a data run where you take additional .csv spectra with the cell at 30C, 40C, 50C, 60C, and 70C. Again, before beginning, make sure the 87a dip is clean and free from mode hops; tweak the current to make sure it looks good.

Note that changing the temperature set point from 23C to 30C will cause the controller to overshoot to ~35C, and it takes long time to cool back down to 30C. A better strategy is to change the set point to 27C, then take a data point when the system stabilizes a bit, which should be around 30C. Things generally go faster at higher temperatures, because the cooldown times are shorter when the cell temperature is far above ambient. Having the temperature exactly at 30C is not important; just record the temperature (near 30C) when you take a spectrum. This temperature

series takes time, so do not proceed unless you are confident that the system is working well. Make a plot like Figure 8 for each temperature point, and include all your plots in your e-notebook. When you are finished with the full series, turn the cell down to 60C for the next task (below), or always turn it down below 20C if you leave the lab.

Exercise 7. Note how some of the dips "flatten" at their center frequencies, indicating that $\exp(-\tau_0) \approx 0$ and the laser light is mostly absorbed and rescattered out of the beam when the frequency is on resonance. Make a fit to the 85a dip at 70C and see how well you can estimate τ_0 in this case. If the dip center does not go to zero volts, then this probably indicates some background light or a voltage offset in the detector. No problem; just include that in your model of the 85a dip.

After determining a best-fit τ_0, change your model plot to see when τ_0 is definitely too low (but no lower), and again to see when τ_0 is definitely too high (but no higher). When your brain is certain about its too-low or too-high conclusions, that usually means these values are about 3-sigma from the best fit. (Really, this works better than you might think.) From this toy analysis, report your best-fit τ_0 with an error estimate. Document your work with some plots of your too-low and too-high curves.

The goal here is to make a rough measurement using common sense observations. Often there is no set "rule" for determining measurement uncertainties in experimental physics. A voltmeter gives you a voltage, but it does not give you a measurement uncertainty. Nevertheless, all measurements have uncertainties, so you have to think about that a bit.

[Pro tip: Of course, using a sophisticated nonlinear fitting algorithm may produce more accurate numbers. But it can take a *lot* more time to get that working compared to a chi-by-eye fit. It's easy to spend days getting the fitting algorithm just right, only to find that your data were no good to begin with. Better to do a quick analysis first, see how the results look, and go back for a more in-depth analysis later.]

Exercise 8. To better understand the physics responsible for your 87a data, note that τ_0 is proportional to the Rb vapor density in the cell, and this is described by the Arrhenius equation, which is a special case of the Clausius-Clapeyron equation. You may have seen some of this when you studied statistical mechanics, but you can look it up online as well. Stat-mech theory gives the result that

$$\tau_0(T) \sim \rho_{Rb}(T) \sim \exp(-T_A/T_K) \tag{5}$$

where τ_0 is the on-resonance optical depth at temperature T, ρ_{Rb} is the Rb vapor density in the cell, T_K is the cell temperature in Kelvin, and T_A (also in Kelvin) is a measured Arrhenius constant that derives from the material properties of rubidium.

Display your results in a semi log plot like that shown in Figure 9, as this is customary with Arrhenius-like data. If all went well, your data should fall on a reasonably straight line, so use this to extract the value of T_A for rubidium. Include a screenshot showing a result you found on the internet, or from your favorite AI.

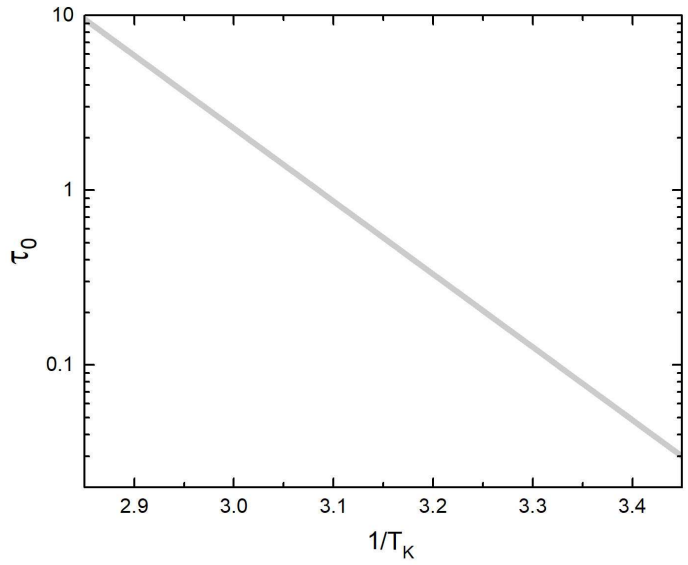

Figure 9. Plot your 87a absorption data with axes like this and use a linear fit to extract the Arrhenius constant for rubidium.

Saturated absorption

Moving on, your next task is to look at the atomic absorption as a function of the laser intensity. It is beneficial to observe the phenomenon before discussing the analysis, so first turn the cell temperature up to 60C and wait until the temperature stabilizes. With the OD4 attenuator in place and the secondary Rb cell removed, again tweak the laser current to make sure you have a nice stable absorption spectrum, with the 87a feature not too close to the edge of the sweep. And, as usual, ask your TA if you have questions. Tweak the laser current to make sure that mode hops or other distortions of the 87a dip will not be problematic.

Exercise 9. When everything looks good, save a series of oscilloscope traces (.csv files) covering a range of attenuation values: OD0 (no attenuation), OD1, OD2, OD3, OD4, and OD5 (note that

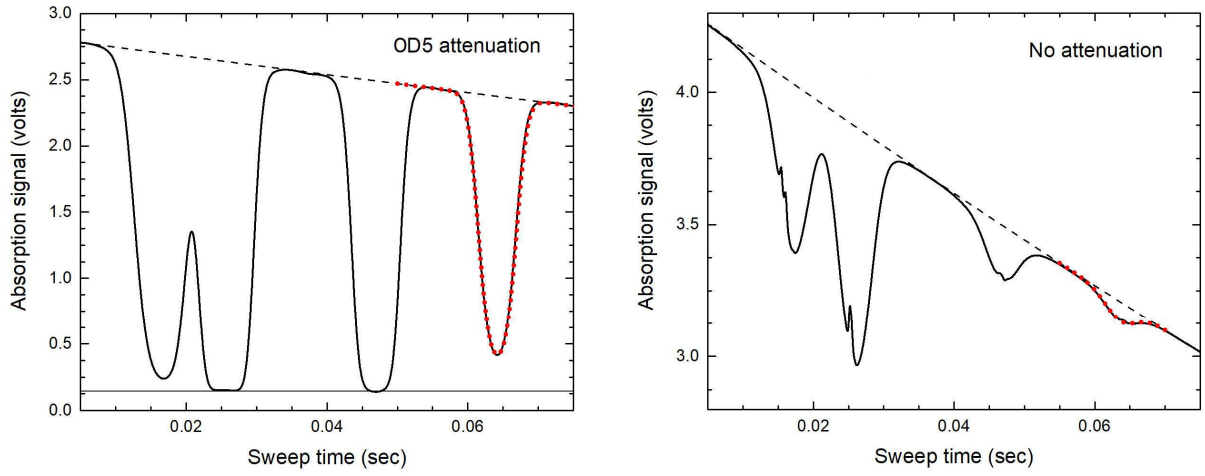

Figure 10. Example spectra showing the Rb absorption dips at 60C, with OD5 laser attenuation (left) and no attenuation (right).

OD3 = OD2 + OD1, so use combinations of attenuators as needed). For each value, you will need to readjust the PD gain and oscilloscope gain to get a good trace. As before, it is best to keep the 'scope at about 500mV/division and turn the PD gain up as needed, if this is possible. Figure 10 shows some example data when things are working well. Add all six of your plots to your e-notebook.

Exercise 10. Model your data to produce measurements of $\tau_{0,eff}$ as a function of the filter OD number (recall that ODn means an attenuation factor of 10^{-n}.) We call this $\tau_{0,ef}$ because this number is not just proportional to the Rb vapor density. For the no-attenuation case in Figure 10, note that the background is now utterly negligible, because the laser light hitting the PD is very bright. But the background signal (from room lights and/or voltage offsets in the PD) is not negligible in the OD5 case. Note the suppressed zero in OD0 graph, indicating that the absorption dips are much smaller than at OD5. The no-attenuation absorption profiles also show some fine structure that comes from the saturated-absorption phenomenon, which you will examine in detail later in this lab. Your data points will probably look something like those in Figure 11.

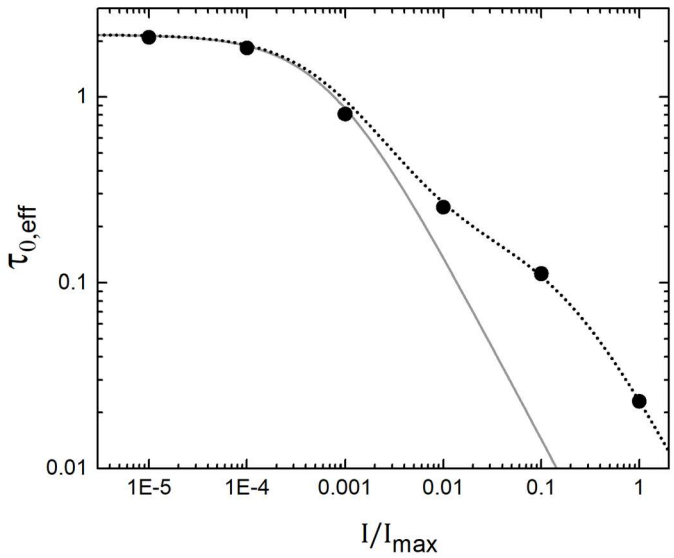

Figure 11. A sample plot of the effective optical depth of the 87a absorption feature as a function of beam attenuation. The simple theory (solid line, described in the text) does not describe the full range of data well. The higher-than-expected $\tau_{0,eff}$ values at high beam intensity can be explained by the finite size of the laser beam, while the theory model assumed an infinite plane wave.

Modeling saturated absorption

To understand what is going on here, first consider a collection of 2-level atoms at rest. The relative populations of the ground and excited states will depend on laser intensity, which can be seen by considering the (simplified) rate equations for atoms driven by on-resonance radiation, given by

$$\frac{dP_1}{dt} = \Gamma P_2 - \alpha I(P_1 - P_2) \tag{6}$$

$$\frac{dP_2}{dt} = -\Gamma P_2 + \alpha I(P_1 - P_2)$$

where P_1 and P_2 are the relative populations of the ground and excited states and $P_1 + P_2 = 1$. The first terms in both these expressions come from spontaneous emission, with Γ^{-1} being the excited state lifetime, and the second terms are from stimulated emission, with α being a normalization constant. Note that stimulated emission is proportional to the intensity I, while spontaneous emission does not depend on I. In steady state we have $dP_1/dt = dP_2/dt = 0$, and plugging this in gives

$$P_2 = \frac{\alpha I/\Gamma}{1 + 2\alpha I/\Gamma} \tag{7}$$

This result is usually written in the form

$$P_2 = \frac{s/2}{1+s} \tag{8}$$

for on-resonant radiation, where $s = I/I_{sat}$ and the value of I_{sat} is given by

$$I_{sat} = \frac{2\pi^2 hc\Gamma}{3\lambda^3} \tag{9}$$

For the case of rubidium, $\Gamma \approx 6$ MHz, which gives $I_{sat} = 2$ mW/cm^2. Atomic physicists call I_{sat} the *saturation intensity*, because the transition begins to "saturate" (P_2 becomes comparable to P_1) when I is increased substantially above I_{sat}. However, the transition is only fully saturated with $P_2 = P_1 = 1/2$ at infinite intensity. We are not deriving these atomic physics results here, because we do not have time for a full course on atomic spectroscopy. A more complete derivation of this result and its ramifications can be found in quantum mechanics texts examining atom-photon interactions in detail, but this gives you the basic idea.

Assuming that the laser intensity is independent of position in the vapor cell (not a terribly good assumption, but make it anyway), we can then write

$$\begin{aligned} \tau_{0,eff} &\approx \tau_0(P_1 - P_2) \\ &\approx \tau_0\left(1 - \frac{s}{1+s}\right) \\ &\approx \tau_0 \frac{1}{1+s} \end{aligned} \tag{10}$$

where τ_0 is the optical depth at line center in the limit of zero laser intensity (so all atoms are in the ground state).

Exercise 11. Compare your measurements with the simplified theory above, and plot the theory with your data (as illustrated by the solid line in Figure 11, where the normalization of I_{max} was adjusted to fit the data at low I). What value do you get for τ_0? How does this compare with your previous measurement of τ_0 at 60C?

We see from all this that the data fit the theory reasonably well when the laser intensity (after the OD filter) is low, but clearly the measure values of $\tau_{0,eff}$ are substantially higher than theory at high laser intensity. (The dotted line in Figure 11 was just added to guide the eye.) The descrepency arises simply because our simple model does not incorporate the the nonzero velocities of the Rb atoms, the finite size of the laser beam, and the optical depth of the Rb cell (meaning the laser intensity at the front of the cell is higher than at the back of the cell).

Looking at the high I/I_{max} behavior, note that a Rb atom placed at rest in the center of an unattenuated laser beam would be highly saturated. But when a moving Rb atom first enters the laser beam, it starts out in the ground state and requires time to reach a state of high excitation. For a brief time, therefore, these entering atoms absorb light and increase $\tau_{0,eff}$ above the expected value. The dotted line in Figure 11 shows a crude attempt at including these transient atoms, but making an accurate numerical model is quite involved, and beyond what we want to do in this lab.

Although we cannot easily model the measurements, there is an important take-away message: lasers are super bright, making it easy to "saturate" an atomic transition. When this happens, you can no longer assume that atoms are mostly sitting in the ground state, absorbing and rescattering any impinging laser light. Instead the highly saturated atoms hardly aborb any light (technically, the loss of light from absorption is offset by stimulated emission into the laser beam). Highly saturated atoms, therefore, act as if they were transparent. If you find yourself working in many fields of AMO (Atomic/Molecular/Optical) physics, you will likely encounter laser spectroscopy in some of its many forms, and you will find that I_{sat} is an integral part of the theoretical framework.

Laboratory Exercises – Assignment 2

Saturated-absorption spectroscopy

One of the most important scientific applications of AMO physics is in the area of precision atomic and molecular spectroscopy. Spectroscopy is used not only to better understand the structure of atoms and molecules, but also to define standards in metrology. For example, the second is defined from atomic clocks using the 9192631770 Hz (exact, by definition) hyperfine transition frequency in atomic cesium, and the meter is (indirectly) defined from the wavelength of lasers locked to atomic reference lines. Furthermore, precision spectroscopy of atomic hydrogen and positronium has shown that quantum electrodynamics describes quantitative experiments to better than a part in 10^8, a better absolute accuracy than any other physical theory. The technique of *Saturated-Absorption Spectroscopy* (SAS) is a simple and frequently used technique for measuring narrow atomic spectral features, limited only by the natural linewidth of the transition (for the rubidium D lines $\Gamma \approx 6$ MHz), even using atomic vapor cells with large Doppler broadening of $\Delta f_{Doppler} \sim 300$ MHz.

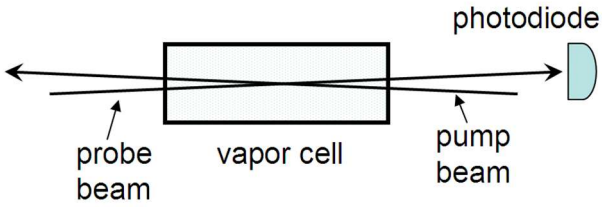

Figure 12. The basic saturated absorption spectroscopy setup. The two laser beams are tilted here for clarity; in practice the beams overlap.

Qualitative picture of saturated absorption spectroscopy – 2-level atoms. To see how SAS works, consider the experimental setup shown in Figure 12. Two lasers are sent through an atomic vapor cell from opposite directions; one, the "probe" beam, is very weak, while the other, the "pump" beam, is strong. Both beams are derived from the same laser and therefore have the same frequency. As the laser frequency is scanned, the probe beam intensity is measured by a photodetector. If one has 2-level atoms in the vapor cell, one might record spectra like those shown in Figure 13. The left plot gives the probe beam absorption without the pump beam, showing simple Doppler-broadened absorption. Here the Doppler width is much larger than the natural linewidth, of $\Delta f_{Doppler} \gg \Gamma$, and the optical depth of the vapor is fairly small, so the probe spectrum is essentially a simple Gaussian profile.

The second plot in Figure 13 illustrates the spectrum including the pump beam, showing an additional narrow spike at the atomic resonance frequency. The reason this spike appears is as follows: If we write the laser frequency as $f_{laser} = f_0 + \Delta f$, where f_0 is the atomic resonance frequency, then the probe beam is absorbed only by atoms moving with longitudinal velocity $V \approx c\Delta f/f_0$ moving toward the probe beam, because these atoms see the probe beam blue shifted into

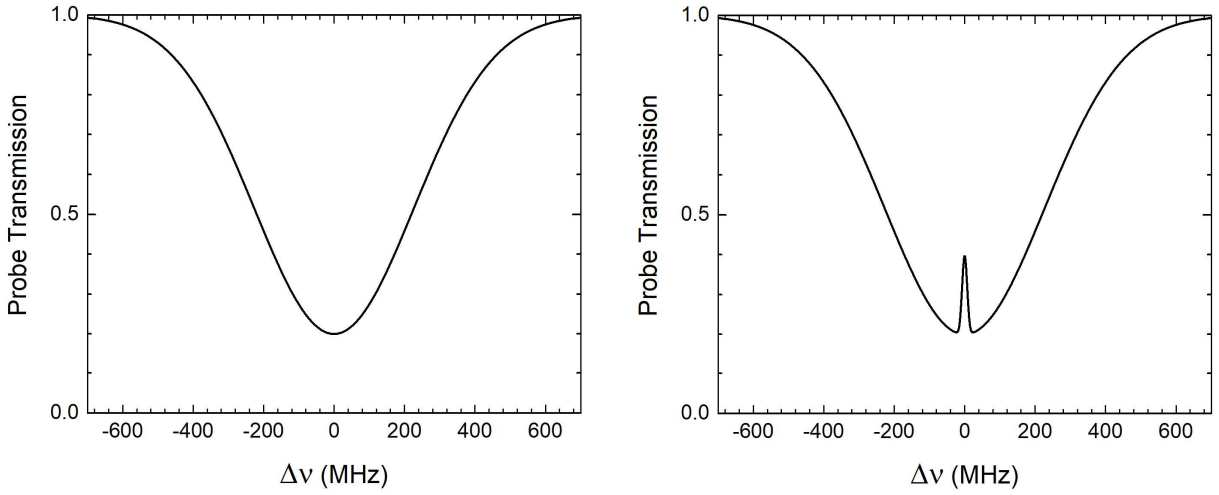

Figure 13. A qualitative model of saturated-absorption spectroscopy for 2-level atoms, illustrating the probe transmission without (left) and with (right) the pump beam. Because zero-velocity atoms see both the pump and probe lasers on resonance simultaneously, the probe absorption is diminished for only the zero-velocity atoms.

resonance. Other atoms are not in resonance with the probe beam, so they do not contribute to the probe absorption. These same atoms see the pump beam redshifted away from resonance (because the pump beam is in the opposite direction), so they are unaffected by the pump beam. Thus, for laser frequencies $\Delta f \neq 0$, the probe absorption is the same with or without the pump beam.

If $\Delta f = 0$, however, then only atoms with $V = 0$ contribute to the probe absorption. But these $V = 0$ atoms also see an on-resonance pump beam, which is strong enough to keep a significant fraction of the atoms in the excited state (so the transition is partially *saturated*), where they do not absorb the probe beam. Thus at $\Delta f = 0$ (and only at $\Delta f = 0$) the probe absorption is *less* than it was without the pump beam, because only the $V = 0$ atoms can "see" both lasers on resonance at the same laser frequency. The advantage of this form of spectroscopy becomes clear … one can measure sharp Doppler-free features in a Doppler-broadened vapor. Note that the "saturated" part of saturated-absorption spectroscopy involves the same physics you saw in your measurements of intensity-dependent absorption – a strong laser beam "saturates" the atomic transition, which in this case reduces the absorption of the probe laser.

Qualitative picture of saturated absorption spectroscopy – multi level atoms. The picture becomes somewhat richer if the atoms in the absorption cell have a single ground state and two excited states (typically an electronic level split by the hyperfine interaction), and this case is illustrated in Figure 14. Importantly, the separation of the excited states is less than the Doppler width here, as is the case for rubidium. The peaks on the left and right in Figure 14 are the ordinary saturated absorption peaks at f_1 and f_2, these being the two atomic resonance frequencies. The origin of these two peaks is the same as in Figure 13, except there are now two atomic resonances instead of just one. The two resonances have different strengths in this example, so the center of the overall Doppler dip is closer to one resonance than the other.

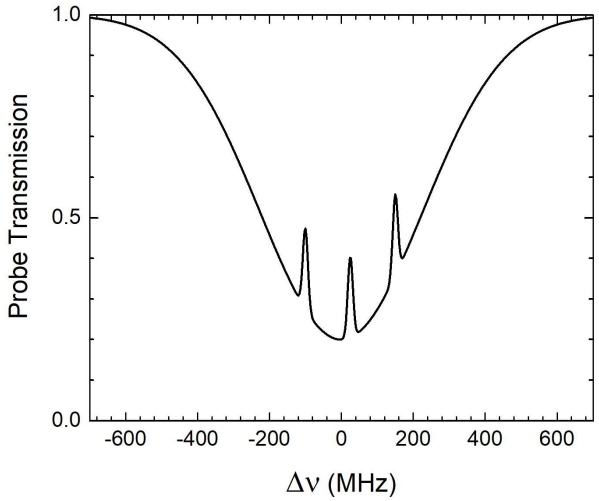

Figure 14. A qualitative model of saturated-absorption spectroscopy for atoms with one ground state and two excited states. The central "crossover peak" appears because at that Δν value one laser is redshifted and excites the lower-energy transition, while the other laser is blue-shifted and excites the higher-energy transition. The crossover peak is exactly midway between the other two peaks.

The middle peak at $(f_1 + f_2)/2$ is called a "crossover resonance," and it arises from atoms moving at velocities $V = \pm c\Delta f_{res}/2f_0$, where $f_{res} = (f_1 - f_2)$. For the first case of atoms moving toward the probe laser, these atoms see the probe laser *blue*shifted by $f_{res}/2$, so the probe laser then hits the higher-energy f_1 resonance. And these same atoms see the pump laser *red*shifted by $f_{res}/2$, so the pump laser then hits the lower-energy f_2 resonance. Once again, the strong pump laser depletes the ground state for these atoms, so the weak probe laser experiences less absorption. You have to think about it a while before it really makes sense, but then it does.

Beyond this qualitative discussion of SAS, one can make some progress developing a toy model using the rate equations described above. But we will skip past that for now, in part because real atoms exhibit behaviors that are much more complex than the toy model would predict. The laser polarization also factors into the atomic dynamics, and so does optical pumping, the atomic velocity distribution, the finite size of the laser beam, and a host of other factors. SAS peaks sometimes turn into dips, and the whole phenomenon is quite fascinating when you get into the details. Schmidt et al. (1994) do a deep dive into this subject as they attempt to model all the various peaks, dips, and other spectral features that emerge during careful experiments. Once you begin investigating complex physical systems (in this case multi-level atoms interacting with laser fields), whole new research areas soon appear, and numerous Nobel prizes have been awarded for atomic and molecular spectroscopy over the years.

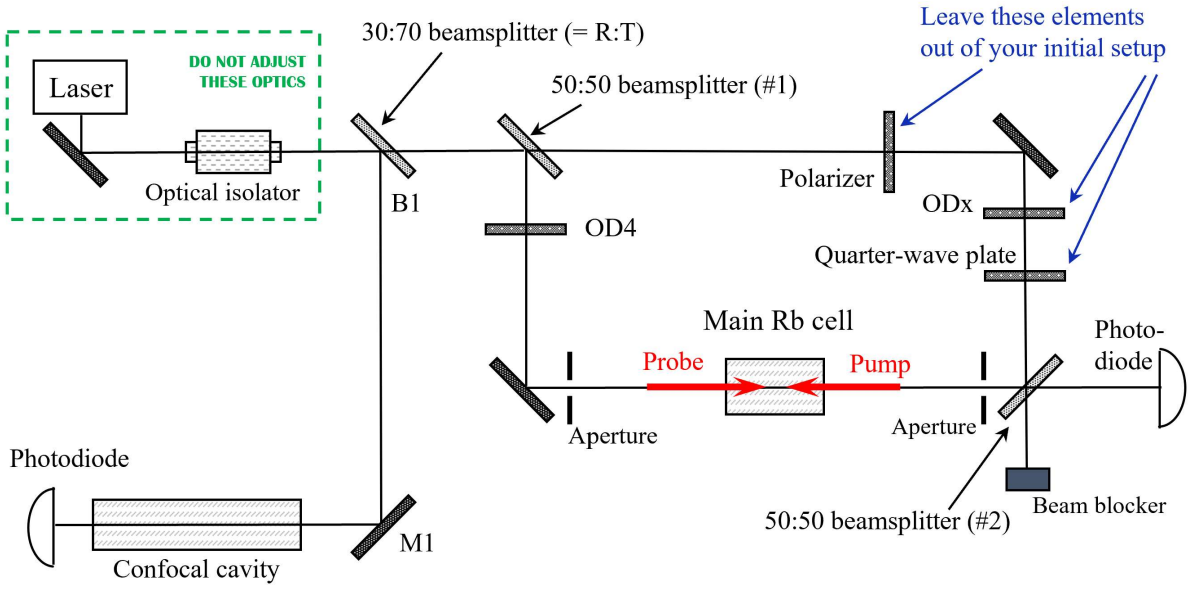

Figure 15. The optical layout for observing saturated-absorption spectroscopy and measuring the rubidium hyperfine splittings. The ***polarizer, ODx, and quarter-wave plate*** *are not necessary for observing the SAS phenomenon, so leave these out of your initial setup. Once you have created some saturated-absorption spectra, you can use these optical elements to adjust the pump-beam polarization and intensity to better observe detailed SAS spectral structures [see Schmidt et al., 1994].*

Lab setup for observing saturated-absorption spectroscopy

Figure 15 shows the overall optical layout for this task, so your first step is to place the various optical elements and get everything aligned and working. Before beginning, however, consider these...

Rules for aligning laser optics

Laser optical setups in research labs can easily involve hundreds of optical components, all of which must be carefully aligned. It is useful, therefore, to learn some of the accepted standard practices for setting such systems up, which researchers have learned through much trial-and-error. If you follow these best practices, your optical alignments will go much faster, *and* they will work much better.

1) As you are handling and aligning the optics, **do not touch any optical surfaces**. This is important, because optical coatings are fragile and easily damaged. The mild acids in fingerprints tend to etch the coatings away, thereby ruining the expensive optical surfaces. Some people insist that gloves be worn while handling optics, but we will not go that far. If you accidentally touch any optical surfaces, or notice any fingerprints, please let us know so we can clean them (which involves standard practices as well).
2) Keep all laser beams at a **constant height above the optical table**. In our lab, we use a fixed height of four inches, and many optical elements are already set to this height. Following this rule eliminates one degree of freedom for all optical elements (height), and that simplifies the alignment process.
3) Use **right-angle beam steering** whenever possible, as illustrated in Figure 15. This keeps all the beams on a rectangular grid, again reducing degrees of freedom and thus simplifying the alignment process. Use the hole pattern on the optical table to define your coordinate system. This rule has become so commonplace that many ultra-high-reflectivity mirrors are designed to be used with right-angle reflections.
4) Keep laser beams **away from the edges** of optical components. Laser beams should always reflect from, or pass through, the central regions of every mirror, beamsplitter, filter, polarizer, lens, etc. Hitting the edge of an optical mount is called *vignetting*, and clearly this should be avoided.
5) Keep mirror mounts near the **middle of their pointing range**. The pointing knobs are for fine adjustments only, so avoid running these fine-pitch screws all the way to the ends of their travel. If you ever see a screw that is all the way in or out, stop what you are doing, set it to the middle of its range, and then proceed with your alignment.
6) Use **apertures** to help guide the alignment process. We will see an example of this below.
7) Finally, **make a plan** for aligning your optics. Usually that means making a detailed sketch of what goes where, and we have already done this for you in Figure 15. Sure, you can just wing it when the number of optical elements is small, like it is here. But when the number gets into the hundreds (as it quickly does in research labs), you need a detailed plan.

The confocal cavity

Even with just a handful of optical elements, it is best to place and align the components a few at a time, making sure each section is working properly before proceeding to the next. (It's like software: writing the whole code and expecting it to work is a bad strategy. Better to write the code bit by bit, testing and debugging along the way.) With this in mind, begin with the 30:70 beamsplitter B1, mirror M1, confocal cavity, and photodiode as shown in in Figure 15, and consider how the various degrees of freedom are managed. Because B1 is far away from the cavity, adjusting the pointing of B1 mainly adjusts the **position** of laser beam when it hits the front face of the cavity. Meanwhile M1 is positioned close to the cavity, so adjusting the pointing of M1 mainly changes the **angle** of the beam as it enters the cavity. Thus these degrees of freedom are largely decoupled… by design. Decoupling the various degrees of freedom is often used to simplify the optical alignment, and this is the kind of thing one plans out in advance when designing an optical layout.

So adjust B1 so the beam hits the center of the first cavity mirror, and then adjust M1 to get the beam to retroreflect from the cavity. In other words, the beam should strike the center of the first cavity mirror and then reflect back on itself. Accomplish this by sending the beam through a white card with a hole in it, and adjust M1 to send the reflected beam back through the hole in the opposite direction. Place the card right in front of M1 to begin, because the M1 pointing may start out way off. Use the camera to view the beam on the card. You can place the card closer to the beamsplitter once you have M1 with nearly the correct pointing.

Now you need to iterate this procedure: adjust B1 to center the beam on the cavity entrance, then adjust M1 to get the retroreflected beam going in the right direction, then go back and tweak the beamsplitter, then back to M1 again to tweak the retroreflection. This iteration is needed because the position and pointing degrees of freedom are not completely decoupled, but only approximately so.

When the beam accurately hits the center of the first cavity mirror and is nicely reflected back into the optical isolator, then the cavity should be adequately aligned. Now place the photodiode close to the back end of the cavity and look at the light transmitted through the cavity as the laser frequency is scanned. Send the 2nd (cavity) PD to ch2 on the 'scope, and set ch2 to Invert, as you did with ch1. Turn up the PD gain to find the cavity transmission signal. Once you see a signal on the oscilloscope, move the photodiode around to maximize it. The end result should look like the lower trace in Figure 16.

If your cavity peaks are not nice and uniform like those shown in Figure 16, this might be because the laser is not running in a single mode as it does its frequency scan. As you may have seen above, tweaking the laser current often helps, and changing the Piezo DC Offset may help as well. The underlying problem is that 10 GHz is a long sweep for a diode laser to make without any mode hops or other odd behavior. Often the best strategy is to do a shorter sweep, because then you can usually adjust the laser parameters to obtain single-mode operation over the reduced sweep. But there is usually a combination of adjustments that will give you a clean 10 GHz sweep if you work at it.

Note that setting up these four optical elements (B1, M1, the cavity, and the photodiode) provides a basic example of how optical alignment should be planned and executed. The layout in Figure 15 was designed to provide enough degrees of freedom to accomplish the needed alignment,

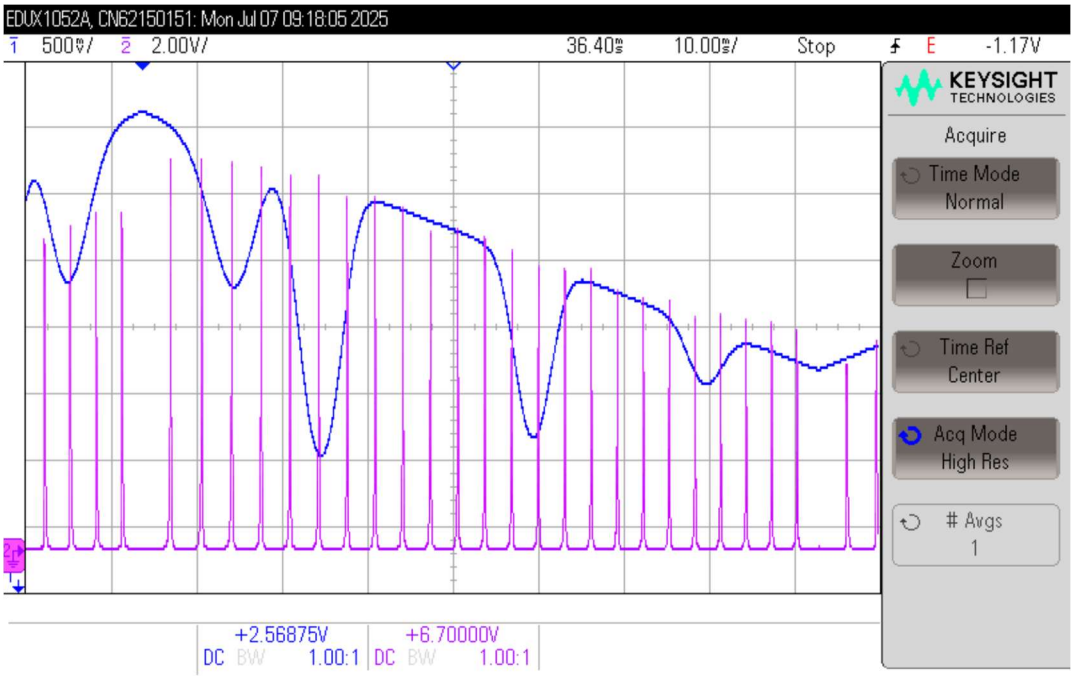

Figure 16. A screenshot showing light from the probe laser beam transmitted through the Rb cell (top trace) along with light from a separate beam transmitted through the confocal cavity. The cavity peaks are separated by 375 MHz. If you see these extra peaks in your scan, you might get rid of them by tweaking the laser current and piezo voltage to reduce mode hops.

with suitable decoupling between positional and angular variables. Having a good layout plan ahead of time speeds things up considerably.

Exercise 12. Tweak the various 'scope parameters to get a nice signal that looks close to what you see in Figure 16 (but with the cavity spectrum only) and add it to your e-notebook. It is not trivial to get all this working well, so take your time and try to work through it systematically.

The cavity peaks you see in Figure 16 happen when the laser frequency coincides with one of the cavity resonances. The physics behind this is described in the Optics Track, which you either did already or might do later. We will not repeat the full discussion here; suffice it to say that these cavity resonance peaks differ in frequency by the free-spectral-range (FSR) $\Delta f_{FSR} = c/4L$, where c is the speed of light and L is the spacing between the cavity mirrors. In this case $L \approx 20$ cm, so $\Delta f_{FSR} \approx 375$ MHz. This spacing supplies a "ruler" you can use to measure the Rb hyperfine splittings.

Another feature you may notice is that the cavity peaks are not perfectly uniformly spaced on the oscilloscope. Use the 'scope cursors to measure the peak spacing on the left and right sides of the display to see this for yourself. What's happening here is that the laser frequency scanning is not simply proportional to the ramp signal you are using in the laser controller. Ideally the scan would be perfectly linear, but there are some nonlinearities in the system. What this means is that while the spacing between adjacent "ticks" on your confocal-cavity ruler always correspond to 375

MHz, the spacing on the ticks on the oscilloscope is not perfectly uniform. It's just something you have to deal with, so be aware.

Your next task is to apply 100 MHz sidebands to the laser and use these to measure the cavity free-spectral range Δf_{FSR}. Turn on the signal generator on top of the laser controller and set it to produce a 100 MHz sine-wave signal at an amplitude of 1Vrms. The cabling should already be set up for you, sending this signal into the laser box, where it modulates the laser injection current. Next turn down the laser scan amplitude (Ramp Generator amplitude) and you should see fewer cavity peaks flanked by small "sideband" peaks, as shown in Figure 17. If your sidebands are small, turn up the amplitude of the signal generator (but do *not* remove the inline attenuator… this protects the laser).

The theory behind these sidebands is discussed in detail in both the Electronics Track and the Optics Track, so we will not repeat it here. If you end up working in just about any field that involves sinusoidal signals (at radio frequencies, optical frequencies, or anything in between) you will likely encounter these kinds of modulation techniques being used.

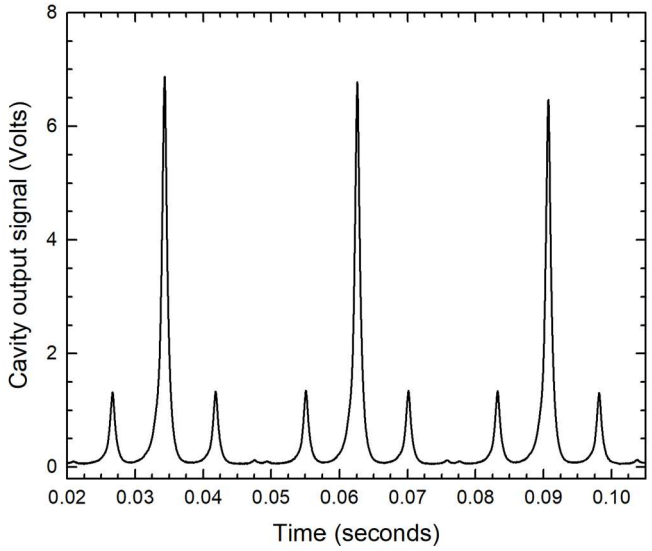

Figure 17. Here a 100-MHz modulation of the laser current adds "sidebands" to the cavity peaks. These new peaks are ±100 MHz away from the central peaks, and this spacing is very precise, because the function generator frequency accuracy is better than 1ppm. You can use this known peak spacing to measure the spacing between cavity peaks.

Exercise 13. This is a good time to align the laser beam going into the cavity a bit more, to produce peaks that are sharp and symmetrical. When you have a nice-looking trace, save the data to a .csv file. Plot the data and measure the positions of the various peaks. Use the known sideband spacing to measure the free-spectral range of the cavity, which is expected to be around 375 MHz. Document your work and report a measured FSR with an uncertainty estimate.

The saturated absorption setup

The next step is to set up the probe laser beam in Figure 15, and you can begin by turning the Rb cell temperature up to 35C, so this has time to stabilize. Next examine the two apertures already placed on the table. Although these are nothing more than empty holes, they are extremely useful for aligning the probe/pump beams. Because the utility of these apertures might not be immediately

obvious if you are new to laser optics, they have already been placed on the table in the correct positions. (Our assumption is that you have not done much laser optics before, and it is often best to see it done correctly before you do it entirely on your own.) Place a mirror right next to the first aperture as shown in Figure 15, and then place the first 50:50 beamsplitter so everything is aligned along the xy grid. Choose the beamsplitter labeled "1", which should be available in the lab.

Note that the laser position at the first aperture is mainly determined by the beamsplitter, while the angle of the beam is mainly determined by the mirror. And, as you saw with the cavity alignment, and these degrees of freedom are again nearly decoupled. So you know what to do: iterate the alignment of these optical elements until the probe beam passes through both apertures, which you can easily do without moving either aperture. The apertures are pre-aligned on the rectangular grid, so everything goes together neatly. And you can slide the Rb cavity down along the rail so it sits between the apertures, as shown in Figure 15.

Next set up the second 50:50 beamsplitter, and this time choose the one labeled "2". To see why these two beamsplitters are labeled, set "2" next to "1" as if you were going to swap the two. With "1", the transmitted beam does not suffer from any vignetting. But if you used "2" in its place, the transmitted beam then might get blocked by the beamsplitter mount. The difference is just in where the steel post is attached to the mirror mount, and you can see that "1" and "2" are different in this regard. This is an easy problem to fix (just move the post), but you can avoid this added step if you just use "1" in the first position and "2" in the second position, as shown in Figure 15. This mirror-mount chirality is just a tidbit of experimental knowledge you soon pick up when you work with laser optics.

Next add the photodiode to complete the probe-laser path (still without the pump laser) and align it so the probe beam signal can be seen on the 'scope. If the absorption dips are not visible, remember the power-dependent absorption you studied already; this would be a good time to add the OD4 attenuator in the probe beam.

Exercise 14. Produce a screenshot that looks similar to that in Figure 16, with both traces present now. Plot the data from a .csv file as well, with your plot showing both channels, much like the screenshot. Document your setup progress so far in your e-notebook.

Exercise 15. Use your plot to measure the large Rb *ground-state* hyperfine splittings for both isotopes. Using your .csv file, measure both of the ground-state splittings shown in Figure 3, and estimate your measurement uncertainties. This means measuring the centers of all four absorption dips (again fitting the data like you did in Figure 8) and then using the confocal "ruler" to convert arbitrary oscilloscope time to transition frequencies in MHz. Try to measure the splittings as accurately as you can, to a fraction of a cavity FSR. And think a bit about how you can model the non-uniform spacing of the ticks if appropriate. Do not just report the numbers; document your work.

Saturated-absorption spectra

Your next setup task is to add the pump laser as shown in Figure 15, again using the two apertures to guide the alignment. Before you start working on the optics, turn the Rb cell temperature up to

45C, so you will not have to wait for it to stabilize later. As noted in the figure, leave out the polarizer, ODx attenuator, and quarter-wave plate during the first stage of this setup.

When adjusting one laser beam (either the pump or probe), you will often find it beneficial to block the other to cut down on potential confusion. Here you can see the utility of using apertures as alignment aids. If two laser beams both pass through two separate points in space, then the beams are guaranteed to be colinear and overlapping. Once both beams are in place, you should see some saturated-absorption (SA) peaks on the probe laser signal. Tweak the alignment of the pump laser to maximize the SA signal, using both the pump-beam mirror and beamsplitter. It is often beneficial to double-check all the optical alignments at this point, because things tend to shift every time you add a new piece of glass. Recenter the probe-laser beam on the photodiode as well.

Exercise 16. Once you see some SA peaks, zoom in on each of the four absorption dips and see how well you can bring out the SA peaks. Including crossover resonances, there should be six peaks for each absorption dip, and it is usually possible to see all six if you zoom in and optimize each set separately, as illustrated in Figure 18. Some absorption features are easier than others, so start with the 87b resonance, because the absorption is strong and the upper-level splittings are large.

Obtaining clean spectra takes some time and perseverance, and this is the time to try different ODx attenuators, different quarter-waveplate rotations, and different polarizer rotations for optimal results. A larger ODx usually makes the peaks sharper, although shorter also. Often an OD2 attenuator in the pump beam is a good choice for the Rb-85 peaks, as they are quite closely spaced. Rotating the quarter-wave plate can turn dips into peaks, and the polarizer can be helpful in this

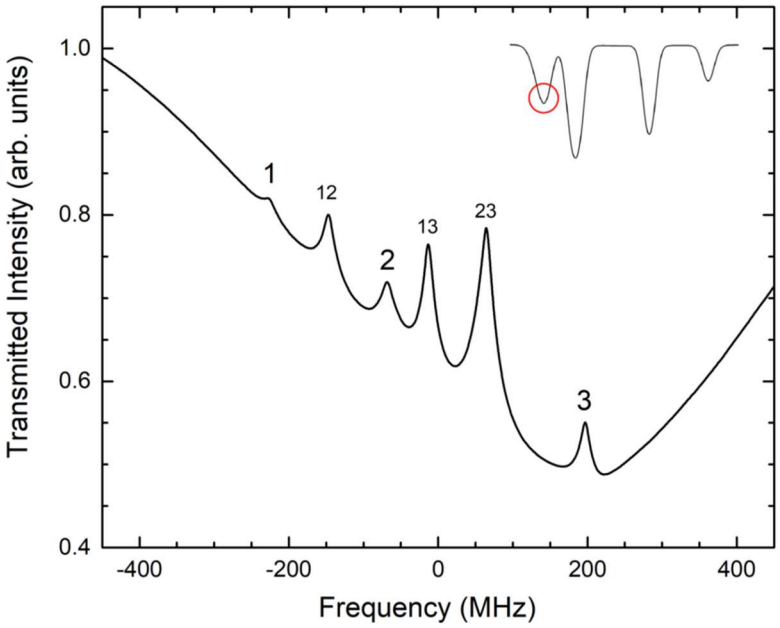

Figure 18. An example spectrum showing the Rb 87b absorption feature, including saturated-absorption peaks. The confocal-cavity peaks were used to convert oscilloscope time to changes in the laser frequency Δf during the sweep. Numbers show the primary transitions and their associated crossover resonances. Note that each crossover peak is exactly midway between its primary peaks, as theory predicts.

regard also. It seems to work well keeping the OD4 in the probe beam throughout, as the probe beam always wants to be weak.

For each of the four absorption dips, record screen shots showing as many SA peaks as you could produce. If you cannot bring out all six peaks for every absorption dip, that's fine. Some days it can be difficult to see the smaller features. Be sure to label your plots with the resonance feature (e.g. 87b). (In the lab, be sure to keep careful track of which data files correspond to which absorption dips, to avoid any possible confusion.)

Exercise 17. Produce an 87b spectrum that shows all six hyperfine peaks, save it to a .csv file, and make a plot for your e-notebook. Identify and label the six peaks and crossover resonances, as described in Figure 14 and the associated discussion. Include the cavity peaks in your plot. Using the cavity peaks, convert your data to absorption as a function of Δf, the change in laser frequency in MHz.

From the plot, measure the Rb-87 upper-state hyperfine splittings shown in Figure 3, and again estimate your measurement uncertainties. Again, do not just report your numbers; document your work. [Optional. Use your other data to also measure the Rb-85 upper-state hyperfine splittings.]

Congratulations! You have finished our rubidium spectroscopy lab, so now you have some experience working with tunable lasers, atomic spectroscopy, and general optics and electronics hardware. Laser technology is a large and booming business, and atomic spectroscopy is a foundational element in our understanding of quantum mechanics. Our sincere hope is that you will find these new experiences useful in your future endeavors.

References

[1] Schmidt, O., Knaak, K.M., Wynands, R., and Meschede, D. 1994, "Cesium Saturation Spectroscopy Revisited: How to Reverse Peaks and Observe Narrow Resonances," Appl. Phys. B, 59, 167.