

Laser Cooling and Trapping

1 INTRODUCTION

Laser cooling and trapping of neutral atoms is a rapidly expanding area of physics research that has seen dramatic new developments over the last decade. These include the ability to cool atoms down to unprecedented kinetic temperatures (as low as one micro Kelvin) and to hold samples of a gas isolated in the middle of a vacuum system for many seconds. This unique new level of control of atomic motion is allowing researchers to probe the behavior of atoms in a whole new regime of matter where deBroglie wavelengths are much larger than the Bohr radius. Undoubtedly one of the distinct appeals of this research is the leisurely and highly visible motion of the laser cooled and trapped atoms.

In this experiment you will operate a laser trap that is equal or superior in performance to what is used in many current research programs. This experiment uses the lasers and saturated absorption spectrometers used in the laser spectroscopy experiment and thus you should have done that experiment before doing this one. A small fraction (~ 10%) of the beams of each of the two lasers goes to their respective saturated absorption spectrometers. This allows for precise detection and control of the laser frequencies, which is essential for cooling and trapping. The remainder of the laser light goes into the trapping cell.

Section 2 of this write-up provides a brief introduction to the relevant physics of the atom trap, section 3 discusses the laser stabilization, section 4 explains the optical layout for sending the laser beams into the cell to create the trap, section 5 explains the trapping cell construction and section 6 discusses the operation of the trap, measurement of the number of trapped atoms, and measurement of the time the atoms remain in the trap.

2 THEORY AND OVERVIEW

We will present a brief description of the relevant physics of the vapor cell magneto-optical trap. For more information, a relatively non-technical discussion is given in Ref. 2, while more detailed discussions of the magneto-optical trap and the vapor cell trap can be found in Ref. 3 and Ref. 4, respectively.

2.1 LASER COOLING

The primary force used in laser cooling and trapping is the recoil when momentum is transferred from photons scattering off an atom. This radiation-pressure force is analogous to that applied to a bowling ball when it is bombarded by a stream of ping pong balls. The momentum kick that the atom receives from each scattered photon is quite small; a typical velocity change is about 1 cm/s. However, by exciting a strong atomic transition, it is possible to scatter more than 10^7 photons per second and




Figure 1. Atomic scattering rate versus laser frequency.

produce large accelerations (10^4 g). The radiation-pressure force is controlled in such a way that it brings the atoms in a sample to a velocity near zero ("cooling"), and holds them at a particular point in space ("trapping").

The cooling is achieved by making the photon scattering rate velocity-dependent using the Doppler effect.[5] The basic principle is illustrated in Figure 1. If an atom is moving in a laser beam, it will see the laser frequency ν_{laser} shifted by an amount $-(v/c)\nu_{\text{laser}}$, where v is the velocity of the atom along the direction of the laser beam. If

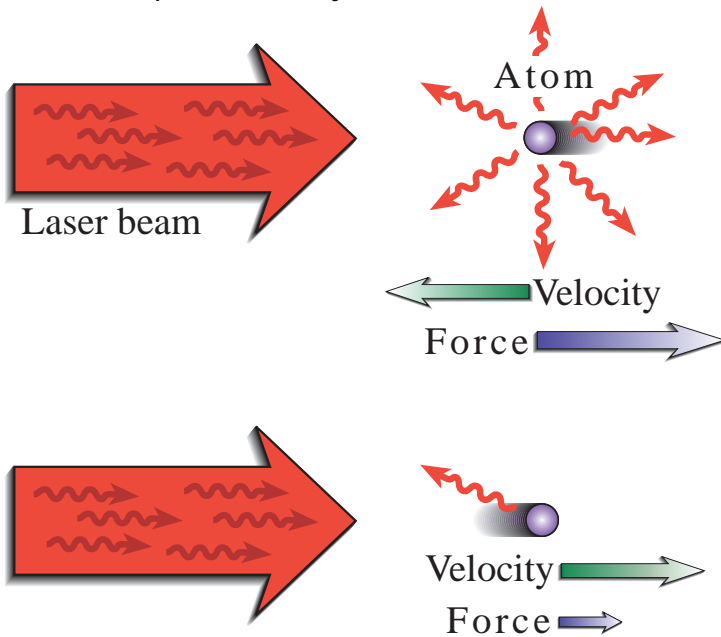


Figure 2. With the laser tuned to below the peak of atomic resonance. Due to the Doppler shift, atoms moving in the direction opposite the laser beam will scatter photons at a higher rate than those moving in the same direction as the beam. This leads to a larger force on the counter-propagating atoms.




Figure 3. One dimensional explanation of the MOT. Laser beams with opposite helicity polarizations impinge on an atom from opposite directions. The lasers excite the $J = 0$ to $J = 1$ transition. The laser beam from the right only excites the $m = -1$ excited state, and the laser from the left only excites the $m = +1$ state. As an atom moves to the right or left, these levels are shifted by the magnetic field thereby affecting the respective photon scattering rates. The net result is a position-dependent force that pushes the atoms into the center.

beam from the left can only excite transitions to the $m = +1$ state, while the beam from the right can only excite transitions to the $m = -1$ state. The magnetic field is zero in the center, increases linearly in the positive x direction, and decreases linearly in the negative x direction. This field perturbs the energy levels so that the $\Delta m = +1$ transition shifts to lower frequency if the atom moves to the left of the origin,

the laser frequency is below the atomic resonance frequency, the atom, as a result of this Doppler shift, will scatter photons at a higher rate if it is moving toward the laser beam (v negative), than if it is moving away. If laser beams impinge on the atom from all six directions, the only remaining force on the atom is the velocity-dependent part, which opposes the motion of the atoms. This provides strong damping of any atomic motion and cools the atomic vapor. This arrangement of laser fields is often known as "optical molasses".[6]

Will scatter photons at a higher rate than those moving in the same direction as the beam. This leads to a larger force on the counter-propagating atoms.

2.2 MAGNETO-OPTICAL TRAP

Although optical molasses will cool atoms, the atoms will still diffuse out of the region if there is no position dependence to the optical force. Position dependence can be introduced in a variety of ways. Here we will only discuss how it is done in the "magneto-optical trap" (MOT), also known as the "Zeeman shift optical trap," or "ZOT." The position-dependent force is created by using appropriately polarized laser beams and by applying an inhomogeneous magnetic field to the trapping region. Through Zeeman shifts of the atomic energy levels, the magnetic field regulates the rate at which an atom in a particular position scatters photons from the various beams and thereby causes the atoms to be pushed to a particular point in space. In addition to holding the atoms in place, this greatly increases the atomic density since many atoms are pushed to the same position. Details of how the trapping works are rather complex for a real atom in three dimensions, so we will illustrate the basic principle using the simplified case shown in Figure 2.

In this simplified case we consider an atom with a $J = 0$ ground state and a $J = 1$ excited state, illuminated by circularly polarized beams of light coming from the left and the right. Because of its polarization, the




Figure 4. Schematic of the MOT. Lasers beams are incident from all six directions and have helicities (circular polarizations) as shown. Two coils with opposite currents produce a magnetic field that is zero in the middle and changes linearly along all three axes.

while the $\Delta m = -1$ transition shifts to higher frequency. If the laser frequency is below all the atomic transition frequencies and the atom is to the left of the origin, many photons are scattered from the σ^+ laser beam, because it is close to resonance. The σ^- laser beam from the right, however, is far from its resonance and scatters few photons. Thus the force from the scattered photons pushes the atom back to the zero of the magnetic field. If the atom moves to the right of the origin, exactly the opposite happens, and again the atom is pushed toward the center where the magnetic field is zero. Although it is somewhat more complicated to extend the analysis to three dimensions, experimentally it is simple, as shown in Figure 4. As in optical molasses, laser beams illuminate the atom from all six directions. Two symmetric magnetic field coils with oppositely directed currents create a magnetic field that is zero in the center and changes linearly along the x, y, and z axes. If the circular polarizations of the lasers are set correctly, a linear restoring force is produced in each direction. Damping in the trap is provided by the cooling forces discussed in section 2.1. It is best to characterize the trap "depth" in terms of the maximum velocity that an atom can have and still be contained in the trap. This maximum velocity v_{max} is typically a few times $\Gamma\lambda$ ($\Gamma\lambda$

is the velocity at which the Doppler shift equals the natural linewidth Γ of the trapping transition, where λ is the wavelength of the laser light).

A much more complicated three-dimensional calculation using the appropriate angular momentum states for a real atom will give results which are qualitatively very similar to those provided by the above analysis if: (1) the atom is excited on a transition where the upper state total angular momentum is larger than that of the lower state ($F \rightarrow F' = F + 1$) and (2) $v \geq (\hbar\Gamma / 2m)^{1/2}$ where m is the mass of the atom. This velocity is often known as the "Doppler limit" velocity. If the atoms are moving more slowly than this, "sub-Doppler" cooling and trapping processes become important, and the simple analysis can no longer be used.[7] We will not discuss these processes here, but their primary effect is to increase the cooling and trapping forces for very slow atoms in the case of $F \rightarrow F + 1$ transitions.




Figure 5. ^{87}Rb energy level diagram showing the trapping and hyperfine pumping transitions. The atoms are observed by detecting the 780 nm fluorescence as they decay back to the ground state.

We will now consider the specific case of rubidium (Figure 5). Essentially all the trapping and cooling is done by one laser which is tuned slightly (1-3 natural linewidths) to the low frequency side of the $5S_{1/2} F = 2 \rightarrow 2P_{3/2} F' = 3$ transition of ^{87}Rb . (For simplicity we will only discuss trapping of this isotope. The other stable isotope, ^{85}Rb , can be trapped equally well using its $F = 3 \rightarrow F' = 4$ transition.) Unfortunately, about one excitation out of 1000 will cause the atom to decay to the $F=1$ state instead of the $F=2$ state. This takes the atom out of resonance with the trapping laser. Another laser (called the "hyperfine pumping laser") is used to excite the atom from the $5S F = 1$ to the $5P F' = 1$ or 2 state, from which it can decay back

to the $5S F = 2$ state where it will again be excited by the trapping laser.

In a vapor cell trap, the MOT is established in a low pressure cell containing a small amount of rubidium vapor.[4] The rubidium atoms in the low energy tail ($v < v_{\max} \approx 20 \text{ m/s}$) of the Maxwell-Boltzmann distribution are captured in the laser trap. If the trap is fumed on at $t=0$, the number N of atoms in the trap will increase with the same functional form as that of a capacitor charging,

$$N(t) = N_0(1 - e^{-t/\tau}), \quad (1)$$

where τ is the time constant for the trap to fill to its steady state value N_0 and is also the average time an atom will remain in the trap before it is knocked out by a collision. This time is just the inverse of the loss rate from the trap due to collisions. Under certain conditions, collisions between the trapped atoms can be important, but for conditions that are usually encountered, the loss rate will be dominated by collisions with the room temperature background gas. These "hot" background atoms and molecules (rubidium and contaminants) have more than enough energy to knock atoms out of the trap. The time constant τ can be expressed in terms of the cross sections σ , densities n , and velocities of rubidium and non-rubidium components as:

$$\frac{1}{\tau} = n_{\text{Rb}}\sigma_{\text{Rb}}v_{\text{Rb}} + n_{\text{non}}\sigma_{\text{non}}v_{\text{non}}. \quad (2)$$

The steady-state number of trapped atoms is that value for which the capture and loss rates of the trap are equal. The capture rate is simply given by the number of atoms which enter the trap volume (as defined by the overlap of the laser beams) with speeds less than v_{\max} . It is straightforward to show that this is proportional to the rubidium density, v_{\max}^4 , and the surface area A of the trap. When the background vapor is predominantly rubidium, the loss and capture rates are both proportional to rubidium pressure. In this case N_0 is simply:

$$N_0 = 0.1 \frac{A}{\sigma_{\text{Rb}}} \left(\frac{v_{\max}}{v_{\text{avg}}} \right)^4, \quad (3)$$

where $v_{\text{avg}} = (2kT/m)^{1/2}$, the average velocity of the rubidium atoms in the vapor. If the loss rate due to collisions with non-rubidium background gas is significant, Eq. (3) must be multiplied by the factor $n_{\text{Rb}}\sigma_{\text{Rb}}v_{\text{avg}} / (n_{\text{Rb}}\sigma_{\text{Rb}}v_{\text{avg}} + n_{\text{non}}\sigma_{\text{non}}v_{\text{non}})$. The densities are proportional to the respective partial pressures. Finally, if the loss rate is dominated by collisions with non-rubidium background gas, the number of atoms in the trap will be proportional to the rubidium pressure divided by the non-rubidium pressure, but τ will be independent of the rubidium pressure.

As a final note on the theory of trapping and cooling, we emphasize certain qualitative features that are not initially obvious. This trap is a highly over-damped system; hence damping effects are more important for determining trap performance than is the trapping force. If this is kept in mind it is much easier to gain an intuitive understanding of the trap behavior. Because it is highly over-damped, the critical quantity v_{\max} is determined almost entirely by the Doppler slowing which provides the damping. Also, the cross sections for collisional loss are only very weakly dependent on the depth of the trap, and therefore the trap lifetime is usually quite insensitive to everything except background pressure. As a result of these two features, the number of atoms in the trap is very sensitive to laser beam diameter, power, and frequency, all of which affect the Doppler cooling and hence v_{\max} . However, the number of trapped atoms is insensitive to factors which primarily affect the trapping force but not the damping, such as the magnetic field (stray or applied) and the alignment and polarizations of the laser beams. For

example, changing the alignment of the laser beams will dramatically affect the shape of the cloud of trapped atoms since it changes the shape of the trapping potential. However, these very differently shaped clouds will still have similar numbers of atoms until the alignment is changed enough to affect the volume of the laser beam overlap. When this happens, the damping in three dimensions is changed and the number of trapped atoms will change dramatically. Of course, if the trapping potential is changed enough that there is no potential minimum (for example, the zero of the magnetic field is no longer within the region of overlap of the laser beams), there will be no trapped atoms. However, as long as the damping force remains the same, almost any potential minimum will have about the same number of atoms and trap lifetime.

2.3 OVERVIEW OF THE TRAPPING APPARATUS.

Figure 6 shows a general schematic of the trapping apparatus. It consists of two Vortex™ diode lasers, two saturated absorption spectrometers, a trapping cell, and a variety of optics. The optical elements are lenses for expanding the laser beams, mirrors and beamsplitters for splitting and steering the

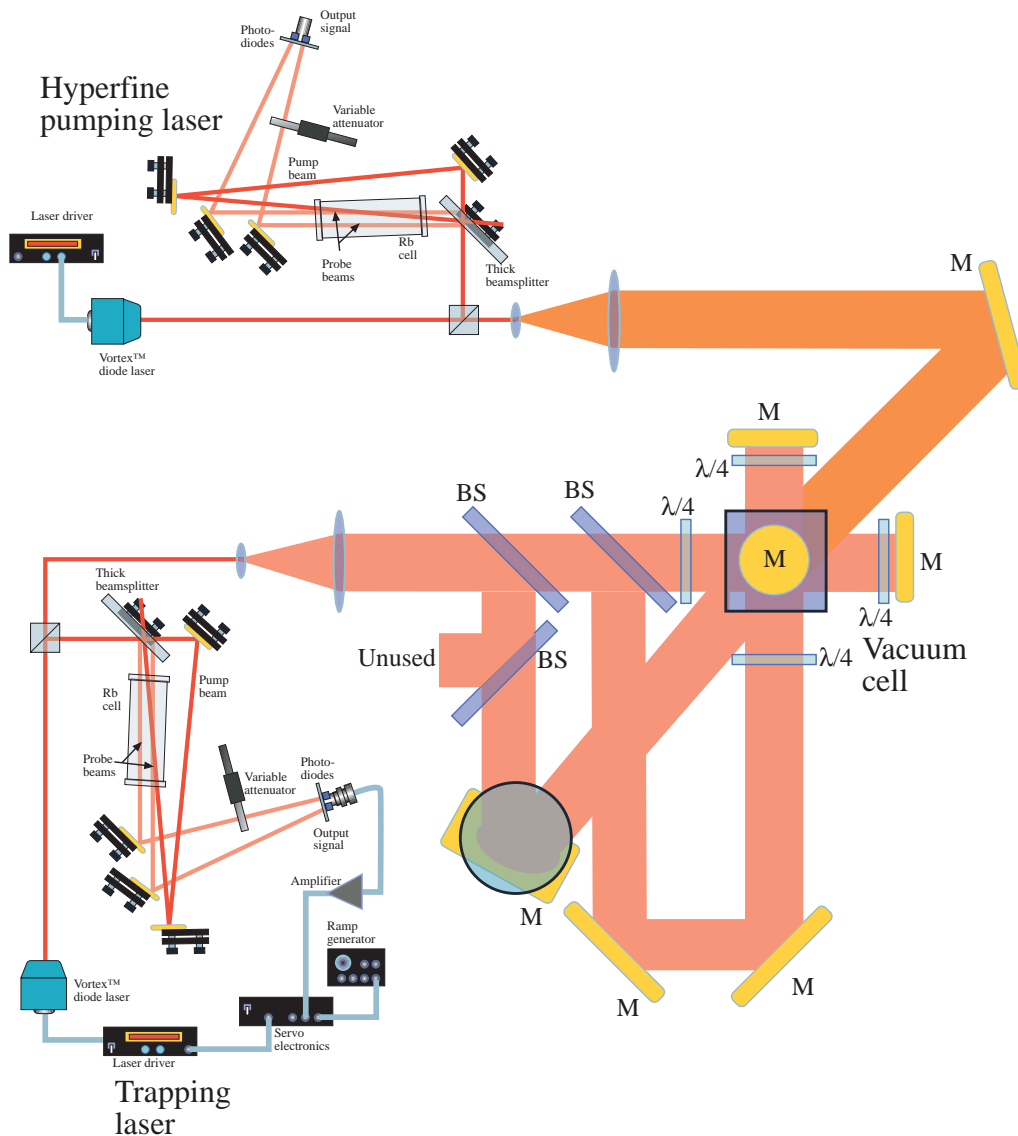


Figure 6. Overall optical layout for laser trap experiment including both saturated absorption spectrometers.




Figure 7. Detail of laser beams that are sent through the trapping cell. To simplify the figure, the $\lambda/4$ waveplates are not shown. Two beam paths are in the - horizontal plane and a third is brought down and is then reflected up vertically through the bottom of the cell. The retro-reflected beams are tilted slightly to avoid feedback to the diode laser.

beams, and waveplates for controlling their polarizations. To monitor the laser frequency, a small fraction of the output of each laser is split off and sent to a saturated absorption spectrometer. An electronic error signal from the trapping laser's saturated absorption spectrometer is fed back to the laser to actively stabilize its frequency. The trapping cell is a small vacuum chamber with an ion vacuum pump, a rubidium source, and windows for transmitting the laser light. In the following sections we will discuss the various components of the apparatus and the operation of the trap.

The key to the cooling are the three mutually perpendicular laser beams that cross within the trapping cell, from which another three are derived by retro-reflecting them (see the detail of Figure 7). Figure 6 shows the three laser beams derived using three beamsplitters. The arrangement may be somewhat different in your setup.

3 LASER STABILIZATION

As mentioned above, two lasers are needed for the trap. A few milliwatts of laser power are plenty for hyperfine pumping ($F = 1 \rightarrow F' = 1, 2$), but the number of trapped atoms is nearly proportional to the

amount of power in the trapping laser ($F = 2 \rightarrow F' = 3$). Setting up and using the trap is much easier with at least 5 mW of trapping laser power, although it can work with less. The trapping laser must have an absolute frequency stability of a few megahertz. This requires you to actively eliminate fluctuations in the laser frequency, which are usually due to changes in the length of the laser cavity caused by mechanical vibrations or temperature drifts. This is accomplished by using the saturated absorption signal to detect the laser frequency and a feedback loop to hold the length of the laser cavity constant.

The laser servo control system is far less complicated as first appears. As intimidating as the servo box inputs, outputs, switches and dials seem, the principles are relatively straightforward. The Vortex™ laser frequency is controlled by a piezo voltage. As its frequency is tuned and scanned near a hyperfine level of interest the output from saturated spectroscopy setup (that is, from the photodiode preamplifier)

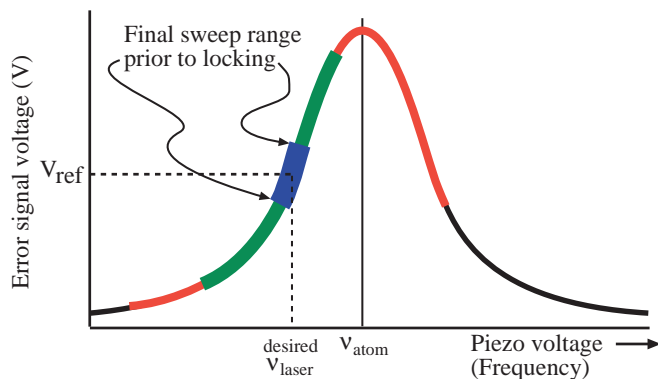


Figure 8. Photodiode preamplifier output voltage (error signal voltage) as a function of Vortex™ piezo (PZT) voltage near the hyperfine line of interest. The servo system is meant to stabilize the laser to the desired frequency.

looks more or less like Figure 8 (or perhaps an upside-down and/or backwards version of it). The purpose of the laser servo controller is to maintain the laser frequency to the “red” side of atomic resonance, about halfway down the curve shown in the figure. At that point the photodiode preamplifier produces an output that has been labeled as V_{ref} . The servo controller box stabilizes the laser frequency by adjusting the piezo voltage such that the photodiode amplifier output remains at V_{ref} , or nearly so. Briefly it does so by decreasing the piezo voltage when the photodiode output falls above V_{ref} and increases the piezo voltage when the photodiode rises above V_{ref} .

The servo needs some help from you to get started, however. The insides of the servo

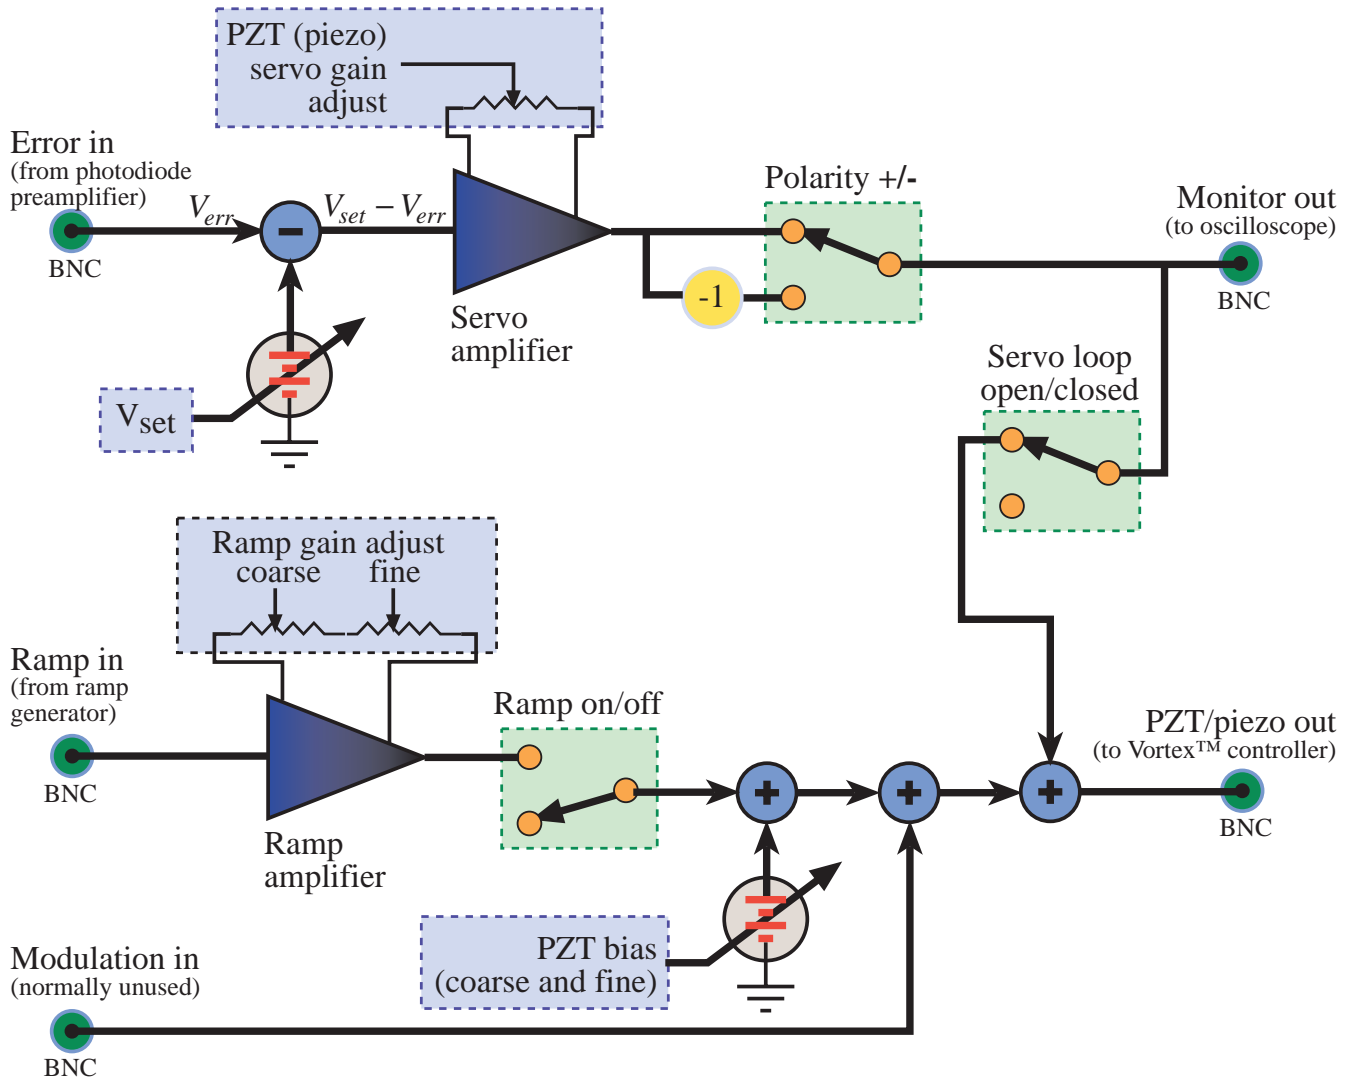


Figure 9. Laser frequency stabilization servo control box. Shaded and dashed boxes indicate the knobs and switches on the controller unit.

controller box are schematically illustrated in Figure 9. The shaded, dashed-outline boxes label the various switches and knobs that are on the controller.

3.1 SWEEP AND SETPOINT ADJUSTMENTS

The first order of business is to produce a triangular wave input using the function generator. Start with a low amplitude and turn all biases and offsets to zero. A few tens of Hertz, up to about 100 Hz is a good frequency range to work with. You will connect the function generator output to the BNC on the servo controller box (labeled “Ramp In”) but before you do: 1. Switch the servo loop to “open” (switch down) and the “Ramp” switch to “on” (switch up). 2. Turn the coarse ramp gain adjust knob all the way down (counter-clockwise). 3. Connect the piezo output to an oscilloscope and *not* to the Vortex™ controller. Now connect the function generator to to the ramp input connector on the servo controller unit. Turn the ramp coarse gain all the way up and then adjust the output of the function generator so that the maximum piezo voltage allowed by the Vortex™ controller is not exceeded (± 4.5 V). From now on you should not have to adjust anything on the function generator except perhaps its frequency. Set the piezo voltage on the Vortex™ controller to midway (50 V). Turn down the ramp gain on the servo controller box and connect the piezo output to the Vortex™ piezo input.

The next item of business is to zoom in on the rubidium hyperfine atomic transition of interest (e.g., the $5S F = 2 \rightarrow 5P F' = 3$ transition of ^{87}Rb .) To do so you will use the ramp gain adjust and PZT bias controls. Connect the photodiode preamplifier output to the “Error In” BNC of the servo controller. Also observe the photodiode preamplifier output with the scope. Set the ramp gain to a relatively high value until you see the various rubidium lines as described in the Doppler Free Saturated Absorption Spectroscopy experiment. Center the line of interest on the oscilloscope. You can do so by adjusting the PZT bias knob (coarse, at first) on the servo controller—you can also do it by setting the piezo voltage on the Vortex™ controller directly. Actually, the latter is a bit safer, as the input specifications will not inadvertently be exceeded. To zero-in on the line requires an iterative process. The idea is to reduce the sweep voltage of the piezo by decreasing the ramp gain, then re-centering the scan range to the desired frequency with the piezo offset or the Vortex™ controller piezo setting, then reducing the ramp gain again, and so on. Using the second channel of the oscilloscope used to monitor the photodiode preamplifier output, connect the “Monitor output” of the servo controller box. Be sure to note the position of the scope trace zero for this channel. At the same time, notice the value V_{ref} of the photodiode preamplifier output when it is halfway up the transition. The next step is to adjust the “Setpoint” value of the servo controller. It determines value of the voltage V_{set} shown in Figure 9. Notice what the circuit does—the input to the servo amplifier is the difference, $V_{err} - V_{set}$, where V_{err} is the photodiode preamplifier output voltage. This difference is amplified and then (will later be) fed back to the piezo voltage input of the Vortex™ controller. You should adjust the Setpoint knob so that the monitor output passes through zero exactly at the point where the laser is halfway up the transition curve. You can overlap the monitor and photodiode preamplifier traces on the oscilloscope, then adjust the setpoint knob so former traces crosses zero at the appropriate place.

3.2 ACQUIRING LASER LOCK

You are almost ready to “close the loop” and let the frequency servo do its job. First turn the “Servo gain” knob to midway between its minimum and maximum. Second you must further zero-in on frequency region of interest by adjusting the piezo sweep using the same iterative process as you did above. Continue reducing the ramp gain in steps until the laser frequency sweep covers just a small region around the laser frequency of interest. Figure 8’s illustration is meant to be suggestive of this process—the increasing thickness of lines is meant to show the decreasing frequency scans of the piezo voltage.

Now you can close the servo loop by putting the loop switch in the up position. If the laser is properly locked the photodiode preamplifier output should be moving up and down a bit but stay near the value V_{ref} . However, it is not entirely likely that the servo was able to “catch” lock on the first try. Both the servo gain and the polarity may need to be adjusted. The polarity switch adjust the sign of the feedback—you can see from Figure 9 that it merely enters or removes a -1 in the feedback loop. In principle one can track through the electronics and know which polarity, $+$ or $-$, is appropriate for the feedback signal. In practice, however, it is usually easier just to flip the polarity switch to discover which position is correct. If you choose the wrong sign, then the laser becomes stabilized to the wrong side of the transition and laser cooling will not work!

If the servo is not working, it is also possible that the servo gain is either too low or too high. Too low and there is simply not enough piezo voltage to return the laser to the desired frequency when it drifts away. If the servo gain is too high the servo can go “crazy” by oscillating, or the servo can suddenly jump and lock to some other (undesired) transition. It is usually effective to turn the servo gain low and then turn it up until the servo goes, or nearly goes, into oscillation. Then set the gain just a bit lower and redo the locking procedure.

Detailed discussions of the saturated absorption spectrometer, the feedback loop circuit, and the procedure for locking the laser frequency are given in Ref. 1, along with examples of spectra obtained. You should already be familiar with all but the locking of the laser frequency. After a little practice one can lock the laser frequency to the proper value within a few seconds.

Here we just mention a few potential problems and solutions. Check that the baseline on the saturated absorption spectrum well off resonance is not fluctuating by more than several percent of the $F = 2 \rightarrow F' = 3$ peak height. If there are larger fluctuations, they are likely caused by light feeding back into the laser or by the probe beams vibrating across the surface of the photodiodes. Changing the alignment and making sure that all the optical components are rigidly mounted will normally solve these problems. Under quiet conditions and reasonably constant room temperature, the laser should stay locked for many minutes and sometimes hours at a time. Bumping the table or the laser or making a loud noise will likely knock the laser out of lock.

3.3 STABILITY REQUIREMENT

[Optional note: Even under the best conditions, the system we have described is likely to have residual frequency fluctuations of around 1 MHz. The trap will work fine with this level of stability but these fluctuations will cause some noise in the fluorescence from the trapped atoms. For most experiments this is not serious, but it can limit some measurements. If one has a little knowledge of servo systems, it is quite straightforward to construct a second feedback loop that adjusts the laser current. This current feedback loop should have a roll-on filter and high gain. (Because the laser current loop can be much faster than the PZT loop, it can have a much higher gain.) This results in the current loop feedback dominating for frequencies above a few hertz. However, the PZT loop gain is largest for very low frequencies, and thus handles the DC drifts. The combination of the two servo loops will make the laser frequency much more stable and will also make the lock extremely robust. We have had lasers with combined PZT and current servo loops remain frequency-locked for days at a time and resist all but the most violent jarring of the laser. A circuit diagram for the combined PZT and current servo loop can be obtained from C. Wieman. This extra servo loop is a complication, however, which is not necessary for your experiments.]

The requirements for the frequency stability of the hyperfine pumping laser are much less stringent than those for the trapping laser. For many situations it is adequate to simply set the frequency near the peak of the $5S F = 1 \rightarrow 5P F' = 2$ transition by hand. Over time it will drift off, but if the room temperature does not vary too much, it will only be necessary to bring the laser back on to the peak by slightly adjusting the PZT offset control on the ramp box every 5 to 10 minutes. If better control is desired, the frequency of this laser can be locked to the peak of the $F = 1 \rightarrow F' = 1, 2, \text{ and } 3$ Doppler-broadened (and hence unresolved) absorption line by modulating the laser frequency and using phase-sensitive detection with a lock-in amplifier. The output of the lock-in amplifier is then fed back to the PZT to keep the laser frequency on the peak. There is a lock-in amplifier in the lab that can be used if desired. This provides a relatively crude frequency lock that is quite adequate for the hyperfine pumping laser.

4 OPTICAL SYSTEM

The beams from the lasers must now be sent into the cell to form the trap. The basic requirement is to send light beams from the trapping laser into the cell in such a way that the radiation-pressure force has a component along all six directions. To motivate the discussion of the optical design for this lab, we first mention the design used in most of the traps in research programs. In these research traps, the light from the trapping laser first passes through an optical isolator and then through beam shaping optics that make the elliptical diode laser beam circular and expand it to between 1 and 1.5 cm in diameter. The beam is then split into three equal intensity beams using dielectric beam splitters. These three beams are circularly polarized with quarter-wave plates before they pass through the trapping cell where they intersect at right angles in the center of the cell. After leaving the trapping cell each beam goes through a second quarter-wave plate and is then reflected back on itself with a mirror. This accomplishes the goal of having three orthogonal pairs of nearly counter-propagating beams, with the reflected beams having circular polarization opposite to the original beams.

4.1 GEOMETRY

For this lab you use a modified version of this setup, as shown in Fig. 5. A number of the expensive optical components (optical isolator, dielectric beam splitters, and large aperture high quality quarter-wave plates) have been eliminated from the research design. The light from the trapping diode laser is sent into a simple two-lens telescope that expands it. This is then split into three beams. The operation of the trap is insensitive to the relative amounts of power in each of the beams. [You should check how much the number of trapped atoms change with a factor-of-two difference in relative powers.] However, the beam size as set by the telescope is of some importance. The number of trapped atoms increases quite rapidly as the beam size increases and, as discussed below, the larger the beam, the less critical the alignment. However, if the beams are too large they will not fit on the mirrors and it becomes harder to see the beams due to their reduced intensity. Beams of about 1.5 cm diameter work well, but avoid using beams much smaller than 1.5 cm in diameter because of the decrease in the number of trapped atoms and the increased alignment sensitivity. The lenses in the lab will make a simple telescope that expands the beam from the diode laser to this size.

Two of the three beams remain in a horizontal plane and are sent into the cell as shown in Figure 7. The third is angled down and reflects up from the bottom of the cell. To simplify the adjustment of the polarizations, the light should be kept linearly polarized until it reaches the quarter-wave plates. This will be the case as long as all the beams have their axes of polarization either parallel (p) or perpendicular (s) to the plane of incidence of each mirror. This is easy to achieve for the beams in the horizontal plane, but difficult for the beam that comes up through the bottom of the cell. However, with minimal effort to be close to this condition, the polarization will remain sufficiently well linearly polarized. If you think this might be a problem you can check the ellipticity of the polarization by using a photodiode and a rotatable linear polarizer. An eccentricity of 10 or greater on the polarization ellipse is adequate.

After they pass through the cell, the beams are reflected approximately (but not exactly!) back on themselves. The reason for having an optical isolator in the research design is that even a small amount of laser light reflected back into the laser will dramatically shift the laser frequency and cause it to jump out of lock. In the absence of an optical isolator, this will always happen if the laser beams are reflected nearly back on themselves. Feedback can be avoided by insuring that the reflected beams are steered away from the incident beams so that they are spatially offset by many (5-10) beam diameters when they arrive back at the position of the laser. Fortunately, for operation of the trap the return beam need only overlap most of the incident beam in the cell, but its exact direction is unimportant. Thus by making the beams large and placing the retro mirrors close to the trap (within 10 cm for example) it is possible to have the forward and backward going beams almost entirely overlap even when the angle between them is substantially different from 180° . This design eliminates the need for the very expensive (\$2500) optical isolator and, as an added benefit, makes the operation of the trap very insensitive to the alignment of the return beams.

It is easy to tell if feedback from the return beams is perturbing the laser by watching the signal from the saturated absorption spectrometer on the oscilloscope. If the amplitude of the fluctuations is affected by the alignment of the reflected beam or is reduced when the beam to the trapping cell is blocked, unwanted optical feedback is occurring.

4.2 POLARIZATION

The next task is to set the polarizations of the three incident beams. The orientations of the respective circular polarizations are determined by the orientation of the magnetic field gradient coils. The two transverse beams that propagate through the cell perpendicular to the coil axis should have the same circular polarization, while the beam that propagates along the axis of the coils should have the opposite circular polarization. Although in principle it is possible to initially determine and set all three polarizations correctly with respect to the magnetic field gradient, in practice it is much simpler to set the three polarizations relative to each other and then try both directions of current through the magnetic field coils to determine which sign of magnetic field gradient makes the trap work. To set the relative

polarization of the three beams, first identify the same (fast or slow) axis of the three quarter-wave ($\lambda/4$) plates. For the two beams that are to have the same polarization, this axis is set at an angle of 45° clockwise with respect to the linear polarization axis when looking along the laser beam. For the axial beam, the axis is oriented at 45° counterclockwise with respect to the linear polarization. This orientation needs only to be set to within about $\pm 10^\circ$. The orientation of the $\lambda/4$ plates through which the beams pass after they have gone through the trapping cell ("retro $\lambda/4$ plates") is arbitrary. No matter what the plates' orientations are, after the beams have passed through them twice, the light's circular polarization will be reversed.

An optional experiment is to see what happens when you replace one or more of the three retro $\lambda/4$ plates with retro-reflecting right angle mirrors. [1] Although this combination of mirrors does not provide ideal $\lambda/2$ retardance, it is fairly close. An added benefit of this approach is that two reflections off a mirror usually result in much less light loss than one reflection and two passes through a $\lambda/4$ plate.

4.3 HYPERFINE PUMPING LASER OPTICS

Minimal optics is needed for the hyperfine pumping laser. You need only send it through the two lens beam expanding telescope in the lab to make a large roughly collimated beam (typically with a 2-3 cm diameter) and send it into the cell from a direction which will minimize the scattered light into the detectors that observe the trapped atoms. The trap is insensitive to nearly everything about the hyperfine pumping light, including its polarization. Why should you expect this to be the case based on how the trap works?

5 TRAPPING CELL CONSTRUCTION

The primary concern in the construction of the trapping cell is that ultrahigh vacuum (UHV) is required. Although trapped atoms can be observed at pressures of 10^{-5} Pa ($\sim 10^{-7}$ Torr), trap lifetimes long enough for most experiments of interest require pressures in the 10^{-6} to 10^{-7} Pa range. There are three main elements in the trapping cell: (1) a pump to remove unwanted background gas - mostly water, hydrogen, and helium (helium can diffuse through glass), (2) a controllable source of rubidium atoms, and (3) windows to transmit the laser light and allow observation of the trapped atoms. The major components are depicted in Figure 10.

5.1 VACUUM PUMP

The cell is attached to an ion pump. In this pump the atoms are ionized in a high voltage discharge and then embedded in the electrodes, thereby removing them from the system. The 5000 V needed to make the pump operate is provided by a power supply that is attached to the back of the pump by a wire. Although this wire has a grounded outer shield, and therefore in principle quite safe, it is wise to always avoid touching it or the pump near the high voltage connector. The ion pump has the minor drawback that it requires a large magnetic field, which is provided by a permanent magnet. The fringing fields from this magnet can extend into the trap region and will affect the trap to some extent. Although the trap will usually work without it, we put a layer of 0.75 mm (0.03 in) magnetically permeable steel sheet around the pump to shield the trap from this field. For this same reason it is advisable to avoid having magnetic bases very near the trap.

5.2 RUBIDIUM SOURCES

We will now discuss how to produce the correct pressure of rubidium vapor in the cell. The vapor pressure of a room temperature sample of rubidium is about 5×10^{-5} Pa. This is much higher than the 10^{-6} to 10^{-7} Pa ($\sim 10^{-8}$ to 10^{-9} Torr) of rubidium vapor pressure that is optimum for trapping. At higher pressures the trap will still work, but the atoms remain in the trap a very short time, and it is often




Figure 10. Drawing of trapping cell. The tubes on the cross have been elongated in the drawing for ease of display. The Rb dispenser is inside a stainless steel vacuum tube; we show a cut-away view in this drawing.

difficult to see them because of the bright fluorescence from the un-trapped background atoms. Also, the absorption of the trapping beams when passing through the cell will be significant. Thus you should maintain the rubidium at well below its room temperature vapor pressure. Because it is necessary to continuously pump on the system to avoid the buildup of hydrogen and helium vapor, it is necessary to have a constant source of rubidium to maintain the correct pressure. Through chemical reactions and physisorption, the walls of the cell usually remove far more rubidium than the ion pump does and the rate of pumping by the walls depends on how well they are coated with rubidium.

In this experiment the rubidium vapor is produced by a commercial “rubidium dispenser” (locally sometimes also referred to as a “rubidium getter”. The dispenser is several milligrams of a rubidium compound that is contained in a small (1.0 x 0.2 x 0.2 cm) stainless steel oven. Two of these ovens are spot-welded onto two pins of a vacuum feed-through, one of which is illustrated in Figure 10. When current (3-5 A) is sent through the oven, rubidium vapor is produced. The higher the current, the higher the rubidium pressure in the chamber. With this system it is unnecessary to coat the entire surface with rubidium, and is in fact undesirable since the dispenser is likely to be exhausted before the surface is entirely saturated. You should be able to produce enough rubidium pressure to easily see the background fluorescence with as little as 3.4 A through the dispenser oven. The current through the dispenser is controlled by the power supply under the table. Be careful not to turn the current higher than 6.5 amps under any circumstances! If you do it will burn out the dispenser and ruin the experiment! You can see how much rubidium is in the cell by sending the laser beams into the cell, setting the laser frequency to the trapping transition (the saturated absorption cell should glow brightly) and look for fluorescence from the atoms in the cell. If you can see any glow from inside the cell that goes away when the laser frequency is tuned off the transition, there is more than enough rubidium for trapping.

After the current through the dispenser has been turned on, the rubidium vapor comes to an equilibrium pressure with a time constant of about 5 min. When the current through the dispenser is turned off, the rubidium pressure drops with a time constant of about 4 s if the dispenser has been on for only a short while. With prolonged use of the dispenser, this time constant can increase up to a few minutes as rubidium builds up on the walls, but it decreases to the original value if the dispenser is left off for several days.

The superior rubidium pressure control provided by the dispenser makes it possible to use high pressures to easily observe fluorescence from the background vapor. This allows one both to check whether the laser and cell are operating properly, and to have rapid response in the number of atoms while optimizing trapping parameters. Once these tasks are complete, simply reducing the dispenser current provides low pressures almost immediately. Low pressures are desirable for many experiments because they yield relatively long trap lifetimes (seconds) and little background light from the fluorescence of un-trapped atoms.

5.3 MAGNETIC FIELD GRADIENT

With the cell installed and the laser beams aligned, the final ingredients for trapping are the magnetic field coils. A gradient of up to about 0.20 T/m (20 G/cm; normal trap operation is at 10-15 G/cm) is needed. This is provided by two freestanding coils 1.3 cm in diameter with 25 turns each of 24 gauge magnet wire and a separation of 3.3 cm. The coils are mounted on either side of the cell such that the current travels through the loops in opposite directions and the coil axes are collinear with one of the laser beam axes.

6 OPERATION OF THE TRAP AND MEASUREMENTS ON IT

6.1 OBSERVATION SYSTEM

An inexpensive CCD TV camera and monitor are used to observe the trapping cell. This will show the cloud of trapped atoms as a very bright white glow in the center of the cell. Because of the poor response of the eye at 780 nm, the trapped atoms can be seen by eye only if the room is quite dark. For aligning the trapping laser beams you can use the IR phosphorescent card or a piece of white paper if the room is darkened. A 1 cm² photodiode with a simple current-to-voltage amplifier is used for making quantitative measurements on the trapped atoms. It is placed at any convenient position that is close to the trap, has an unobstructed view, and receives relatively little scattered light from the windows. The photodiode is used to detect the 780 nm fluorescence from the atoms as they spontaneously decay to the ground state from the 5P_{3/2} level. This measurement can be quantified and used to determine the number of trapped atoms. The same or a similar photodiode can also be used to look at the absorption by the trapped atoms and to monitor the rubidium pressure in the cell by measuring the absorption of a probe laser beam.

Although the trap fluorescence is large enough to easily detect with the photodiode, it can be obscured by fluorescence from the background rubidium vapor or by scattered light from the cell windows. Over a large range of pressures, the fluorescence from the background gas will be smaller than that from the trapped atoms. However, the scatter from the windows is likely to be significant under all conditions. The scattered light background will simply be a constant offset on the photodiode signal. For studying small numbers of atoms in the trap, however, the noise on this background can become a problem. In this case, you can use a lens to image the trap fluorescence onto a mask which blocks out the unwanted scattered light but allows the light from the trapped atoms to reach the photodiode.

6.2 TRAP OPERATION

When the cell and all the optics are in place, the first step is to turn on the dispenser to put rubidium into the cell. Initially, monitor the absorption of a weak probe beam through the cell to determine the rubidium pressure. Although the trap will operate over a wide range of pressures, a good starting point is to have about 1%/cm absorption on the F=2 to upper states transition in the region of the trap. At this pressure it is possible to see dim lines of fluorescence where the trapping beams pass through the cell when the trapping laser is tuned to one of the rubidium transitions. It is often easier to identify this fluorescence by slowly scanning the laser frequency and looking for a change in the amount of light in the cell. While absorption measurements are valuable for the initial setup and for quantitative

measurements, in the standard operation of the trap, use the setting on the dispenser current and/or the observation of the fluorescence to check that the pressure is reasonable.

After an adequate rubidium pressure has been detected, the magnetic field gradient is turned on and the lasers are set to the appropriate transitions. If the apparatus is being used for the first time, it will be necessary to try both directions of current through the field coils to determine the correct sign for trapping. The trapped atoms should appear as a small bright cloud, much brighter than the background fluorescence. Pieces of dust on the windows may appear nearly as bright, but they will be more localized and can be easily distinguished by the fact they do not change with the laser frequencies or magnetic field. If the trap does not work (and the direction of magnetic field and the laser polarizations are set correctly) the lasers are probably not set on the correct transitions. The trapped atom cloud may vary in size; it can be anywhere from less than 1 mm in diameter to several millimeters. Blocking any of the beams is also a simple method for distinguishing the trapped atoms from the background light.

6.3 MEASUREMENTS

Although many other more complicated measurements could be made with the trap, you should start by making the two most straightforward measurements: the number of trapped atoms and the time that atoms remain in the trap. If you have time after completing these, you can consider other experiments. These measurements are both made by observing the fluorescence from the trapped atoms with a photodiode. The number of atoms is determined by measuring the amount of light coming from the trapped atoms and dividing by the amount of light scattered per atom, which is calculated from the excited state lifetime. The time the atoms remain in the trap is found by observing the trap filling time and using Eq.(1). To make a reliable measurement of the number of trapped atoms, it is crucial to accurately separate the fluorescence of the trapped atoms from the scattered light and the fluorescence of the background vapor. To do this one must compare the signal difference between having the trap off and on. Therefore, the trap must be disabled in a way that has a negligibly small effect on the background light. We have found that turning off or, even better, reversing the magnetic field is usually the best way to do this. The magnetic field may alter the background fluorescence, but this change is generally smaller than the signal of a typical cloud of trapped atoms. Check how big an effect this is. Once you have determined the photocurrent due to just the trapped atoms, the total amount of light emitted can be found using the photodiode calibration of 0.3 mA/mW and calculating the detection solid angle. The rate R at which an individual atom scatters photons is given by

$$R = \frac{\pi\Gamma(I/I_s)}{1 + (I/I_s) + 4(\Delta/\Gamma)^2} \quad (4)$$

where I is the sum of the intensities of the six trapping beams, Γ is the 6 MHz natural linewidth of the transition, Δ is the detuning of the laser frequency from resonance, and I_s is the 4.1 mW/cm² saturation intensity. The simplest way to find Δ is to ramp over the saturated absorption spectrum and, when looking at the locking error signal, find the position of the lock point (zero crossing point) relative to the peak of the line. The frequency scale for the ramp can be determined using the known spacing between two hyperfine peaks. A typical number for R is 6×10^6 photons / (s \times atoms). One can optimize the number of atoms in the trap by adjusting the position of the magnetic field coils, the size of the gradient, the frequencies of both trapping and hyperfine pumping lasers, the beam alignments, and the polarization of the beams. More than 10^7 trapped atoms have been obtained when the rubidium pressure is large enough to dominate the lifetime.

The filling of the trap can also be observed using the same photodiode signal. This is best done by suddenly turning on the current to the field coils to produce a trap. The fluorescence signal from the photodiode will then follow the dependence given by Eq. (1), as shown in Figure 11. This can be most easily observed by sending the photodiode signal into a storage scope. The value of the 1/e trap lifetime τ




Figure 11. Typical curve of number of trapped atoms versus time after the trap is turned on at 0 sec, as shown on oscilloscope.

(the characteristic time an atom remains trapped) can then be determined from this curve. Eq. (1) tells us that it is the same as the time for the trap to fill to $1/e$ of its final value. Lifetimes between a fraction of a second to a few seconds are reasonable.

By changing the current through the dispenser, you should vary the rubidium pressure and see how it affects the number and lifetime of the trapped atoms. This can be compared with the predictions of Eqs. (1) and (2) and can be used to determine the collision cross sections if you also measure the rubidium density in the cell. There are many other experiments that can be done with the trapped atoms. The choices are only limited by your imagination and/or available equipment. You might look at how the number of atoms and the lifetime depends on the various parameters such as magnetic field, laser intensity, etc. You could also think about how to measure the spring constant and

damping constant for the trap. You can also turn off the magnetic field and watch the atoms spread out. In this situation, known as optical molasses and discussed in references 6 and 7, the laser light provides damping of the velocity but no trapping. In these conditions the atoms reach the lowest possible temperatures, but it is necessary to use additional field coils to cancel the magnetic fields from the earth and the ion pump. If you do, you will be able to watch the atoms spread out slowly (in a fraction of a second or longer) if the laser is detuned well to the red of the resonance line.

We conclude with both a challenge and a warning about all such possible experiments: you are likely to observe phenomena that have not yet been studied or are just being studied in research labs around the world. Often you will not find explanations (or will find incorrect explanations) in the current research literature, your instructor will not be able to provide much help, and the subject is too new for you to find textbooks providing answers. On the other hand, this means that if you do observe new phenomena and can explain them you will be able to publish your results in a scientific journal.

7 REFERENCES

1. K. B. MacAdam, A. Steinbach, and C. Wieman, "A narrow-band tunable diode laser system with grating feedback and a saturated absorption spectrometer for Cs and Rb," *Am. J. Phys.*, 60, 1098-1111 (1992) and the Senior Lab manual.
2. S. L. Gilbert and C. E. Wieman, "Laser Cooling and Trapping for the Masses," *Optics and Photonics News*, 4, 8-14 (1993).
3. The initial demonstration and discussion is E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. Pritchard, "Trapping of neutral sodium atoms with radiation pressure," *Phys. Rev. Lett.* 59, 2631-2634 (1987). More detailed analysis of the trap is given in A. M. Steane, M. Chowdhury, and C. J. Foot, "Radiation force in the magneto-optical trap," *J. Opt. Soc. Am B* 9, 2142 (1992), while a discussion of many of the novel aspects of the behavior of atoms in the trap is given in D. Sesko, T. Walker and C. Wieman, "Behavior of neutral atoms in a spontaneous force trap," *J. Opt. Soc. Am. B* 8, 946-958 (1991).
4. C. Monroe, W. Swann, H. Robinson, and C. Wieman, "Very cold atoms in a vapor cell," *Phys. Rev. Lett.* 65, 1571-1574; K. Lindquist, M. Stephens, and C. Wieman, "Experimental and theoretical study of the vapor-cell Zeeman optical trap," *Phys. Rev. A* 46, 4082-4090 (1992).
5. T.W. Hansch and A. L. Schawlow, "Cooling of gases by laser radiation," *Opt. Commun.* 13, 68-69 (1975).

6. S. Chu, L. Hollberg, J. Bjorkholm, A. Cable, and A. Ashkin, "Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure," *Phys. Rev. Lett.* 55, 48-51 (1985).
7. J. Dalibard and C. Cohen-Tannoudii, "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models," *J. Opt. Soc. Am. B* 6, 2023-2045 (1989); P. Ungar, D. Weiss, E. Riis, S. Chu, "Optical molasses and multilevel atoms: theory," *J. Opt. Soc. Am. B* 6, 2058-2072 (1989).
8. L. Orozco, SUNY at Stony Brook, private communication.

8 Prelab Exercises

1. (Section 2.1)

Calculate the force on, and the acceleration of, the atoms when they are scattering photons at the maximum rate from a laser beam. Include this result in your write-up. An atom can absorb and reemit a photon every 2 natural lifetimes when it is in very intense laser light that is at the resonance frequency of the atom.

2. (Section 2.2)

a) What is the purpose of the re-pumping laser? What would happen in the MOT without it? b) Your MOT uses the $5S_{1/2} F=2 \rightarrow 5P_{3/2} F'=3$ transition for the trapping and cooling transition and requires a re-pumping laser for the $5S_{1/2} F=1 \rightarrow 5P_{3/2} F'=2$ transition. What other transition scheme might work for trapping and cooling ^{87}Rb with two lasers? Explain.

3. (Section 3)

Calculate how large a change in the length of a 3 cm laser cavity will shift the frequency by 1 MHz. Hint: an integer number m of half wavelengths must fit in the cavity so $3 \text{ cm} = m\lambda / 2$, and $c = \lambda\nu$.

4. (Section 5.3)

Recall that the magnetic field is provided by two freestanding coils 1.3 cm in diameter with 25 turns each of 24 gauge magnet wire and a separation of 3.3 cm. How much current do you need to get a gradient of 15 G/cm? How much power? (Hint, the coils should get warm, but not burn up!).