Inexpensive laser cooling and trapping experiment for undergraduate laboratories

Carl Wieman and Gwenn Flowers
Joint Institute for Laboratory Astrophysics and Department of Physics, University of Colorado, Boulder, Colorado 80309

Sarah Gilbert
National Institute of Standards and Technology, Boulder, Colorado 80303

(Received 14 July 1994; accepted 15 December 1994)

We present detailed instructions for the construction and operation of an inexpensive apparatus for laser cooling and trapping of rubidium atoms. This apparatus allows one to use the light from low power diode lasers to produce a magneto-optical trap in a low pressure vapor cell. We present a design which has reduced the cost to less than $3000 and does not require any machining or glassblowing skills in the construction. It has the additional virtues that the alignment of the trapping laser beams is very easy, and the rubidium pressure is conveniently and rapidly controlled. These features make the trap simple and reliable to operate, and the trapped atoms can be easily seen and studied. With a few milliwatts of laser power we are able to trap 4×10^7 atoms for 3.5 s in this apparatus. A step-by-step procedure is given for construction of the cell, setup of the optical system, and operation of the trap. A list of parts with prices and vendors is given in the Appendix. © *1995 American Association of Physics Teachers.*

I. INTRODUCTION

Laser cooling and trapping of neutral atoms is a rapidly expanding area of physics research which has seen dramatic new developments over the last decade. These include the ability to cool atoms down to unprecedented kinetic temperatures (as low as 1 μK) 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. Undoubtedly one of the distinct appeals of this research is the leisurely and highly visible motion of the laser cooled and trapped atoms. Because of this visual appeal and the current research excitement in this area, we felt that it was highly desirable to develop an atom trapping apparatus that could be incorporated into undergraduate laboratory classes.

This paper presents a detailed discussion of how to build a simple and inexpensive atom trapping apparatus for atomic rubidium (Rb). Our principal goal was to develop an apparatus which could be built and operated reliably with minimal expense and technical support. In most respects, however, this trap's performance is equal or superior to what is

 © 1995 American Association of Physics Teachers

achieved with the "traditional" designs used in many research programs, and some innovations have advantages over these designs. Thus portions of this paper are likely to be of interest to the researcher already working (or considering working) in the field of laser trapping.

This paper is written in the same style as the previous paper on grating-stabilized diode lasers and saturated absorption spectroscopy.[1] It is intended to provide a "cookbook" discussion that will allow a relative novice to construct and operate an optical trap. This paper is essentially a continuation of the work presented in Ref. 1, and thus the end of that paper is used as a starting point. Without further discussion, we assume that one has two diode lasers of the type discussed in Ref. 1 which produce 5 mW or greater of narrow band tunable light. A small fraction (~10%) of each laser beam is split off and sent to its saturated absorption spectrometer of the type discussed in Ref. 1. This allows for precise detection and control of the laser frequencies, which is essential for cooling and trapping.

The remainder of this paper will discuss how to use the light from these lasers to trap Rb atoms. Section II provides a brief introduction to the relevant physics of the atom trap; Sec. III covers laser stabilization; Sec. IV explains the optical layout; Sec. V details the construction of the trapping cell; and Sec. VI discusses the operation of the trap, measurement of the number of trapped atoms, and measurement of the time the atoms remain in the trap. In the Appendix we provide a list of parts needed in the construction of the apparatus with prices and vendors. In its least expensive version the trapping apparatus, not including the lasers, costs less than $3000, with the ion pump responsible for half of the cost.

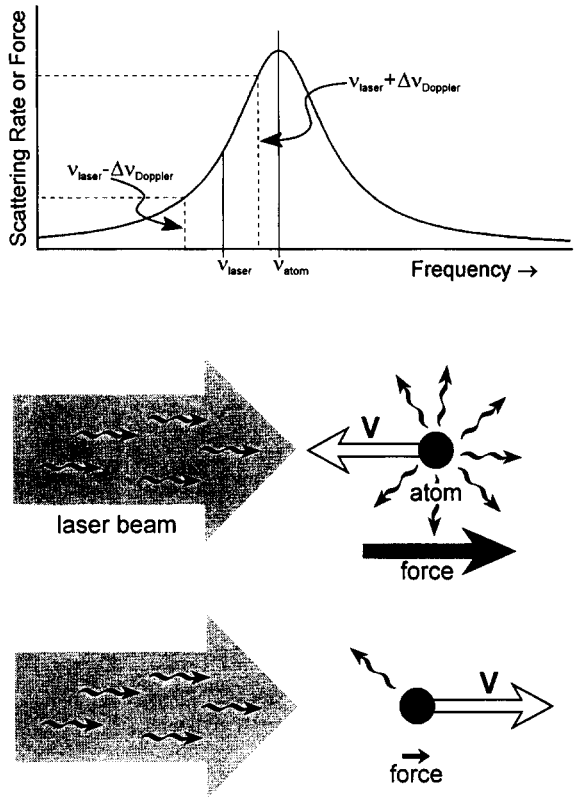

Fig. 1. Graph of the atomic scattering rate versus laser frequency. As shown, a laser is tuned to a frequency below the peak of the resonance. Due to the Doppler shift, atoms moving in the direction opposite the laser beam will scatter photons at a higher rate than those moving in the same direction as the beam. This leads to a larger force on the counter propagating atoms.

II. 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 nontechnical discussion is given in Ref. 2, while more detailed discussions of the magneto-optical trap and the vapor cell trap can be found in Refs. 3 and 4, respectively.

A. 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 produce large accelerations ($10^4 \cdot g$). 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 Fig. 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 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]

B. 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 Fig. 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

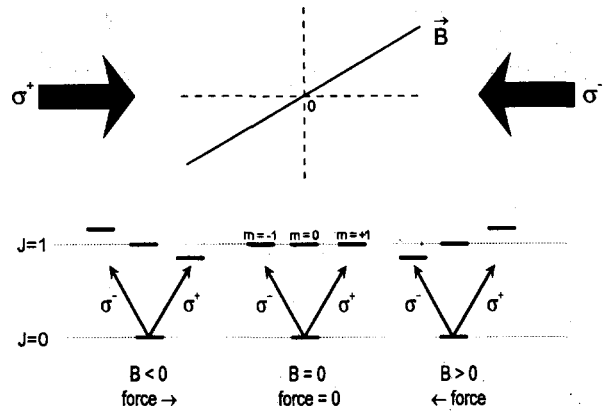

Fig. 2. One dimensional explanation of the MOT. Circularly polarized laser beams with opposite angular momenta impinge on an atom from opposite directions. The lasers excite the $J=0$ to $J=1$ transition. The laser beam from the right only excites to the $m=-1$ excited state, and the laser from the left only excites to 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 which pushes the atoms into the center.

right. Because of its angular momentum, the 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, 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 Fig. 3. As in optical molasses, laser beams illuminate the atom from all

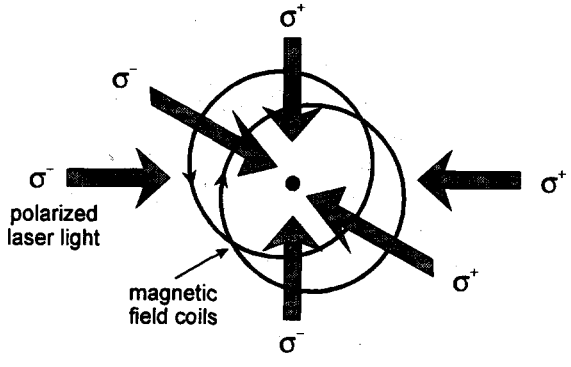

Fig. 3. Schematic of the MOT. Lasers beams are incident from all six directions and have angular momenta as shown. Two coils with opposite currents produce a magnetic field which is zero in the middle and changes linearly along all three axes.

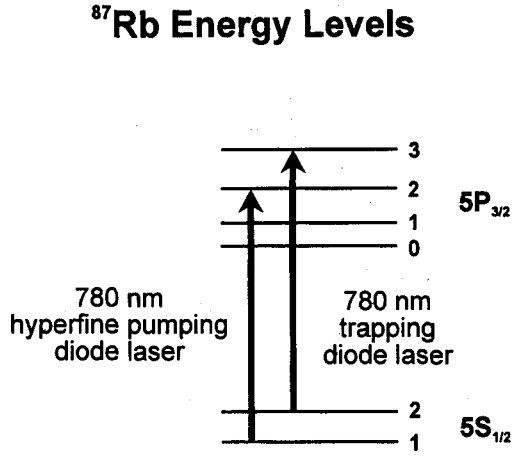

Fig. 4. ^{87}Rb energy level diagram showing the trapping and hyperfine pumping transitions. The atoms are observed by detecting the 780 nm fluorescence as they decay back to the ground state.

six directions. Two symmetric magnetic field coils with oppositely directed currents create a magnetic field which 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 Sec. II A. 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) $V\geq(\hbar\Gamma/m)^{1/2}$, where m is the mass of the atom. This velocity is often known as the "Doppler limit" velocity.[7] 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.

C. Rb vapor cell trap

We will now consider the specific case of Rb (Fig. 4). 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 $5S_{1/2}F=2\rightarrow 5P_{3/2}F'=3$ 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 Rb vapor.[4] The Rb atoms in the low energy tail ($V<V_{max}\approx 20$ m/s) of the Maxwell–Boltzmann distribution are captured in the laser trap. If the trap is turned 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.

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

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 (Rb 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 Rb and non-Rb components as

$$1/\tau=n_{Rb}\sigma_{Rb}V_{Rb}+n_{non}\sigma_{non}V_{non}. \tag{2}$$

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 Rb density, $(V_{max})^4$, and the surface area A of the trap. When the background vapor is predominantly Rb, the loss and capture rates are both proportional to Rb pressure. In this case N_0 is simply[4]

$$N_0=(0.1A/\sigma_{Rb})(V_{max}/V_{avg})^4, \tag{3}$$

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

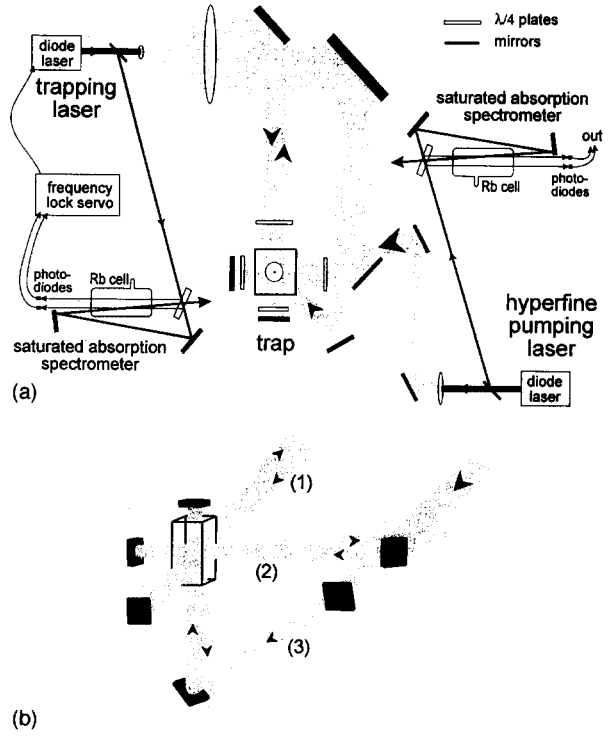

Fig. 5. (a) Overall optical layout for laser trap experiment including both saturated absorption spectrometers. (b) Detail of how laser beams are sent through the trapping cell. To simplify the figure, the λ/4 waveplates are not shown. Beam paths (1) and (2) are in the horizontal plane and beam (3) is angled down and is then reflected up through the bottom of the cell. The retrobeams are tilted slightly to avoid feedback to the diode laser.

D. Overview of the trapping apparatus

Figure 5(a) shows a general schematic of the trapping apparatus. It consists of two 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 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 Rb 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.

III. LASER STABILIZATION

Reference 1 describes how to construct a diode laser system. As mentioned above, two lasers are needed for the trap. Since a few milliwatts of laser power are plenty for hyperfine pumping ($F=1\rightarrow F'=1,2$), the stronger should be used as the trapping laser ($F=2\rightarrow F'=3$) if the two lasers have different powers. Clearly visible clouds of trapped atoms can be obtained with as little as 1.5 mW of trapping laser power (after the saturated absorption spectrometer pickoff). However, the number of trapped atoms is proportional to the laser power, so setting up and using the trap is much easier with at least 5 mW of laser power. The trapping laser must have an absolute frequency stability of a few megahertz. This normally requires one 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 exact laser frequency and then using a feedback loop to hold the length of the laser cavity constant.

Detailed discussions of the saturated absorption spectrometer, the servoloop circuit, and the procedure for locking the laser frequency are given in Ref. 1, along with examples of spectra obtained. 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 a few 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. These problems can normally be solved by changing the alignment and making sure that all the optical components are rigidly mounted. If the fluctuations in the unlocked laser frequency are large it will probably be necessary to improve the vibration isolation of the laser and/or table, or rebuild the laser to make it more mechanically stable. If the fluctuations are stable and synchronized with the 60 Hz line voltage, the problem is probably not mechanical, but rather is the diode laser power supply. Under quiet conditions and reasonably constant room temperature, the laser should stay locked for many minutes and sometimes even hours at a time. Bumping the table or the laser will likely knock the laser out of lock. If the laser frequency does not have excessively large fluctuations before it is locked but jumps out of lock very easily, it is likely that the resonant frequency of the servoloop is too low. This is determined by the response of the diffraction grating mount when driven by the PZT. The resonance limits the gain of the servo and reduces its ability to correct for large deviations. The frequency of this resonance can be determined simply by turning up the lock gain until the servosystem begins to oscillate and measuring the frequency of the sine wave which is observed. We have found that a reasonable value for the resonance frequency is around 1 kHz. If it is much lower than this, the diffraction grating is too massive, the mounting is not stiff enough, or the PZT has not been installed properly.

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 servosystems, it is quite straightforward to construct a second feedback loop which adjusts the laser current. The 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 servoloops 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 servoloops 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 servoloop can be obtained by sending a request and a self-addressed stamped envelope to C. Wieman.) This extra servoloop is a complication, however, which may not be desirable for an undergraduate laboratory.

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 onto the peak by slightly adjusting the PZT offset control on the ramp box every 5 to 10 min. If better control is desired, this frequency can also be locked to the side of the $F=1\rightarrow F'=2$ saturated absorption peak using the above procedure. However, it is often more convenient to lock the frequency of this laser to the peak of the $F=1\rightarrow F'=1$, 2, 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. This is an excessively expensive solution, unless lock-in amplifiers which are too old to use for anything else are available. These work very well for this purpose and this relatively crude frequency lock is adequate for the hyperfine pumping laser.

IV. OPTICAL SYSTEM

With lasers that can be locked to the correct frequency, one is now ready to set up the optical system for 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 our optical design for this student laboratory trap, we first mention the design used in most of the traps in our 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 the 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 counterpropagating beams, with the reflected beams having circular polarization opposite to the original beams.

A. Geometry

Here, we present 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 which expands it to a 4.5 cm×1.5 cm ellipse. This is then split into three beams which are roughly square (1.5 cm×1.5 cm), simply by clipping off portions of the beam with mirrors which interrupt part of the beam, as shown. The operation of the trap is remarkably insensitive to the relative amounts of power in each of the beams; a factor of 2 difference causes little change in the number of trapped atoms. 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 it is difficult to find optical elements to accommodate them and it becomes harder to see the beams due to the reduced intensity. Beams of 1.5 cm diameter work well and fit on standard 2.5 cm optics. With larger optical components it is possible to use larger beams. However, 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.

Two of the three beams remain in a horizontal plane and are sent into the cell as shown in Fig. 5(b). The third is angled down and reflects up from the bottom of the cell. There is nothing special about the layout of the three beams we have shown here. The only requirement is that they all intersect in the cell at roughly right angles. There is a minor benefit in having all of the beams follow nearly equal path lengths to the cell to keep the sizes matched, if the light is not perfectly collimated. To simplify the adjustment of the polarizations, the light should be kept linearly polarized until it reaches the quarter-wave plates. As discussed in any introductory optics text, 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. It is easy to check the ellipticity of the polarization by using a photodiode and a rotatable linear polarizer. First, make sure that the polarizer works at 780 nm by looking at the extinction of the laser light by crossed polarizers; some plastic sheet polarizers do not work at 780 nm. An extinction ratio 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 retromirrors 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. If the retro mirrors are close enough to the cell, they can be tilted by up to 30° without significantly changing the number of trapped atoms.

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.

B. 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 which propagate through the cell perpendicular to the coil axis should have the same circular polarization, while the beam which propagates along the axis 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 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 which 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 need only to be set to within about ±10°. 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 orientations are, after the beams have passed though them twice, the light's angular momentum will be reversed.

We have tried various options for wave plates. The most straightforward option is to simply use six commercial $\lambda/4$ plates which have a clear aperture large enough for the laser beams. Ideally, these should be antireflection coated to reduce the attenuation of the laser beams. The only disadvantage to this approach is the cost. There are many suppliers of $\lambda/4$ plates, but a set of six antireflection coated wave plates will cost over $1000. An alternative to using expensive wave plates is to replace all or some of the commercial $\lambda/4$ plates with inexpensive plastic sheet retardation plates, which are widely available in optics demonstration kits. We have replaced the 780 nm $\lambda/4$ plates with plastic sheet which was nominally $\lambda/2$ at 500 nm (and hence probably about $\lambda/3$ at 780 nm) and observed little change in the number of trapped atoms. One difference in this case is that the trapping is more sensitive to the orientation of the wave plates and, in particular, is now somewhat sensitive to the orientation of the retro wave plates. Another option is to replace the three retro $\lambda/4$ plates with retroreflecting right angle mirrors.[8] Although this

combination of mirrors does not provide an ideal $\lambda/2$ retardance, we found empirically that for five different dielectric mirrors tested, it is quite close. For gold mirrors the retardance is much farther from $\lambda/2$, but still adequate for good trapping. Putting the two right angle mirrors on separate mirror mounts gives more flexibility of adjustment, but is probably unnecessary. Simply gluing the two mirrors together at an angle of slightly more or less than 90° (to insure that the beams overlap in the cell, but the return beams does not go directly back into the laser) works fine. 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.

Finally, there are other options for obtaining inexpensive retardation plates that are approximately $\lambda/4$ at 780 nm. Most transparent plastic sheet is birefringent, and often the retardance can be adjusted by stretching. If one simply tries a number of pieces of plastic sheet, it is likely that some can be found which are not too far from a $\lambda/4$ retarder. The traditional way of making inexpensive $\lambda/4$ plates is to split sheets of mica. We have not tried this, but it should work if the mica does not absorb too much light at 780 nm.

C. Hyperfine pumping laser optics

Minimal optics are needed for the hyperfine pumping laser. Essentially the only requirement for this light is that it cover the region where the trapping laser beams overlap. In its simplest form this means merely sending the laser through a single lens to expand the beam and reflecting it off a mirror into the trapping cell through any window. In practice, it is often worth going to the small additional effort of using two lenses to make a large collimated beam (typically an ellipse with a 2–3 cm major axis) 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.

D. Mirrors and mirror mounts

If low cost $\lambda/4$ retarders are used, the cost of the trapping optics will be dominated by the ten mirrors and mirror mounts required. We have become accustomed to using convenient commercial kinematic mirror mounts (see the Appendix). Similar products are available from many manufacturers, but a price between \$50 and \$100 per mount is standard. However, the control and stability provided by such mounts are actually far superior to what is necessary in this application. Thus if one has severe budget limitations, simple homemade mounts with limited adjustment should be adequate. We use both gold and high reflectance infrared dielectric mirrors (aluminum is lossy at 780 nm). Although gold mirrors are adequate for any of the mirrors in the setup, we prefer to use the dielectric mirrors for the retroreflection, because their higher reflectance makes the trap more symmetric. Also, while gold mirrors are less expensive than the dielectric mirrors, they are more easily damaged. Gold mirrors with protective overcoatings are readily available, but in our limited experience, mirrors with overcoatings which rival the durability of a dielectric mirror end up costing about the same. In an environment where inexperienced students will be handling the mirrors frequently, inexpensive gold mirrors have a rather short lifetime, in contrast to the standard commercial dielectric coated mirrors. In the long run, we think that the least expensive option for mirrors is to buy large dielectric mirrors and cut them into smaller mirrors with a glass saw. (Put a protective cover such as plastic tape on the surface before cutting.) In the short run, however, it is less expensive to use gold mirrors. The surface quality and flatness of any commercial mirror will be more than adequate. Caution should be taken when using dielectric mirrors at an angle; there can be substantial variation of reflectivity with different incidence angles and polarizations.

V. 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 ($\approx 10^{-7}$ Torr), trap lifetimes long enough for most experiments of interest require pressures in the 10^{-6} to 10^{-7} Pa range. Unfortunately, UHV often seems to be synonymous with high cost and specialized technical expertise. Consequently, we have put considerable effort into designing a simple, inexpensive system that can be constructed by someone who does not have experience with ultrahigh vacuum techniques. However, we will present a range of possibilities to best cover the needs of a variety of different users. 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 Rb atoms; and (3) windows to transmit the laser light and allow observation of the trapped atoms. We will discuss each of these elements in turn, and then discuss in detail the actual construction of the cell.

A. Vacuum pump

In all of our systems we have used small ion pumps (2, 8, or 11 ℓ/s). Ion pumps have the virtues that they are small, quiet, and light, use little power, require no cooling, and have low ultimate pressures. These features make them well suited to this application. The 2 ℓ/s pumps have been adequate for a clean system which is evacuated once and remains at high vacuum thereafter. However, in systems which have been cycled up to air a few times the performance of the 2 ℓ/s pumps has not been satisfactory. After very few pumpdowns, they become increasingly difficult to start and soon will not pump at all. We recommend an 8–12 ℓ/s pump if there is any possibility of letting the system up to air more than once. The power supply for a small ion pump is often more expensive than the pump itself. It is possible to avoid this expense by using a dc high voltage supply, if one is available. The pump will typically require a few tenths of 1 μA at 5–6 kV under normal operating conditions, and as much as a few milliamperes at lower voltage for the first few minutes when it is initially turned on. Any power source which will supply sufficient voltage and current will be adequate. It is quite convenient to be able to monitor the pump current as a measure of the pressure. After a few days of exposure to Rb vapor the pump is likely to have some leakage current for which a correction must be made when determing the pressure, but the leakage does not affect pump performance.

These pumps have the minor drawback that they require 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 magnetically permeable iron or steel sheet around the pump to

shield the trap from this field.[9] For this same reason it is advisable to avoid having magnetic bases very near the trap.

B. Rb sources

We will now discuss how to produce the correct pressure of Rb vapor in the cell. The vapor pressure of a room temperature sample of Rb is about 5×10^{-5} Pa. This is much higher than the 10^{-6} to 10^{-7} Pa ($\approx10^{-8}$ to 10^{-9} Torr) of Rb 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 difficult to see them because of the bright fluorescence from the untrapped background atoms. Also, the absorption of the trapping beams when passing through the cell will be significant. Thus we would like to maintain the Rb 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 Rb to maintain the correct pressure.

Before discussing specific ways to generate Rb in the vapor cell, we will present some background on the behavior of Rb in the environment of the cell.[10] This behavior is often counterintuitive. Through chemical reactions and physisorption, the walls of the cell usually remove far more Rb than the ion pump does and the rate of pumping by the walls depends on how well they are coated with Rb. If one has an evacuated stainless steel tube and connects a reservoir of room temperature Rb to one end, initially no Rb vapor will be observed at the other end. Over time the vapor will creep down the tube as it saturates the walls, and in a matter of several days, the pressure differential across the length of the tube will significantly decrease. If the Rb reservoir is then removed and the tube is connected to a pump, the vapor will linger for days as it slowly comes off the walls. If the tube is heated, the vapor will disappear much more quickly. A similar behavior occurs with glass, but the chemical reaction rates vary with different types of glass. With fused silica or Pyrex, the behavior is similar to that observed with stainless steel, and a similar saturation of the chemical reaction rate occurs. With other types of glass (microscope slides, for example) the reaction rate and thus the wall pumping remain large even after extensive exposure to Rb vapor. Since this rapidly depletes Rb and causes large pressure differentials in the cell, we do not recommend making a cell from this type of glass.

Our recommended technique for producing the Rb vapor for a student lab experiment is to use a commercial "getter." This technique was developed specifically for our student lab trap, and, to our knowledge, has not been used before in optical trapping cells. However, as we will discuss, it has advantages over other approaches and is likely to be useful in many research lab traps as well. The getter is several milligrams of a Rb compound which is contained in a small ($1.0\times0.2\times0.2$ cm) stainless steel oven. Two of these ovens are spotwelded onto two pins of a vacuum feedthrough as shown in Fig. 6. When current (3–5 A) is sent through the oven, Rb vapor is produced. The higher the current, the higher the Rb pressure in the chamber. With this system it is unnecessary to coat the entire surface with Rb, and is in fact undesirable since the getter is likely to be exhausted before the surface is entirely saturated. With a glass cell mounted on a 1.33 in. Conflat-type fiveway cross we are able to produce enough Rb pressure to easily see the background fluorescence with as little as 3.4 A through the getter oven. If any

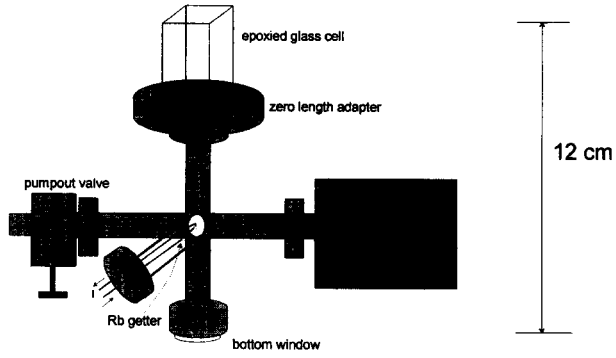

Fig. 6. Drawing of trapping cell. The tubes on the cross have been elongated in the drawing for ease of display. The Rb getter is inside a stainless steel vacuum tube; we show a cut-away view in this drawing. The zero length adapter is a 2.75 in. to 1.33 in. Conflat-type adapter.

fluorescence from the background can be seen, there is more than enough Rb for trapping. According to the data sheets, this is a negligible production rate of Rb, and hence the getter should last a long time. We have operated for 100 h with no sign of depletion of the Rb getter.

The principal advantages of this approach are the simplicity of construction and the ability to quickly and easily adjust the Rb pressure. This is particularly important for a lab class experiment where the students will be able to work for only a few hours at a time. After the current through the getter has been turned on, the Rb vapor comes to an equilibrium pressure with a time constant of about 5 min. The overall pressure in the system rises slightly when the getter is first heated, but after degassing overnight by running a current just below that which produces observable Rb, the pressure rise is less than 10^{-6} Pa, and continues to decrease with further use of the getter. When the current through the getter is turned off, the Rb pressure drops with a time constant of about 4 s if the getter has been on for only a short while. With prolonged use of the getter, this time constant can increase up to a few minutes as Rb builds up on the walls, but it decreases to the original value if the getter is left off for several days. Presumably this recovery would be much faster if the system was heated.

The superior Rb pressure control provided by the getter 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 getter 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 untrapped atoms.

We and others have used two alternative methods to produce Rb vapor. We mention the first only because it is used for many ordinary vapor cells, but is not a very good choice for this trap. In this approach a fraction of 1 g of Rb is distilled into the cell under vacuum and is condensed into a thermoelectrically cooled "cold finger." The vapor pressure in the cell is then adjusted by changing the temperature of the cold finger and allowing the pressure in the remainder of the cell to come to equilibrium. We do not recommend this

approach because the pressure can only be changed very slowly, and (most importantly!) we have found that it is easy to accidently turn off or burn out the thermoelectric cooler. When this happens the entire vacuum system is filled to a relatively high Rb pressure. This is hard on the ion pump, and it is usually difficult and tedious to bake all the Rb back into the cold finger.

A better method is to have the Rb reservoir behind a small stainless steel valve. This approach has been used in most of our research work. A small sealed glass ampoule containing about 1 g of Rb is installed in a thin-walled stainless steel or copper tube. We usually distill the Rb into the glass ampoules ourselves to insure that the Rb does not contain other dissolved gases, but it is also possible to buy such ampoules commercially. One end of the metal tube is sealed, and the other end is welded or brazed onto a 1.33 in. Conflat-type flange which is attached to a valve connected to the main cell. After the system has been thoroughly pumped out and baked, the tube is squeezed in a locking pliers (vice grips) until the metal wall has compressed enough to break the glass and the Rb is released. To get the Rb through the valve and into the cell, we heat the tube and the valve to 60–100 °C. The first time this is done it takes between several hours and several days to see a significant amount of Rb in the cell (depending on the geometry, the surfaces, and the temperature). Once the pressure in the cell reaches the desired level, the valve to the reservoir is closed. At this point, the vapor will slowly begin to leave the cell, and after some time (days to weeks), the heating process will need to be repeated. After the initial Rb coating of the inside of the cell, it only requires several (10–30) minutes to substantially raise the Rb pressure in the cell upon subsequent heating of the reservoir. The length of time between fills depends on the geometry of the system and the properties of the surfaces. For Pyrex and stainless steel systems it will typically be between a day and a few weeks. For systems with more reactive surfaces it can be just a few minutes, in which case one may simply leave the valve open while carrying out experiments.

C. Windows and cell assembly

The third important element of the trap cell is the optical access. This can be either through mounted windows, or through the cell walls themselves if they are transparent and flat. Optical access is necessary for sending the laser beams through the cell and for observing the fluorescence from the trapped atoms. Although there are many ways to assemble a cell with windows, we will discuss the three that are most practical along with their respective advantages and disadvantages. A component common to these designs is the Conflat-type vacuum flange. These flanges are a means of making all-metal vacuum seals which are leakfree, have very low outgassing, and can be baked to high temperatures. The stainless steel flanges are bolted together with a soft copper gasket between them. Knife edges on the flanges press into the copper to make the vacuum seal.

The first technique for making a trap cell is one which we have developed specifically for this paper. This design is shown in Fig. 6. Five pieces of plate glass are sealed together with epoxy. This unit is epoxied onto a commercial zero-length 2.75 in. to 1.33 in. Conflat-type adaptor flange which is bolted onto a fiveway cross with 1.33 in. Conflat-type flanges. The sixth window (the bottom window) is a commercial 1.33 in. Conflat-type viewing port. The viewing port, ion pump, pumpout valve, and Rb getter are attached to the other arms of the cross. The advantages of this design are (1) allows very good optical access and viewing due to the geometrical arrangement and the optical quality of the windows, (2) has relatively low cost, (3) it can be readily built by someone without any special skills in working with glass, and, of least importance, (4) allows the use of antireflection coated windows. Its principal disadvantage is that the ultimate pressure is limited by the temperature to which the epoxy can be baked. However, we have achieved pressures below 10^{-7} Pa in this type of cell using a small ion pump.

We will start by offering some advice on ultrahigh vacuum (UHV) techniques. We strongly recommend, however, that those not familiar with the subject also read about UHV techniques in one of the many reference books on vacuum technology. First, brand new commercial components such as the cross require no additional cleaning and should be kept in their protective wrapping until they are to be assembled. If the components are possibly dirty, they should be cleaned. Our standard cleaning procedure is to put the component in an ultrasonic cleaner with hot water and a good lab detergent for a few minutes. Then it is rinsed in turn with hot tap water, distilled water, and high purity alcohol. After drying in air it is wrapped in clean aluminum foil until it is to be used. Any surface which will be at high vacuum should never be touched with unprotected hands after it has been cleaned. If it must be touched, clean gloves should be used. It is also desirable to avoid having things like hair fall into the vacuum chamber, and it is always wise to carefully inspect the system for such items. Clean new copper gaskets should be used when making the Conflat-type seals, and the flanges should be tightened down uniformly. Graphite or some other high temperature antisiezing compound should be put on the bolts to prevent them from seizing up after they are baked. Also, we find that plate nuts make the assembly of Conflat-type seals much easier. Finally, to reach UHV it is necessary to bake even the cleanest systems under vacuum to remove adsorbed water vapor from the surface.

The glass part of the cell is a parallelepiped with its long axis oriented vertically. Typical dimensions of the vertical windows are about 2.75 cm by 4.5 cm, and the horizontal window on top is a square with side length 2.9 cm. The glass for the cell should be at least 2 mm thick; thinner glass may not be strong enough to withstand the pressure and much thicker glass can limit optical access. The windows should be cut to size with a standard glass saw. It is useful to cover the glass with masking tape while it is being cut to avoid scratching the surfaces. The precise dimensions of the glass pieces are not important, but it is important to have the cell fit together tightly to minimize the gaps that will be filled in with epoxy. To achieve this, one should make the opposite sides of the cell the same width, all four sides the same length, and all the corners of each piece perpendicular. To insure that the opposite sides are the same size, it is simplest to stack the two pieces and cut the stack as a single piece. After the pieces are cut, they should be washed carefully to remove any residual grit from the glass saw, and then the protective masking tape should be removed. The windows should then be cleaned just as the other vacuum components.

The final assembly is done by first bolting the adaptor flange onto the fiveway cross, and then assembling the cell on top of it. It should be done in this order to avoid stressing the epoxy seals while bolting on the flange. Just before assembly, the inside surfaces of the five pieces of glass should

be cleaned with lens tissue and pure alcohol or acetone to remove residual dust, which would scatter the laser light. Residual dust can be easily identified by looking for scattering while illuminating the window using a bright lamp. After the glass slides are clean they are pressed together to form a rectangular tube and placed upright on the flange. Two small rubber bands are placed around the tube to hold the four glass sides firmly pressed together in this position. The square window is then set on top of the piped and a small weight is put on it to hold everything in place while the epoxy is applied. The seams are then sealed with a bead of low-vapor-pressure epoxy. Warming the epoxy slightly makes it somewhat easier to apply, but the glass itself should not be heated. When the epoxy has hardened somewhat (30 min or so) the elastic bands can be cut off without putting any stress on the cell, and epoxy should then be applied to the seams which were covered by the bands. By assembling the cell in this manner, very little epoxy will be exposed to the vacuum system. The epoxy should be allowed to set for 24 h at room temperature.

The rest of the vacuum components can now be bolted in place. There are two items which deserve special mention. First, the pump-out valve should be a bakeable all-metal-seal valve to meet the UHV requirements. For systems which are not going to be let up to air very often, this valve can be replaced by a copper pinch-off tube, which acts as an inexpensive "one use" valve. After the system has been rough pumped and the ion pump started, the tube is pinched off leaving a permanent seal. This works well and avoids the expense of a valve. However, if the system is likely to be opened up a few times or if there is a possibility of leaks in the system, a valve is much more convenient. Second, the Rb getter comes as a powder packed in a small stainless steel boat or oven. It is best to spot weld two of these boats in series and then spot weld them between two 10 A feedthroughs on a 1.33 in. Conflat-type flange. They should be positioned as close to the trapping region as possible without blocking the path of the vertical laser beam. Protective gloves should be used when installing the getter to avoid contaminating the system.

Once the entire cell is assembled it is a good idea to perform a helium leak test if a leak detector is available. In the process, it is quite important to avoid contaminating the system with backstreaming of pump oil from the leak detector. If the leak detector does not have a liquid nitrogen cooled trap to prevent this, install one in the line connecting the trap cell to the leak detector.

D. System pumpdown and bakeout

The next step is to pump the system down to a low enough pressure to start the ion pump. Use whatever pumping options are available; a clean cryogenic sorption pump is usually sufficient and will not contaminate the system. If the rough pumping is done with a mechanical or diffusion pump it must be very well trapped to avoid oil contamination, as with the leak detector. We have found, particularly with the 2 ℓ/s ion pumps, that minimizing the gas load when the ion pump is started improves the long term performance of the pump. If the system was not previously helium leak tested, it should be leak tested after the ion pump has been started and the system closed off from the roughing pump. This testing can be easily done by squirting alcohol on all of the epoxied and copper gasket seals while monitoring the ion pump current. If there is a sudden change in the current when alcohol is sprayed on the seal, it signifies a leak which has just been covered by the alcohol. Although it is important to leak test the system, it is unlikely there will be any leaks if the Conflat seals have been installed properly and care is taken to insure that the epoxy is applied evenly over all the joints.

After confirming that there are no leaks, the system is then baked. One of the drawbacks to this design is the fact that it is quite easy to open up a leak in the glass-to-metal epoxy seal while baking. The reasons for this are the different thermal expansion coefficients of the glass and metal, and more importantly, the fact that the glass and metal tend to heat and cool at very different rates. This results in thermally induced stress which can break the rigid epoxy seal. This problem would be eliminated if a flexible sealant were used, but we have been unable to find one which satisfies the requirements of strength and low vapor pressure, and is also bakeable to 100 °C. After trying a number of baking methods which resulted in leaks, we developed the following procedure, which has not caused leaks the three times we used it. First, the entire system is put inside an oven with a bakeable cable connecting the ion pump to the power supply, which remains outside the oven. The ion pump current (and hence the system pressure) should be monitored while the system is baking. A standard kitchen oven is ideal for this bake. Once the system is in the oven, the temperature should be slowly increased. Normally, the rate at which the temperature can be increased is determined by the maximum operating current of the ion pump and the strong temperature dependence of the outgassing rate in the system. It is not unusual for it to take many (6–8) hours before the temperature can be increased up to the maximum baking temperature of 100 °C. If the epoxy has not been fully cured, it may be advisable to simply leave the system for a day at about 60 °C without attempting to increase the temperature further until the outgassing has dropped. Even if the pressure does not rise to undesirable levels, the temperature should not be changed faster than about 20 °C/h (up or down) to avoid thermal stresses. Once the temperature has been raised to 100 °C the system should be baked for at least a day. Then, it should be slowly cooled by progressively reducing the temperature of the oven. The pressure should drop dramatically as the cell cools and be less than 10^{-6} Pa at room temperature. If the pressure in the system does not drop dramatically as it cools, or even increases with cooling, it almost always indicates a leak. The most likely place for a leak will be in the glass-to-metal epoxy seal. If such a leak has opened, it can often be sealed by applying new epoxy to the offending region without ever letting the system up to air.

This method of baking puts a fairly heavy load of gas into the ion pump. This can be reduced by a "prebake" of the system while it is on the roughing pump. With most facilities, however, it is awkward to insure proper uniform control over the temperature at the rough pumping stage. Once the cell has been baked it is ready to use. The ion pump can be turned off for long periods (up to a few days) without the pressure rising enough to cause problems. This is often convenient while the cell is moved and installed for use.

E. Alternative designs

Before discussing the operation of the trap we will briefly mention two alternative methods of making trapping cells. The first is to have a glassblower make a six-sided glass cross with six windows on the tubes of the cross, and a glass-to-metal seal that allows the cell to be attached to the

ion pump, pump-out valve, and Rb source. The advantages of this system are that it is simple, it can be baked to very high temperatures, and it has very little tendency to leak. The disadvantages are (1) not everyone has a glassblower available, (2) it is very difficult to modify anything about the trap cell, (3) such designs are often rather fragile, (4) the optical access and trap visibility are limited due to the width of the fused glass seams, and (5) the glassblowing often causes distortions and scattering centers in the windows. Perhaps the more serious aspect of (4) is that this often results in the windows being so far from the center of the trap (depending on the skill and effort of the glassblower) that the retroreflecting mirrors are far away from the middle of the cell. As discussed above, this renders the alignment more difficult and can make it very difficult to have adequate overlap of the forward and backward going beams without introducing an unacceptable amount of feedback into the diode laser.

The second alternative is to buy a commercial 2.75 in. or 1.33 in. sixway cross with Conflat-type flanges (or a cube with six 2.75 in. Conflat-type seals), and six viewports (windows mounted on stainless steel Conflat-type flanges). Then one simply bolts five of the windows onto the sixway cross (or cube). A fourway cross is bolted on the sixth side which has the Rb reservoir and pump attached at right angles, and the sixth window attached to the port opposite the sixway cross. The advantages of this system are that all the components are readily available and can simply be bolted together. The principal disadvantages are the cost and the poor viewing and optical access. Also, metal scatters more light, causing a large background which can obscure the signal from the trapped atoms. These problems and the corresponding difficulty in optical alignment are so great in a sixway cross that we do not recommend it for use in an undergraduate lab. A cube is much better in this respect, but cubes with six Conflat-type seals are relatively expensive.

F. Installation on the optical breadboard

After the cell has been baked, it is installed on the optical breadboard. This is simple; one or two ringstand clamps can hold the cell and ion pump in place. It is advantageous to align the six trapping beams before the cell is installed, because testing the overlap of the beams is much easier when the cell is not in the way. The cell must be mounted so that the beams overlap roughly in the center of the glass region and there is enough room below the bottom window to allow the vertical beam to be reflected up from below the cell. Remember also to leave room for the $\lambda/4$ wave plate below the bottom window and room to adjust the orientation of the $\lambda/4$ plate and the bottom mirror. Since the cell is elevated, it will also be necessary to support the ion pump since it is the heaviest part of the system.

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. All other details of the coils are unimportant. We use two freestanding coils 1.3 cm in diameter with 25 turns each of 24 gauge magnet wire. This provides the desired gradients with a current of 2–3 A and a coil separation of 3.3 cm. The coils should be mounted on either side of the cell such that the current travels through the loops in opposite directions and the coil axis is collinear with one of the laser beam axes. Our coils are simply attached to the windows or supporting flange with tape or glue.

VI. OPERATION OF THE TRAP AND MEASUREMENTS

A. Observation system

Although not essential, it is highly desirable to have an inexpensive CCD TV camera and monitor with which to observe the trapping cell. Standard cameras used for security surveillance work well for this purpose. They 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 we usually use a standard IR fluorescent card or a piece of white paper if the room is darkened. (Note that although the laser beam appears weak, it is actually intense enough to cause damage if the beam shines directly into the eye.) A 1 cm^2 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 which 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 Rb 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 Rb 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. We have found that with minimal effort to reduce it, 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, one should use a lens to image the trap fluorescence onto a mask which blocks out the unwanted scattered light.

B. Trap operation

When the cell and all the optics are in place, the first step is to turn on the getter to put Rb into the cell. Initially, monitor the absorption of a weak probe beam through the cell to determine the Rb 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 Rb 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 getter current and/or the observation of the fluorescence to check that the pressure is reasonable.

After an adequate Rb 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. The 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. 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.

C. Measurements

Although many other more complicated measurements could be made with the trap, here we will discuss how to make the two simplest measurements: the number of trapped atoms and the time which atoms remain in the trap. Both of these measurements are 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 much smaller (1/100) than the signal of a typical cloud of trapped atoms. Once one has determined the photocurrent due to the trapped atoms, the total amount of light emitted can be found using the photodiode calibration and calculating the detection solid angle. The rate R at which an individual atom scatters photons is given by[11]

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

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/cm^2 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·atom). 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. With 7 mW from the trapping laser, we obtain nearly 4×10^7 trapped atoms when the Rb pressure is large enough to dominate the lifetime.

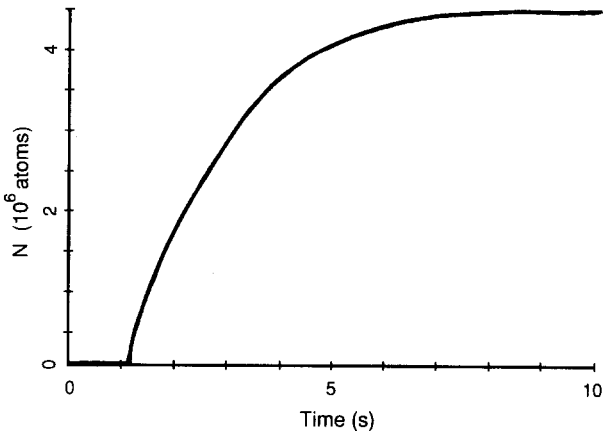

Fig. 7. Number of trapped atoms versus time after the trap is turned on at 1.2 s.

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 Fig. 7. The value of the trap lifetime τ (the characteristic time an atom remains trapped) can then be determined from this curve. We have observed lifetimes which ranged from nearly 4 s down to a small fraction of 1 s depending on the Rb pressure.

There are many other experiments which can be done with the trapped atoms. The choices are only limited by one's imagination and/or available equipment, but here we list a few examples. Although optical traps are being used in a number of experiments, many aspects of the trap performance have not been studied. By changing the current through the getter, one can vary the Rb 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 collision cross sections; to date, very few trap ejection cross sections have been measured. One can also measure the temperature of the trapped atoms, the spring constant of the trap, and how the number of atoms and the lifetime depend on the various parameters. To our knowledge the detailed dependence of these characteristics on the polarizations of the laser beams, the intensity in the various beams, or the magnetic field gradient have not been studied. You can quickly learn that, in fact, the trap will work well with one of the incident beams linearly polarized, and can even work marginally with two linearly polarized beams, but we know of no analysis or studies of these cases in the literature.

It is also relatively easy to do various kinds of high resolution spectroscopy on the trapped atoms because the optical thickness of the trapped atom cloud is substantial and the Doppler shifts are essentially zero. Using another diode laser one can precisely measure the $5S\rightarrow5P$ spectra with the hyperfine components clearly resolved. In addition, there are convenient transitions which allow one to use diode lasers to study excitation from the $5P$ state to higher levels.[12] It is also possible to observe very high resolution microwave transitions between the hyperfine states if a suitable source of microwave power is available.[13]

By turning the laser light off and adding small magnetic fields, it is possible to magnetically trap[4] or "bounce" the

laser-cooled atoms off of an inhomogeneous magnetic field. This dramatically demonstrates the quantization of angular momentum because the different Zeeman levels have different magnetic moments, and hence feel different magnetic forces. This causes the cloud to separate into a number of smaller clouds as it bounces, each of which represent a different projection of the angular momentum on the magnetic field. We conclude with both an invitation and a warning about many experiments one can do with optical traps: this is a new and rapidly changing field; one is likely to observe phenomena which are not explained in the current research literature and there are no textbooks to provide answers. The student and instructor may find themselves in uncharted territory.

ACKNOWLEDGMENTS

This work was supported by the NSF and the University of Colorado. Early design work on this experiment was carried out by K. MacAdam, A. Steinbach, and C. Sackett. We thank M. Stephens for considerable help and advice. G. Flowers acknowledges the support of a tuition scholarship from the NIST PREP Fellowship program.

APPENDIX: COMPONENTS USED IN TRAP CONSTRUCTION

Listed here are the components that we used in construction of the trap and the suppliers of these components. Commercial trade names are identified to make this article useful to the reader. Since we have not made a careful study of other products, we list only those that we actually used; the prices given are current as of 1993 and are meant for guidance only. No endorsement by NIST or the University of Colorado is implied; similar products by other manufacturers may work as well or better. Buyer's guides such as those published by Physics Today (with the August issue), Photonics Spectra, and Laser Focus World are good sources for listings of vendors.

(1) Optical Base # BA1S, 16 @ $8 each, Thorlabs, Inc., P. O. Box 366, Newton, NJ 07860, (201) 579-7227.
(2) Post holder # PH3-ST, 16 @ $12 each, Thorlabs, Inc. (see item 1).
(3) Posts # TR2, 3 @ $5 each; # TR3, 4 @ $6 each; # TR4 1 @ $7; # TR6 14 @ $8 each; # TR10 1 @ $10, Thorlabs, Inc. (see item 1).
(4) Kinematic mirror mounts, # KM-B, nine @ $50 each; # KMS, one @ $41, Thorlabs, Inc. (see item 1).
(5) Right angle post clamps # RA90, seven @ $12 each, Thorlabs, Inc. (see item 1).
(6) Commercial grade quarter-wave plates # RCQ-1.0-0780, 25.4 mm diameter, six @ $184 each (additional charge for antireflection coating). Meadowlark Optics, Inc., 7460 Weld County Road # 1, Longmont, CO 80504, (303) 776-4068. As noted in the text, there are other options for wave plates. These wave plates are destroyed by mild heating.
(7) Near-infrared dielectric mirror, 50.8 mm diameter, to be cut in quarters for four mirrors. Coating BD.2, Substrate Code 20D04, $243, Newport Corporation, 1791 Deere Ave., Irvine, CA 92714, (800) 222-6440.
(8) Flat gold mirror # G41, 801, 76×102 mm to be cut into six 25×25 mm square mirrors and one 25×50 mm mirror, $22, Edmund Scientific Co., 101 East Gloucester Pike, Barrington, NJ 08007-1380, (609) 573-6250.
(9) Rb metal getter Rb/NF/3.4/12/FT 10+10 Code 5g0125, $250, SAES Getters USA, Inc., 1122 E. Cheyenne Mtn. Blvd., Colorado Springs, CO 80906, (719) 576-3200.
(10) Eleven ℓ/s ion pump with magnet and bakeable cable, $880, Duniway Stockroom Corp., 1600 Shoreline Blvd., Mountain View, CA 94043, (800) 446-8811; 2 ℓ/s ion pump with square magnet and bakeable cable, $890, Perkin Elmer Physical Electronics Division, 6509 Flying Cloud Drive, Eden Praire, MN 55344, (612) 828-6100.
(11) High voltage ion pump power supply # 222-0242 Ionpak series 240, $550, Perkin Elmer, Physical Electronics Division, 6509 Flying Cloud Drive, Eden Praire, MN, 55344, (612) 828-6100.
(12) Reducing Flange, 2.75 in. to 1.33 in. Conflat-type # 150001, $55, MDC Vacuum Products Corp., 23842 Cabot Blvd., Hayward, CA 94545-1651, (800) 443-8817.
(13) Fiveway cross with 1.33 in. Conflat-type flange #406000, $160, MDC Vacuum Products Corp. (see above).
(14) Vacuum 1.33 in. Conflat-type viewport # VP-133-075, $91, Duniway Stockroom Corp., 1600 Shoreline Blvd., Mountain View, CA 94043, (800) 446-8811.
(15) Power feedthrough for Rb getter, Moly 10 A/Pin # EFT0024032, $125, Kurt J. Lesker Co., 1515 Worthington Ave., Clairton, PA 15025, (800) 245-1656.
(16) Stainless steel valve. Nupro valve (part SS-4BG-TW) which HPS buys and welds 1.33 in. Conflat-type flanges (HPS # 100881023) on each side and sells for $229. HPS, Division of MKS Instruments, 5330 Sterling Drive, Boulder, CO 80301, (303) 449-9861. These valves have a tendency to leak unless considerable torque has been used to tighten the bonnet seal.
(17) Copper Adapter #953-0706 (copper pinch-off tube as alternative to valve), $131 for pkg. of 4, Varian Vacuum Products, 121 Hartwell Ave., Lexington, MA 02173-9856, (800) 882-7426.
(18) Pyrex plate window #Q720125, 50.8×50.8×3.18 mm, $12 each, ESCO Products, Inc., 171 Oak Ridge Rd., Oak Ridge, NJ 07438-0155, (201) 697-3700.
(19) Torr Seal Epoxy (low vapor pressure epoxy) #953-0001, $41, Varian Vacuum Products (see item 17).
(20) Transmissive IR Viewing Card Model #Q-11-T, $79.20, Quantex, 2 Research Ct., Rockville, MD 20850, (301) 258-2701. A transmissive (as opposed to reflective) viewing card is particularly useful for aligning laser beams.
(21) Photodiode #PIN-10 DP, $55.00, United Detector Technology, 12525 Chadron Avenue, Hawthorne, CA 90250, (310) 978-0516.

[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).
[2]S. L. Gilbert and C. E. Wieman, "Laser cooling and trapping for the masses," Optics and Photonics News **4**, No. 7, 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 (1990); 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. Hänsch 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-Tannoudji, "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, and S. Chu, "Optical molasses and multilevel atoms: Theory," J. Opt. Soc. Am. B **6**, 2058–2072 (1989).

[8]L. Orozco, SUNY at Stonybrook (private communication).

[9]Prevent the shielding from coming into direct contact with the magnet. This can be done by placing small cardboard spacers between the shield and the magnet.

[10]M. Stephens and C. Wieman, "High collection efficiency in a laser trap," Phys. Rev. Lett. **72**, 3787–3790 (1994); M. Stephens, R. Rhodes and C. Wieman, "A study of wall coatings for vapor-cell laser traps," J. Appl. Phys. **76**, 3479–3488 (1994).

[11]For a discussion of saturated transition rates, see W. Demtröder, *Laser Spectroscopy* (Springer, New York, 1982), pp. 105–106.

[12]R. W. Fox, S. L. Gilbert, L. Hollberg, J. H. Marquardt, and H. G. Robinson, "Optical probing of cold trapped atoms," Opt. Lett. **18**, 1456–1458 (1993); A. G. Sinclair, B. D. McDonald, E. Riis, and G. Duxbury, "Double resonance spectroscopy of laser-cooled Rb atoms," Opt. Commun. **106**, 207–212 (1994); T. T. Grove, V. Sanchez-Villicana, B. C. Duncan, S. Maleki, and P. L. Gould, "Two-photon two-color diode laser spectroscopy of the Rb $5D_{5/2}$ state," Phys. Scr.

[13]C. Monroe, H. Robinson, and C. Wieman, "Observation of the cesium clock transition using laser cooled atoms in a vapor cell," Opt. Lett. **16**, 50–52 (1991). This work was done using trapped cesium atoms.

© 1995 American Association of Physics Teachers