

Simplified System for Creating a Bose–Einstein Condensate

H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell

*JILA, National Institute of Standards and Technology and University of Colorado and
Department of Physics, University of Colorado, Boulder, Colorado 80309–0440*

E-mail: lewandoh@colorado.edu

(Received December 23, 2002)

We designed and constructed a simplified experimental system to create a Bose–Einstein condensate in ^{87}Rb . Our system has several novel features including a mechanical atom transfer mechanism and a hybrid Ioffe–Pritchard magnetic trap. The apparatus has been designed to consistently produce a stable condensate even when it is not well optimized.

KEY WORDS: Bose–Einstein condensation; experimental apparatus; magnetic trap; degenerate gas.

1. INTRODUCTION

In the seven years since their first observation, dilute vapor Bose–Einstein condensates (BECs) have been studied extensively. In most cases the condensate properties themselves are the focus of the investigations. Relatively little work has been done using a condensate as a tool to explore questions in other fields. We feel that physicists in other fields such as condensed matter have a different and valuable perspective on possible experiments that could make use of condensates. Therefore, developing a system that could be used for these purposes is worthwhile. The current experimental systems^{1–3} were designed by people with a tremendous amount of knowledge and experience in experimental atomic physics, and until now producing a BEC without expertise in ultracold atom trapping has been a daunting task. We felt, however, that with some modifications to the current experimental design, and with a “cookbook” set of instructions, any experimental physicist, regardless of discipline, could produce a BEC in their lab.⁴

We will describe most of the basic steps in detail on how to build a BEC apparatus. Many of the techniques described have been developed by

others⁵⁻⁷ over the last 20 years but are included so that this paper may serve as a “recipe” to create a BEC. Besides the traditional methods we also describe several new features in our design, which include a new method of atom transfer and a hybrid magnetic trap.

Our design has many features that allow a condensate to be made even if the system is not particularly optimized. We use a magneto-optical trap (MOT) with large beams with a considerable amount of laser power to collect a large number of atoms. The large number of atoms in our MOT is efficiently and reproducibly translated 45 cm to a final magnetic trap via a mechanical transfer mechanism. We are able to place the magnetic trapping magnets and coils very close to our atoms thus producing very strong confinement and high collision rate. All of these features insure that evaporation will work well and produce a condensate consistently.

One of the main concerns when designing any BEC apparatus is the need to optically collect many atoms and yet to have a long lifetime for the atoms in the magnetic trap. These two constraints require orders of magnitude different vapor pressures. It is not easy to change the vapor pressure in a vacuum system by two orders of magnitude in a reasonable time of seconds. There are two conventional ways of solving this problem: a double MOT apparatus⁸ and an atomic beam.⁹ The double MOT system collects atoms in a MOT in one region of the vacuum system, which has a high alkali vapor pressure, and transfers the pre-cooled atoms to a second MOT in an ultra-high vacuum (UHV) chamber, with a pressure in the low 10^{-11} torr, for further cooling. The transfer is done by essentially pushing the atoms between MOTs with a laser beam. There are several disadvantages to this method. Two MOTs are necessary, which requires the system to use more laser power than a single MOT system. Also the optics and optical access needed for the second MOT restrict the space for the magnetic trap coils, so that to create a specified magnetic gradient requires considerable electrical power. Finally, the push beam and MOTs are very sensitive to optical alignment, making optimum performance difficult to sustain. The other main method is to use a laser-cooled atomic beam from a Zeeman slower. Zeeman slowers are large (1–2 m) and require a high temperature oven. Systems with Zeeman slowers have only one MOT, but they suffer from the same optical access problems around the magnetic trap as does the double MOT design.

A design similar to ours, developed by the Hänsch group in Garching,¹⁰ uses a series of electromagnetic coils to move the atoms from one chamber to the other. The atoms are initially transferred from the MOT into a magnetic quadrupole trap; then by ramping the current in successive sets of quadrupole coils lined along the transfer tube, the atoms are transferred between chambers. This design has the benefit of being able

to move the atoms easily around a corner to reduce line-of-sight between the chambers, which reduces trap loss from background collisions. However, the ten overlapping sets of quadrupole coils require a large amount of power to run, take up considerable space around the apparatus, and require effort to design, construct, and optimize.

Our system uses a moving magnetic trap to transfer atoms between regions of the vacuum system. Magnetic coils are mounted on a linear track and translated from one section of the system to the other. This technique has many advantages. It is very easy to use. The track, motor, and controller are commercially available as a turnkey system. The transfer worked the first time and works essentially every time without any maintenance. As with the Garching group's design, our system does not need a second MOT in the UHV region, so we can place our magnetic trap close to the atoms and produce very strong confining fields.

The other new feature in our system is a hybrid magnetic trap. It uses strong permanent magnets to produce radial confining fields and low power electromagnetic coils to produce axial confinement and a bias field. The permanent magnets do not consume power and thus do not need to be actively cooled, as would electromagnetic coils producing the same field. The tight confinement from the permanent magnets insures that we have the necessary collision rate to evaporatively cool the atoms. Permanent magnets are however plagued by long-term stability problems associated with temperature induced field drifts. The longitudinal bias field, the only parameter for which stability is critical, is produced by servo-controlled electromagnetic coils, which are air cooled. Our design allows tight confinement from the permanent magnets and bias-field stability from the coils. The magnetic trap is designed to have tight enough confinement so that we can produce a condensate even if the system is not well optimized (see Sec. 11).

Several other groups around the world have recently unveiled novel condensate-producing technologies, and some of these may in the future prove to be simpler to implement. One such system involves a surface MOT and a magnetic trap based on a wafer with lithographically patterned wires, an "atom chip."¹¹ In our opinion this technology is compact, generates condensates with unprecedented rapidity, and holds promise for eventually being simpler and more robust than traditional condensate machines, but for now it, if anything, requires more expertise to fabricate. A second novel approach is the all optical method championed by the Chapman group¹² and the Thomas group.¹³ By removing the need for two separate vacuum regions, and eliminating the magnetic trap altogether, this method indeed may eventually become the simplest route to BEC. On the other hand, in the intervening two years since this method was first demonstrated

to cool atoms to degeneracy, a number very experienced groups have had difficulty in implementing it. The all-optical method may be more difficult than it sounds.

We will outline in detail the steps required to make a Bose–Einstein condensate using our experimental system. This text¹⁴ begins by giving directions for setting up the necessary vacuum and optical systems. Next we describe the experimental procedure for pre-cooling atoms in a MOT, transferring the atoms to a magnetic trap, and evaporatively cooling them to create a condensate. There is also an extensive section dedicated to imaging the condensate and extracting useful parameters from the images. Attached are some useful appendices that should serve as a reference for working with rubidium and purchasing the parts necessary to construct a BEC apparatus.

In choosing prior publications to cite, we have often preferred to identify useful reference works rather than to trace the history¹⁵ of experimental developments. There are several papers which we particularly recommend for background information on a variety of subjects not covered in depth in this text. For a textbook description of atom cooling and trapping we recommend *Laser Cooling and Trapping* by Metcalf and van der Straten.¹⁶ Two useful papers that describe diode lasers and saturated absorption spectroscopy are “Using diode lasers for atomic physics”¹⁷ and “A narrow-band tunable diode laser system with grating feedback and a saturated absorption spectrometer for cesium and rubidium.”¹⁸ A third paper by Wieman and co-workers describes the components and the process for constructing a MOT for use in undergraduate laboratories.¹⁹ A review of many ideas in evaporative cooling may be found in a paper by van Druten and Ketterle.²⁰ There are good sections on imaging cold atoms and on magnetic trap design in “Making, probing, and understanding Bose–Einstein condensates” by Ketterle and colleagues.²¹ Also an informal overview and timeline of the experimental steps can be requested from ecornell@jilau1.colorado.edu.

2. EXPERIMENTAL RESOURCES

Our experimental apparatus does not require a large amount of real estate compared to other BEC systems, particularly those with Zeeman slower. We have the vacuum system, lasers, and all of the optics on one 122 cm × 244 cm × 30.5 cm optical table. The table is not congested, with a third of a square meter of free space remaining. In addition to the optical table, we have an overhanging canopy the size of the table, which is half full of electronics, and one full free-standing electronics rack. Experimental

control and data acquisition are provided by two computers, which sit on a separate table. A parts list for the apparatus is included in Appendix B. Very roughly, we estimate the cost of building a similar apparatus is \$200,000 with an additional \$75,000 for the three lasers. These figures are in 2002 dollars and would vary enormously depending on how many of the components are homemade, the cost of shop time, etc. The cost of labor (graduate students and postdocs) is not included.

3. LASERS

The experiment requires three different wavelengths of laser light: two to make the magneto-optical trap and one to image the condensate. We use a different laser for each of these tasks. For the MOT's trapping beams we use a commercial external cavity diode laser (ECDL), which is amplified by a single pass through a tapered amplifier chip, in a master-oscillator power-amplifier (MOPA) configuration. This system will nominally produce 500 mW of power at 780 nm. The probe beam and the MOT's hyperfine repump beam are supplied by two separate ECDLs, which each produce ~ 8 mW.

For eleven years now our group has used diode lasers exclusively for our Rubidium trapping and cooling experiments. We like the low cost, the relative ease of use, and the fact that once they are properly set up they require very little attention compared to ring lasers. In the early 90s, we built all our own diode laser systems,^{17, 18, 22} but in the mid 90s we began to replace our home-built systems with commercial systems, which we found worked very well and saved us a lot of effort. Unfortunately, during the late 90s several companies abandoned the scientific diode laser market altogether, and other companies shipped markedly lower quality systems to their scientific customers. It is not yet clear whether there will be a restoration of the availability of high-quality diode laser systems to the level of the mid 1990s. The main alternative technology is Ti-Sapphire ring lasers, which generate plenty of power at the Rubidium wavelength, but which require more money to purchase and more skill to operate. On the other hand, Ti-Sapphire lasers are readily available and their quality has only improved over the last decade.

3.1. Frequency Control and Stabilization

All three lasers are locked to atomic transitions in ^{87}Rb using saturated absorption spectroscopy. This type of frequency stabilization is discussed elsewhere¹⁸ and thus will not be described here in detail. The basic




Fig. 1. Optical set-up for saturated absorption spectroscopy. A small amount of light (0.1–1 mW) is split off from the laser and sent through a 12 mm thick piece of clear plastic to produce two weak beams. The two reflected beams, about 1 to 2 mm wide, pass through the Rb vapor cell and are focused onto two photodiodes. The signals from the photodiodes are subtracted to remove the broad doppler profile then sent through a lock-in amplifier. The saturating beam, which passes through the beamsplitter is sent into the vapor cell counter-propagating with one of the weak beams.

idea is to produce sub-Doppler spectral lines, which can be used as feedback to stabilize the laser. The optical layout for saturated absorption spectroscopy is shown in Fig. 1.

We lock each laser to the peak of an atomic transition. The frequency location of the peak of the transition is relatively insensitive to intensity and broadening effects, which would change the locking set point if the laser were locked instead to the side of the line. Unfortunately a servo can only lock to a region where there is a slope of the line to use as feedback. The standard solution to this problem is to generate a derivative of the saturated absorption signal. We modulate the frequency of the laser, by modulating either the electrical current driving the laser or the radio frequency (rf) driving an acoustic optic modulator (AOM).²³ The modulation has a depth of 5 MHz at a rate of 50 kHz, which is slow enough for the AOMs to respond and fast enough to be above the bandwidth of the servo. The signal from the saturated absorption spectrometer is routed to a homemade lock-in detector,²⁴ which gives the derivative of the original transition lines. The derivative changes sign at the absorption peak, and thus when compared to a zero-volt reference, is a convenient error signal for our servo.

We use AOMs²³ to offset the frequency of the light used in the experiment from the light sent to the saturated absorption spectrometer. The AOMs are driven by the amplified output of voltage-controlled oscillators. A diagram of the optical set-up for the AOMs is shown in Fig. 2. A portion of the trapping beam from the MOPA is sent through a 120 MHz AOM, which is frequency modulated. The negative first-order diffracted beam is

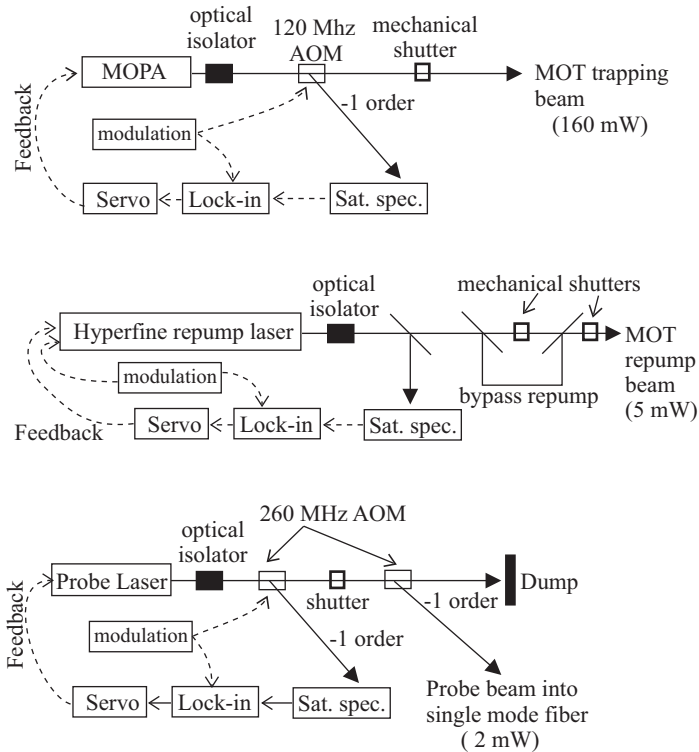


Fig. 2. Laser frequency control for the three lasers in our experiment. Laser frequency control for the three lasers in our experiment. For the less critical repump laser, the frequency modulation for the lock-in detection is applied to the laser frequency itself through current modulation. For the more critical trapping and probing beams, the frequency modulation is applied to the AOMs, through an rf amplifier, and thus the frequency modulation is not on the light used sent to the atoms. Optical isolators are placed at the output of the ECDLs to reduce optical feedback, which can cause frequency noise on the sensitive diode lasers. An optical isolator is also placed at the output of the MOPA. Light reflected back into the MOPA system can cause damage to the amplifier chip.

used in the saturated absorption spectrometer. This scheme allows us to lock the laser to the peak of the ($F = 2 \rightarrow F' = 2$, $F = 2 \rightarrow F' = 3$) crossover saturation line (peak A in Fig. 3) and have the trapping light red detuned by several natural linewidths from the $F = 2 \rightarrow F' = 3$ cooling transition. The unprimed states refer to the $5S_{1/2}$ manifold and the primed states refer to the $5P_{3/2}$ manifold. The repump laser is locked directly to the $F = 1 \rightarrow F' = 2$ transition. We modulate the current of the repump laser to

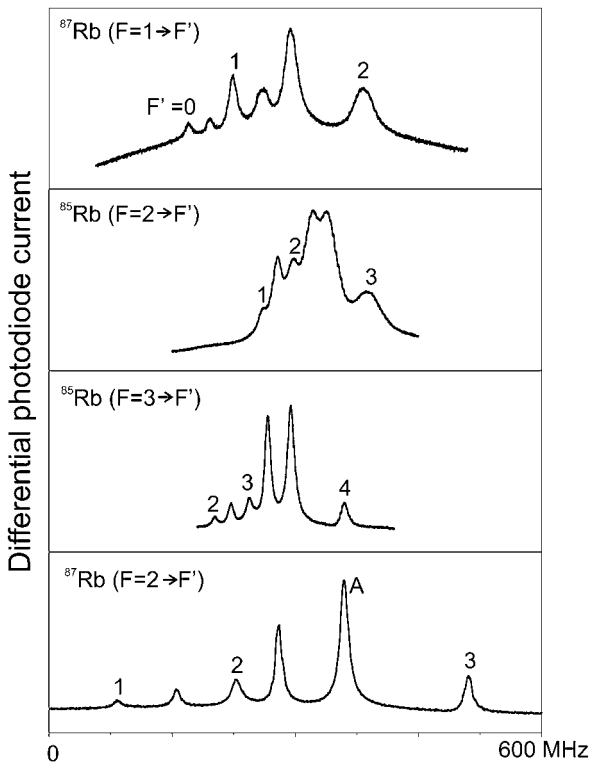


Fig. 3. Saturated absorption spectra of the hyperfine structure on the rubidium $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Widths and relative heights of the peaks are affected by beam alignment, intensity, polarization, and ambient magnetic field. This is the signal we would see from the saturated absorption set up (Fig. 1) if a linear ramp with no rapid modulation was applied to the laser frequency.

produce the locking dither instead of the frequency of an AOM. The probe beam is sent through two AOMs in our design. One AOM is essentially a fast shutter, while the other allows us to modulate the frequency sent to the saturated absorption without imparting the frequency modulation onto the probe beam itself. We need to shift the frequency of the probe beam only a few MHz from its lock point at the peak of the ($F = 2 \rightarrow F' = 2$, $F = 2 \rightarrow F' = 3$) crossover transition. Therefore the two AOMs are set so that their difference frequency is this few MHz, and their center frequency of 260 MHz is arbitrary. For instance, if 80 MHz AOMs are less expensive or more readily available, they will work just as well.

Vibration isolation is important when operating ECDLs, in our case commercial New Focus Vortex lasers.²⁶ Vibration can cause frequency noise at a level that the feedback may not be able to suppress fully. We mount the repump laser on a piece of 6 mm thick sorbathane sheeting to reduce the effect of vibrations from the table, which is in additionally isolated from the floor by air bladders in the table legs. The probe beam laser, on the other hand, is mounted directly on the table because mounting the probe laser on sorbathane could cause the beam pointing to drift and thus decrease the coupling into the fiber we use for spatial filtering. We also mount our mechanical shutters on sorbathane, so that vibrations induced by the solenoids when they open or close are not transmitted to the optical table.

Another concern with the operation ECDLs is electrical ground loops, which can cause noise on the laser. All of the electronics used for the laser and frequency locking should use the same electrical outlet.

3.2. Shutters

We use both mechanical shutters and AOMs to control the timing of the laser light. The mechanical shutters provide an excellent extinction ratio but are slow on the order of 1 ms and can have timing jitter of a few milliseconds. We increase the effective speed of the shutters by placing them at foci of the beams. AOMs are fast, with response times on the order of a 1 μ s, but their extinction ratio is less impressive. For critical light pulses such as the probe pulse, we use an AOM and a mechanical shutter in series. Vibrations associated with the opening or the closing of a mechanical shutter can induce transient noise on the lasers. We are careful to open the mechanical shutter several milliseconds before triggering the AOM for a probe pulse, in order to allow the laser servo time to recover from the transient.

3.3. Spatial Filtering

The spatial beam profile from diode lasers is not Gaussian and may contain high contrast stripes, which must be smoothed before the light is used in the experiment. Depending on the quality of its amplifier chip, the output of a MOPA system can have still worse uniformity than the output directly from a diode laser. While a MOT does not require perfectly uniform beams, high-contrast spatial structure will lower the trapping efficiency. The spatial quality of a beam can be determined by translating a pinhole across the beam and measuring the intensity at each point on a photodiode.

We spatially filter the trapping beam by sending it through a single mode fiber, which reduces the power by about 50%. When we first constructed the experiment we focused the trapping beam, $f/\#$ 16 (i.e., the diameter of the beam is 16 times smaller than the focal length), through a large pinhole (50–100 μm), eliminating the higher-frequency spatial modes while retaining most of the laser power ($\sim 75\%$). We were able to make a factor of 2 larger condensates using the pinhole to filter the light. However, we choose to use a fiber to filter spatially the trapping beam because it reduced our shot-to-shot condensate number variation from 10 to 5%, and reduced the need to adjust the trapping beam alignment from once a day to one every other month. The MOT is largely insensitive to the spatial profile of the repump beam, so we do not spatially filter this light.

The probe beam, on the other hand, must have a very uniform intensity profile. We use a single-mode angle-polished polarization-preserving fiber to filter spatially the probe beam. It is critical to use both an angle-polished and polarization-preserving fiber to reduce temporal intensity fluctuations of the beam. The input and output facets of a flat-polished fiber can form an etalon, which will produce high-frequency intensity fluctuations on the output. A non-polarization-preserving fiber will scramble the input polarization depending on the stress (thermal or mechanical) applied to the length of the fiber. These polarization fluctuations will be converted into intensity fluctuations when the light passes through a polarizer. One good method to align the input light polarization with the axis of the fiber is to first adjust the angle of the initial linear polarization until pure linear light is emitted from the fiber. This procedure may be too coarse to accurately align the axis, so the next step is to tap on the center of the fiber, so as to not change the coupling at the ends or warm it with one's hand, and to watch the transmitted intensity fluctuations on the output after a polarizer. One can now more finely adjust the input polarization until a minimum of polarization fluctuations is observed on the output.

4. VACUUM SYSTEM

4.1. Chamber Design

The vacuum system is comprised of a high vacuum MOT cell (10^{-8} – 10^{-9} torr) and an UHV (10^{-11} torr) science cell. Three vacuum pumps are used in the system, but only one of them is used on a continual basis. A diagram of the system is shown in Fig. 4. The turbo pump, connected to the system by an all-metal valve, is used only during initial pump down and bake out. The Ti-sublimation pump is turned on only every couple of years to remove extra Rb or H from the system. The workhorse

pump is the 40 l/s ion pump, which pumps continuously on the system during normal operations. The pumping speed is conduction limited for some atomic species and thus the ion pump may be larger than needed. However we would recommend the 40 l/s pump to ensure the system has adequate pumping speed. The valve on the ion pump was used only as a diagnostic tool during original testing of the system and would be removed if the system were reconstructed. After bakeout, the turbo pump is valved off and shut down, which improves the ultimate pressure and minimizes vibrations. The sealing surface of the valves and not the bellows should always face the vacuum side; this configuration reduces the surface area in the UHV system.




Fig. 4. Vacuum system layout (top view). The system is suspended with a series of clamps (not shown), so the center line of all the horizontal tubing is 18 cm above the optical table. The MOT cell is supported on the end containing the getters by resting it on a support. The science cell is only supported by the attachment to the flange.

An important consideration in Ti-sublimation pump placement is where the titanium will deposit. The Ti-sublimation pump's filaments should be placed so that there is not a direct line of sight to any valve or pump. The titanium will coat any surfaces with a direct line of sight to the filaments, and this can cause a valve sealing or pumping problem. See Fig. 4 for position of the filaments in our system.

The differential pressure between the two chambers is maintained by placing a small aperture on each side of the bellows to reduce the conductance (Fig. 4). The apertures are 5 mm diameter holes in the oxygen-free solid copper gaskets in the flanges joining the bellows to the system. The diameter of the apertures is not a design parameter that should be modified without careful thought. The aperture diameter of 5 mm was chosen because it allows most of the atoms in the quadrupole trap through when the cloud has a temperature of 200 μK and yet limits the conductance enough to have an adequate pressure in the UHV region. The pressure differential between the two chambers is about a factor of 17. A second ion pump could be used to pump the volume between the two apertures and thus increasing the pressure differential, allowing a shorter MOT loading time without sacrificing science cell lifetime.

The conductance of gas through a tube and an aperture in the molecular-flow regime (i.e., mean free path of a particle is greater than the tube diameter) is

$$C_{\text{tube}} = \frac{3.81D^3}{L} \sqrt{T/MW} \text{ liters sec}^{-1} \quad (1)$$

$$C_{\text{aperture}} = 3.64 A \sqrt{T/MW} \text{ liters sec}^{-1}, \quad (2)$$

where D and L are the diameter and length of the tube in cm.²⁷ A is the cross-sectional area of the aperture in cm^2 , T is the temperature in Kelvin, and MW is the molecular weight of the gas in atomic mass units. The conductance of a system can be found by adding the conductance of each individual part like capacitors in an electrical circuit. ($C_{\text{parallel}} = C_1 + C_2 + \dots$; $1/C_{\text{series}} = 1/C_1 + 1/C_2 + \dots$).

For pumping rubidium, these formulae are not particularly relevant, except perhaps to describe *relative* pumping speeds of different elements in a vacuum system. At room temperature, rubidium atoms adhere essentially each time they collide with a surface, and then remain on the surface for a variable length of time having to do with the material and with degree of existing surface coverage. As a result pumping speeds can be so slow that one can observe rubidium partial pressure differences of six orders of magnitude at different locations within a typical vacuum system. For this reason, ion gauges and residual gas analyzers are seldom very useful.




Fig. 5. Diagram of the MOT cell showing the getter assembly and glass-to-metal seal welded onto a knife-edge seal flange.

The science and MOT cells are cylindrical glass cells attached to glass-to-metal seals. Quartz cells are more permeable than Pyrex to atmospheric helium and should be avoided in UHV regions of the system. The science cell is a 10 cm length of 1.4 cm outside diameter 1.3 mm thick pyrex tubing with a window on one end and a glass-to-metal seal on the other. The MOT cell is a 25 cm length of 5 cm outside diameter 1.9 mm thick pyrex tubing necked down on each end. One end is attached to a glass-to-metal seal, while a getter assembly is fused to the other end (Fig. 5).

The getter assembly consists of a current feedthrough and two Rb dispensers (getters). The feedthrough, called a pin press, is a glass section with Tungsten pins inserted. These items are commercially available. The Rb getters are spot welded to the pins of the pin press. The Rb getters are a controllable source of Rb vapor. A getter is a small foil container of a Rb salt, which releases Rb when a moderate current of 2 to 6 A is run through the device. The getter assembly can be seen in Fig. 6.

Special care must be taken with the getters to insure they will produce clean Rb vapor. First, the getter material can easily absorb water, so we

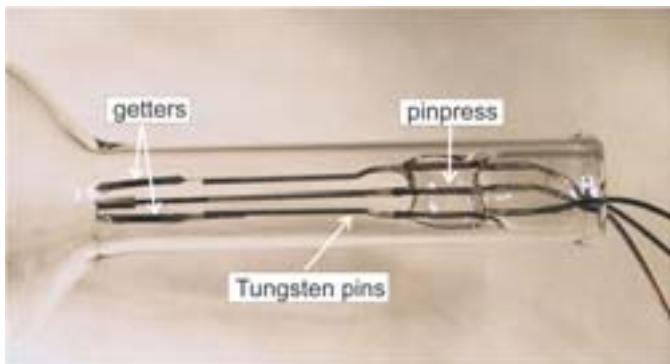


Fig. 6. Getter assembly showing two getters and reentrant glass with pins fed through.

store them under vacuum with desiccant and flow dry gas during the glass fusing process. Second, they release Rb as a double exponential function of temperature. Thus, we make sure that while the pin press is being fused to the cell, the getters are not heated significantly by the fusing torch. Avoiding moisture and heat, we are usually able to install getters that produce fairly clean Rb vapor. In our current system when we turn on the getters the number of atoms in the MOT decreases, presumably due to contaminants being released from the getter. Our mode of operation is to turn on the getters at 3.5 A for 10 minutes to supply the MOT cell with a day's worth of Rb, and then to allow 10 to 20 min for the contaminants to pump out of the system before taking data. One getter in our system has been used in this manner each day for over 4 years without any sign of reduced production of rubidium. Getters that are less contaminated can be run continuously throughout the day at a lower current.

One is aiming to have a partial pressure of Rubidium of something less than 10^{-9} torr and a partial pressure of all impurity gases lower than the Rubidium pressure by at least a factor of two. Note that because of the sticky nature of rubidium, its pumping speed is extremely low, so the pressure read for instance on the ion-pump controller current will have little to do with the rubidium pressure in the MOT cell. Rubidium pressure can be determined locally by looking at absorption on a beam through the cell, but pressures are best understood and measured in terms of inverse lifetimes of trapped atoms. One would like the lifetime in the MOT cell to be about 5 to 10 seconds, and in the science cell to be in excess of 100 seconds.

It is worth discussing why we chose to use cylindrical glass cells instead of square cross-section cells. One reason is it is easier to make cylindrical cells. Anyone with a small amount of glassblowing training can fabricate the cylindrical cells, whereas constructing square cells usually requires significantly more equipment and expertise. An equally important advantage to using a cylindrical MOT cell, however, is that the interference fringes on the trapping beams have higher spatial frequency and are spaced less regularly than with a square cell. These relatively fine structure intensity fringes have little effect on the trapped atoms. One can shake the glass MOT cell around by as much as 1 cm and see little movement of the MOT cloud. MOT alignment with a square cell can be more difficult because it is important to place the minimum of the broad intensity fringes away from the center of the trap. This type of alignment requires more frequent adjustments and is more a trial-and-error process than simply overlapping the beams at the correct angles, which is all that is required for a cylindrical cell. A cylindrical cell does distort the trapping beam, but this is not a large problem for us as we do not retro reflect our trapping beams and the diameter of the cylinder is large.

The cylindrical science cell, on the other hand, is less desirable. The probe beam is focused by the cell, which acts as two cylindrical lenses. This is generally not a problem for absorption imaging, but can be for phase contrast imaging. Phase contrast imaging requires placing a material in the fourier plane to shift the phase of the light. The large astigmatism induced by the cylindrical cell requires the use of a phase shifting line rather than a dot, thus making alignment more difficult. If we were to reconstruct our experiment, we would replace the cylindrical science cell with a commercial square cell, which can be obtained from companies that specialize in producing spectroscopic cells. Interference fringes generated by reflections off the uncoated walls of the science cell are relatively unimportant, as they are far from the object plane of the imaging system.

4.2. Chamber Construction

Obtaining UHV pressures requires careful assembly of the vacuum components. The most important thing is to make sure all of the components are clean. We start the cleaning process by placing the submersible parts (no valves, pumps, or cells) in an ultrasonic cleaner with strong soap for 1 hour. If the valves are cleaned in the ultrasonic cleaner they must be re-greased before they are used. The long cleaning time allows the strong soap to remove residual oil from the factory. Typically when a stainless steel vacuum part is baked in air it will become a golden color, which we assume is residual burnt factory oil. A one hour bath in a strong basic soap will remove this coating, and it will not return with subsequent air bakes. The long bath is not absolutely necessary and may be reduced to a few minutes to just remove any particulates from the parts. After the ultrasonic bath the parts are rinsed first with deionized water, then acetone, and finally spectroscopic-grade methanol. Next the parts are baked in air for 4 hours at 400°C to drive off any residual solvents. Once the parts have cooled they are wrapped in *oil-free* aluminum foil until assembly.

It is important to avoid contamination of the vacuum system during assembly. We always wear powder-free latex gloves and change them often. All copper gaskets are wiped with ultra-pure methanol before installation to remove any factory residue.

We use silvered bolts on the knife edge flanges to reduce the possibility of bolts seizing in the flanges during the bake out. If silvered bolts are not available we place some molybdenum disulfide powder suspended in methanol on the threads of the bolts for lubrication. Suspending the lubricant in methanol reduces the chance that it will fall into the vacuum system during assembly and become a contaminant.

After the entire vacuum system has been assembled it is pumped out and checked for leaks. We use a small turbo pump backed by a dry, four-stage diaphragm pump to initially pump out the system. We use a diaphragm pump rather than a standard oil-filled roughing pump, because the diaphragm pump does not contain any oil, which could backflow into the system. Once the turbo pump has spun up to full speed we spray a small amount of spectroscopic-grade methanol on all the flanges and cells. If there is a large leak, the pressure in the tubing connecting the turbo pump to the diaphragm pump, read by a thermocouple gauge, will change when methanol is applied. We avoid using commercial leak detector apparatuses, as they are frequently contaminated with heavy hydrocarbons, which can backflow into our clean system. The system is pumped on overnight before the bake out is started. A carefully cleaned, leak-free system should pump out overnight with the pressure reaching around several 10^{-8} torr, read from the ion pump current. We turn on the ion pump briefly to determine the pressure in the system. The ion pump will not be turned on to pump for extended periods of time until the bake is underway.

4.3. Chamber Bake Out

The vacuum system must be baked at high temperatures under vacuum to remove contaminants to obtain UHV pressures. We bake most of the vacuum system at 300°C for several days. Before the bake out, we run about 6 A through each getter for 30 seconds to verify the presence of Rb, which can be seen by either laser absorption or fluorescence. The high current also degasses the getters. It is important to not run the getters for more than several minutes at 6 A or all of the rubidium contained maybe be released from the getter. We mount the system loosely to the optical table such that, when the system expands, there is minimal stress on the system's joints. The thermal expansion may cause enough torque to cause the flanges to leak. We also place microscope slides under the mounts to decrease thermal contact between the system and the optical table.

The system can now be prepared for the bake. The first step in the bake out process is to wrap the glass cells with clean fiberglass cloth. The cloth will protect the cells from anything that may melt onto the cells during the bake. Next we place thermocouples on the vacuum system at critical places such as the cells, glass-to-metal seals, and pumps. We then wrap the system with resistive heater tape. The aim in wrapping the heater tape is not to cover the entire surface of the vacuum system with tape, but rather to have a constant tape-to-chamber surface area ratio. Heater tape is applied to different objects proportional to their surface area and not the mass of the object. The mass only defines the time constant for thermal

TABLE I
Vacuum Component Temperature Limits

Vacuum element	Maximum baking temperature (°C)
Turbo pump inlet flange	120
Ion pump magnets	350
Ion pump body	400
Ion pump cable	250
Bakeable valve, open	450
Bakeable valve, closed	300
Ti-sublimation pump	350
Glass/metal seals	300

equilibration, whereas the ultimate temperature is determined by the heat flow in and out of the region, which is proportional to the surface area. The tape should never overlap itself, or the intense heat will cause the tape to burn. Several short tapes are used to wrap the system so each section may be controlled independently. The turbo pump is not baked because it is not part of the final system and can not handle high temperatures. We do however bake the entire ion pump with the magnets in place. Typical maximum baking temperatures for different components are listed in Table I. After the tapes are in place, the system is wrapped loosely with strips of fiberglass tape and then aluminum foil to provide thermal insulation.

It is tempting to bake the main chambers to less than 300° to eliminate any chance of breaking a glass cell. This precaution could cost more time than replacing a broken cell. It could take several weeks to make a MOT, transfer atoms into the quadrupole trap and determine that the vacuum pressure is not adequate because the system was not baked at a high enough temperature. On the other hand, replacing a broken cell and rebaking the system will usually take only one week.

The system is slowly brought up to the final temperature over 6 to 8 hours. The ion pump is off during the warm up. There is a large amount of material driven off the walls of the vacuum system during the initial several hours of the bake. We prefer to have the turbo pump remove the bulk of the material rather than the finite-lifetime ion pump. Each heater tape is powered by a variable AC transformer (Variac) to adjust the temperature of the corresponding section of the system. Generally we increase the temperature by at most 50°C per hour. Temperature gradients can apply significant stress to the system. We prefer to keep the temperature gradients to under 30°C across the glass cells and glass-to-metal seals, which are the most susceptible components to failure. Caution must be taken as the

system approaches its final bake temperature because some parts of the system could overshoot in temperature due to long thermal time constants. During the bake 3 A are run continuously through the getters to clean them. If the getters are left off they will be the coldest part of the system because of the thermal conduction through the leads, and contaminants will accumulate on them. We also run 25 A through the Ti-sublimation filaments during most of the bake. Throughout the warming up process, thermocouple readings, Variac settings, and pressure readings are recorded to facilitate future bakes. Once the system is at the desired temperature we bake with just the turbo pump on for 12 hours. At this point we degas the getters and Ti-sublimation filaments. To degas the getters, we increase the current in each getter for 30 seconds to 5 A to drive off any surface contaminants. After the degassing we turn on the ion pump and valve off the still-running turbo pump. When the valve is above room temperature we close it only finger tight. Therefore, we do not allow the turbo pump to spin down until the system is back to room temperature, and the valve has been properly closed with a torque wrench. We allow the system to bake with the ion pump on for ~ 2 days, or until the pressure on the ion pump reads in the low 10^{-8} torr. We cool the system down slowly over 4 to 6 hours. At this point the ion pump should read the lowest possible pressure, which is 10^{-10} torr for most pumps, if there is no leakage current. In our experiment, including a separate ion gauge in the system is more likely to do harm than good. The ultimate test of the vacuum pressure will be the lifetime of the atoms in a magnetic trap.

5. MAGNETO-OPTICAL TRAP

Perhaps the single most important predictor of success in an evaporative cooling experiment is the number of atoms one can collect in the MOT. The larger the MOT number, the more forgiving every other aspect (vacuum, beam alignment and balance, magnetic trap strength, etc) of the experiment becomes. 1×10^9 is good, but 5×10^9 is better. To collect a large number of atoms in a MOT, one needs basically lots of power, and large diameter trapping beams.^{28,29} At very large atom number, the trapped atoms cast such a dark shadow in the MOT beams that it is no longer wise to use three retroreflected beams, but rather one should split the MOT power into six independent beams. We designed our optical layout to support large diameter beams for this reason, and we encourage other groups, especially groups with less experience in successfully creating condensates, to do the same. With our apparatus we made our largest condensates by applying only minimal optical filtering to the trapping beam laser so as to have the most available power, and then spreading that

power out over as wide beams as we could get through our optics and into our glass cylinder. That said, once one actually has condensates in one's machine, one can afford to be a little more picky. For most of our experiments we are much more interested in maximizing condensate reproducibility than condensate size. In our current mode of operation, we use more aggressive spatial filtering, getting cleaner beams at the expense of laser power. With less power in the beams, there is less to be gained from expanding them as far spatially, so we operate with beam diameters of approximately 3 cm (FWHM = 1.6 cm). Our condensates are smaller, but the overall performance of the machine is still very robust (see Sec. 11).

The optical layout is shown in Fig. 7 and Fig. 8. We use 5 cm diameter polarizing beamsplitting cubes and waveplates and 7.5 cm mirrors to accommodate the large beams. The repump beam enters the system via the backside of a polarizing cube; therefore the polarization of the repump beam will not be optimum when it enters the trapping region. This is not a problem because a Rb MOT requires very little repump power.




Fig. 7. Schematic of the MOT optical layout (top view). The linear track and servo motor are shown in their approximate locations on the table. Some of the mirrors and waveplates for the vertical beams are not shown. There are enough degrees of freedom to adjust the position and angle of each beam.



Fig. 8. Picture of the vertical-beam optics for the MOT omitted from Fig. 7. Figure shows the location of the track (black arrow) and the vertical MOT beams (white arrows). The coils of the MOT/quadropole trap are translated left, towards the science cell, out of the field of view of this photograph.

The MOT coils (Fig. 12), which also serve as the quadrupole trap coils, are each made of 24 turns of square hollow copper tubing coated with Kapton. The wire has a square cross-section of 4.15 mm on a side with a round 2.5 mm diameter hole in the center. The coils are cooled by running water through the center region of the wire. The wire is wound onto a phenolic spool and secured with epoxy. Phenolic was chosen as the spool material because it will not support eddy currents when the current is abruptly changed in the coils. The inner diameter of the coils is 5 cm, and their centers are separated axially by 10 cm. The current in the coils, run in series, is controlled by a simple servo circuit (Fig. 9), which controls three power MOSFETs. We use a 580 A, 8 V switching power supply, run in voltage-controlled mode, to supply current to the MOT coils. We are limited to running a maximum of 250 A through the coils due to the

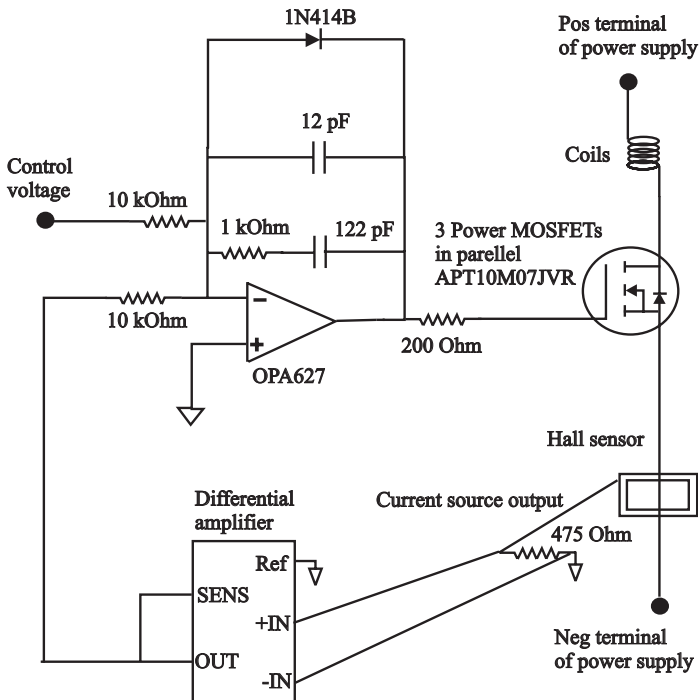


Fig. 9. Schematic of the MOT/Quadrupole trap servo circuit. The circuit has standard proportional-integral loop gain. We place a 200 Ω resistor on the gate of *each* MOSFET. We use 2/0 gauge welding cable to carry 300 A from the power supply to the coils and MOSFETs. The three MOSFETs are mounted on a water cooled copper plate.

limited voltage produced by the power supply. The coil configuration produces a magnetic field gradient of 1 Gauss/cm/A along the axis of the coils.

5.1. MOT Alignment

The alignment of our MOT is not as sensitive as it would be for a MOT with smaller beams. We start aligning the MOT by placing an iris in the trapping beam before it is split into six. Closing the aperture to a 2 mm diameter allows us to align the centers of the beams. Once we get all the beams roughly aligned with respect to the magnetic coils and each other, we balance the power in the beams. We have $\lambda/2$ plates mounted on rotating

mounts before every polarizing beamsplitting cube to adjust the power in each beam.⁴ We measure the power in each beam just before it enters the MOT cell and adjust the waveplates until the power is equal in each of the six beams to better than 10%. At this point we align the repump beam. Next we open the iris and attempt to see a trapped cloud. A simple security camera can image the fluorescence from the cloud, which can be viewed on a monitor. It may take several hours with a getter running for the cell to be coated with a monolayer or two of Rb and a significant Rb vapor pressure to be established. Once a cloud is visible, an easy way to adjust finely the alignment is to reduce the beam size, optimize the cloud for number and roundness and then iterate with ever smaller beam diameters, all the time making sure the beam pairs are kept counter-propagating. We typically get 5×10^9 atoms in our MOT with a loading rate of 8×10^8 atoms/s. The full width at half max of our MOT cloud is about 3 mm.

The position of the trapped cloud should be centered with the quadrupole magnetic trap to minimize energy gained by the cloud when it is transferred into the magnetic trap. To check this, we increase the current in the MOT coils until the cloud size is reduced greatly by the large magnetic field-induced detuning. The position of a cloud, in a very large magnetic field gradient is reliably at the null of the magnetic field, and thus the center of the magnetic trap. We then decrease the field, and adjust the beam alignment and power balance until the cloud center is in the same location at high and low magnetic fields.

Although well-optimized optical molasses³⁰ is not required for our system, optical molasses is a good diagnostic of MOT alignment. To examine the quality of the alignment we quickly turn off the magnetic field of the MOT and look at the expanding cloud. If it moves rapidly in one direction this could be a sign of beam imbalance (from incorrect polarization or intensity splitting between the beam), poor alignment, or stray magnetic fields. The goal is have slow, spatially-uniform expansion during optical molasses. We adjust the beam balance and alignment until the cloud expands fairly uniformly in the optical molasses. While the MOT is collecting atoms, the magnetic field gradient dB/dz should be set to optimize the number of atoms collected. The optimum value depends weakly on beam diameter and intensity. We use 8 G/cm.

In traditional systems, it is necessary to use shim coils to cancel ambient magnetic fields so that the atoms in the optical molasses expand uniformly. A feature of our system is the lack of need for shim coils. We transfer atoms to our quadrupole trap (see Sec. 6.1) at a relatively high temperature, where the small reduction in energy from the shim coils would not make a great improvement in the phase-space density or collision rate of the magnetically trapped cloud.

5.2. MOT Characterization

We determine the number of atoms in the MOT by imaging the fluorescence induced by the trapping lasers onto a photodiode. Some care must be taken in selecting the location of the collection lens. Ideally the line of sight from the collection lens through the glass wall, to the center of the MOT cloud, and onto the far glass wall, should not include any section of the glass wall that is illuminated by a trapping beam, as this results in too much scattered laser light hitting the photodiode. The side of the mount for our collection lens is visible in Fig. 8, near the far right end of the cylindrical glass cell. The photodiode itself is off the right edge of the photo. The photodiode is shielded by a tube of black paper so that it can “see” only the collection lens. Collecting some scattered light is unavoidable. Most of this comes from stray light scattering from imperfections in the glass cell; at our vapor pressures, essentially none of the scattered light comes from the background Rubidium vapor in the cell. The beams are not visible in the cell. In any case we subtract out the background scattered light level, which we establish by turning off the MOT magnetic coils. The number of atoms in the MOT is

$$N = \frac{4\pi(\text{photodiode current})}{(\text{solid angle})(\text{responsivity})(\text{energy of a photon})(R)(0.96)^k}, \quad (3)$$

where solid angle refers to the solid angle subtended by the collection lens, the responsivity refers to the current produced for a given power incident on the photodiode, and k to the number of uncoated glass surfaces between the atoms and the detector. R the photon scattering rate in photons/sec/atom, is

$$R = \frac{\frac{I_0}{I_s} \pi \Gamma}{1 + \frac{I_0}{I_s} + 4 \left(\frac{\Delta}{\Gamma} \right)^2}, \quad (4)$$

where I_0 is the total intensity of the six beams impinging on the atoms, I_s is the saturation intensity, which is 4.1 mW/cm^2 for random polarization for Rb. Γ is the natural linewidth of 6 MHz for Rb, and Δ is the detuning from resonance. In our experience, using the I_s appropriate for random polarization gives the most accurate number of atoms in a MOT.

We servo the MOT load to increase reproducibility in condensate number, which we do by measuring the voltage output from a photodiode collecting light from the MOT; when a desired value is reached we stop the

loading and proceed to transfer atoms to the quadrupole trap. This also allows us to vary the number of atoms in the final evaporatively cooled cloud by adjusting the initial MOT load level. An easier method of setting the number in MOT is to simply load for a set period of time, but this method can cause the number in the final cloud to drift throughout the day due to change for instance in rubidium pressure.

6. FROM MOT TO IOFFE-PRITCHARD TRAP

6.1. Transfer into the Quadrupole Magnetic Trap

There are three main steps to transferring atoms from the MOT into the quadrupole magnetic trap: compressed MOT (CMOT),^{31, 32} optical pumping, and magnetic trap turn on. Our goal is to transfer the atoms into the quadrupole trap with the highest possible phase-space density. When the atoms are caught in the magnetic trap, most of the resultant energy of the atoms comes from the added potential energy due to the Zeeman energy from the magnetic field. The larger the cloud is when the magnetic trap is turned on the greater the potential energy gained by the atoms. We can not adiabatically ramp on our magnetic trap from zero gradient, because at low magnetic gradients the trap center is significantly offset from the cloud center due to gravity. The center offset induces slosh in the trap, which turns into thermal energy. Therefore reducing the initial spatial extent of our CMOT cloud is more important than obtaining the coldest temperature in the CMOT.

Our CMOT step consists of a MOT with increased red detuning of the trapping laser and greatly reduced repump laser power. MOTs with large numbers of atoms have a maximum density of around 10^{10} atoms/cm³, which is limited by reradiation pressure. The CMOT has the effect of reducing radiation pressure in the trap and thus creating a denser cloud of atoms. Reducing the repump power reduces the time the atoms spend in the state ($F = 2$) resonant with the trapping light. Increasing the detuning of the trapping laser decreases the scattering rate and thus the absorption of reradiated photons. The CMOT stage not only reduces the overall spatial extent of the atoms in the MOT, it also cleans up much of the irregular shape.

In sodium experiments, some groups use a dark spot MOT to compress atoms from a MOT before transfer to a magnetic trap.³³ In Rubidium-87, this strategy is usually not worth the effort.

We use a short CMOT stage in preparation for transfer to the magnetic trap. The CMOT has a much smaller loading rate than a regular MOT.

Therefore we want to minimize the time spent in the CMOT stage and just go briefly to a CMOT configuration before the magnetic trap is turned on. The repump power is reduced from several mW to 50 μ W for the CMOT stage. We have two separate overlapping repump beams entering the MOT cell as shown in Fig. 7. One beam is the main repump beam with several mW of power, and the other, which we call the bypass beam, has only 50 μ W of power. Using two shutters (Fig. 2) we are able to have either full repump power or reduced power for the CMOT stage. Simultaneously with the repump power decrease, we jump the detuning of the trapping laser 50 MHz red of resonance. This frequency jump is accomplished by unlocking the laser,²⁴ applying an additional voltage to the laser piezo electric tuner during the CMOT and optical pumping stages, and then, after the shutters are closed, turning off the additional applied voltage and relocking the laser. We keep the magnetic field gradient constant at the MOT value during the CMOT stage. The entire CMOT stage lasts about 10 ms and is not very sensitive to changes in trapping laser detuning on the order of 10 MHz. The optimal CMOT parameters may be different depending on the exact experimental configuration. For instance, it is sometimes necessary to change the magnetic field gradient to optimize for the CMOT stage in the case of much smaller or larger atom numbers in the MOT.

We can characterize the atom cloud in the CMOT using fluorescence imaging. The position of the cloud in the CMOT may be very different from the position of the cloud in the MOT or the magnetic trap because of beam imbalances or misalignment. We adjust the alignment and half-wave plates controlling the power in the beams to overlap the position CMOT with that of the magnetic trap, using the same high magnetic field gradient alignment technique used for the MOT/magnetic trap alignment. A large offset of the centers will increase the temperature of the magnetically trapped cloud.

After the CMOT stage we optically pump the atoms into the lower hyperfine ground state with arbitrary population in the magnetic sublevels. An atom has a small chance of being excited to the $F' = 2$ state and decaying to the $F = 1$; typically the atoms will be pumped into the $F = 1$ state in less than one ms if the repump light is turned off. The magnetic trap will confine only the $m_f = -1$ Zeeman sublevel. One might think that we would trap only 1/3 of the $F = 1$ state atoms but, we often do better. We can, in certain circumstances, get over 50% of the atoms in the right Zeeman state, depending on the MOT beam polarization. The population distribution in a MOT is not a well controlled parameter but can be adjusted with small random tweaks of the MOT beams. We check to see that we are effectively pumping the atoms into the $F = 1$ state by attempting to take a fluorescence

TABLE II
Parameters for Trapping and Loading into a Magnetic Trap

Stage	Trapping/repump	Detuning	Magnetic gradient	Time
MOT	On/3 mW	-2.5Γ	8 G/cm	~ 10 sec
CMOT	On/50 μ W	-10Γ	8 G/cm	20 ms
Optical pumping	On/Off	-10Γ	8 G/cm	1 ms
Magnetic catch	Off/Off	—	100 G/cm	200 μ s
Magnetic trap ramp	Off/Off	—	100 \rightarrow 250 G/cm	500 ms

image of the cloud (see Sec. 6.2) with the trapping beams alone (no repump beam). If the atoms fluoresce, they have not been fully pumped into the $F = 1$ state, and the optical pumping time must be increased.

The atoms are now ready to be caught in the magnetic trap. The quadrupole magnetic trap is formed simply by turning up the current to the MOT coils (described in Sec. 5).⁶ The null in the field at the point exactly between the centers of the two MOT coils becomes the potential minimum of this simple magnetic trap. As stated before, we can not slowly ramp the magnetic field up from zero because of the effect of gravity. On the other hand, we also do not want to turn on the magnetic trap at the highest gradient possible because this will add an excess amount of energy to the cloud. Our procedure consists of diabatically turning on the magnetic trap to a point where gravity has a minimal effect and yet the magnetic trap adds as little potential energy as possible. The optimal catch point is determined empirically to be around 100 G/cm in the axial (vertical) direction. After the initial catch we adiabatically ramp the magnetic field gradient to 250 G/cm in 0.5 s. We optimize various parameters of the MOT-CMOT-quadrupole trap transfer by imaging the atoms after they have been loaded into the magnetic trap, since we ultimately care about the temperature and number of atoms in the magnetic trap. We typically get 2 to 4×10^9 atoms at 250 to 400 μ K in the fully compressed (250 G/cm) quadrupole trap.

6.2. Fluorescence Imaging

We use fluorescence imaging to characterize the cloud in the MOT region. We use fluorescence imaging because, although it has less absolute accuracy than absorption imaging, it is easy to set up and gives us the information we require. Imaging in the MOT cell is useful for a variety of diagnostics, such as loading efficiency into and temperature in the magnetic trap, transfer efficiency to the science cell, and magnetic trap lifetime in




Fig. 10. Top and side views of the optical layout for fluorescence imaging. The objective lens is apertured with an iris to 10 mm in diameter, and the second lens is 30 mm in diameter. The mirror is used to direct the light to the camera, which can not be placed in direct line of sight due to limited free space near the MOT cell.

different regions of the vacuum system. For these diagnostics it is not important to measure the absolute temperature and number of atoms in the trap but rather relative quantities. Later, when we require an accurate measure of the cloud parameters after evaporation, we will use absorption imaging, which is discussed in Sec. 8.

To capture a fluorescence image, we turn off the quadrupole trap, open the camera shutter, turn on the repump beam, flash the MOT trapping beams for less than 1 ms, and image the fluorescence from the cloud onto a CCD camera. We image the atoms directly out of the magnetic trap without allowing any additional time for expansion. The optical layout is shown in Fig. 10.

We extract parameters from the image of the cloud using a Gaussian fitting routine. The density profile is not Gaussian in a linear-potential trap such as a quadrupole trap. However, the cloud's profile is not far from Gaussian, and all we really need is a measure of the cloud that is monotonic with respect to size and fluorescence intensity. For calculational convenience, we use a Gaussian surface fit to extract the full width half max size of the cloud and then calculate the temperature and density of the cloud using Eqs. (5) and (6), which take as an input parameter the size which should be extracted from the more elaborate functional form for the projected density of atoms in a linear trap. The associated systematic error is only a few percent. The correct functional forms for the temperature T and peak density n_0 in a quadrupole trap are

$$T = \frac{2}{5} \frac{\mu_B g_f h}{k_b} B'_x \sigma_{FWHM} \quad (5)$$

$$n_0 = 1.27 \frac{N}{\sigma_{FWHM}^3}, \quad (6)$$

where g_f is the Landé g factor, μ_B is the Bohr magneton, h is Plank's constant, k_B is Boltzmann's constant, B'_x is the radial magnetic field gradient, and σ_{FWHM} is the radial full width at half max size of the cloud using the correct functional form of atoms in a spherical quadrupole trap. We calibrate the number N using the photodiode (see Sec. 5.2). Unfortunately our cylindrical glass cell causes some problems with imaging. The cell lenses the scattered light so that the size of the cloud is distorted by about 25% in the vertical direction. We get the temperature from the horizontal direction, on which the cylindrical cell has no effect.

Vignetting³⁴ is such a common imaging systematic for fluorescence imaging that it deserves to be elaborated on here. Vignetting occurs in a multiple lens system imaging an extended object, and can be a problem whenever there is more than one effective aperture in the system. For example, see the lens configuration in Fig. 11 in which some rays of light that pass through the objective lens do not make it through the second lens. The rays that do not make it through the second lens come primarily from the edge of the object as seen in Fig. 11. The decrease in imaged light from the edge of the cloud decreases the apparent size of the cloud. The larger the cloud one attempts to image, the more likely vignetting will arise.

There are a few easy ways to check if the image of a particular cloud size suffers from vignetting. First reduce the diameter of the objective lens by a factor of 2 with an iris. If vignetting is not a problem, the reduction should decrease the total intensity of the image by a factor of 4 without changing the apparent width. Alternatively one can also measure how close the system is to being affected by vignetting by reducing the size of the second lens with an iris. The size of the image will remain the same until

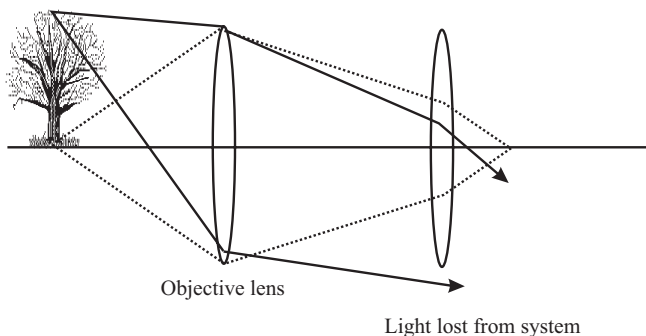


Fig. 11. Illustration of vignetting. Vignetting occurs when rays of light from the edge of an extended object are removed from the imaging system by a second aperture, in this case the second lens.

the second lens begins to become an aperture in the system. There are several ways to eliminate vignetting: replace the second lens with a larger diameter lens, aperture the objective lens, or move the lenses closer together.

6.3. Transfer from Vapor Cell to UHV Region

We use moving magnetic coils to transfer the atoms from the relatively high pressure MOT cell to the UHV region, where we evaporatively cool to BEC. The quadrupole coils are mounted on a linear stage that is driven by a servo motor and controlled by a computer (Fig. 12). The maximum possible acceleration of the coils is about 3.3 m/s^2 , which is much less than the trapping acceleration (40 m/s^2) from the magnetic trap; the atoms are therefore not heated any detectable amount. We do not see any atom loss from moving the atoms. Other similar systems have seen a loss of atoms from fringing magnetic fields from a weld in the vacuum system. It is important to avoid creating stray magnetic fields near the chamber from items such as ion pumps, magnetic bases and magnetic screws. We want to get the atoms out of the MOT cell as quickly as possible because collisions with the background gas limit the lifetime to 5 to 15 s depending on the Rb vapor pressure. However, we must slow down the coils as the trapped cloud enters the fringing fields of the permanent magnets of the Ioffe–Pritchard (IP) trap in order to adiabatically compress the cloud. We move the atoms out of the MOT cell, into the UHV region, and to within 4 cm from the center of IP trap in about 1 s. We decelerate to a speed of 1 cm/s as the atoms enter the permanent magnetic region and are adiabatically compressed.

We purchased a commercial servo-linear track to move the coils from one end of the vacuum system to the other. The track consists of a table




Fig. 12. Illustration of the motion of quadrupole trap coils from the MOT cell to the science cell. The low coil is hidden from view by the upper coil. (Top view)

mounted on a ground ball screw, which can accommodate higher speeds and has less backlash than a traditional lead screw. The servo motor in our system has reproducibility of 5 μm , which is measured by a rotary encoder in the motor housing. Although the encoder signal is sent through a shielded cable, electrical noise is radiated from the cable; typically we see 20 kHz spikes of 3 μs duration coming from the cable. So far this radiation has not caused any problems with other equipment or with our ability to make a condensate.

The ability to move the atoms in the magnetic trap to different regions of the vacuum system allows us to measure the background pressure and also to localize possible places where near-resonant stray light impinges on the system. We need long, background-gas-limited lifetimes in the UHV region to be able to efficiently evaporate and form a condensate. We do not know the lower limit on the necessary lifetime, but we do know that our 170 s lifetime is much more than sufficient. We measure the lifetime in the magnetic trap by loading atoms into the quadrupole magnetic trap, moving the cloud to the desired position in the vacuum system, waiting a variable length of time, moving the atoms back to the MOT region, and imaging the cloud. We fit an exponential to the number of atoms remaining as a function of waiting time. The exponential time constant gives us the lifetime (inversely proportional to the pressure) at various regions of the system.

Beyond collisions with background gas, there are two additional loss mechanisms that could reduce the lifetime in the quadrupole magnetic trap. One is resonant light impinging on the atoms. An atom absorbing a single photon has a large probability of falling back to an untrapped state and thus being ejected from the trap. We place a large (137 cm \times 124 cm \times 53 cm) box made from opaque plastic panels on an aluminum frame around the MOT optics and vacuum system. This has the added benefit of also blocking room lights from the experiment and thus allowing the experiment to be run with the room lights on.

The other loss mechanism is due to Majorana spin flips.³⁵ Majorana or diabatic spin flips happen in a magnetic trap only when the trap has a zero of the magnetic field. Atoms can undergo a spin flip if the time rate of change of the magnetic field is not much smaller than the Larmor frequency. In a quadrupole trap, atoms which pass through an ellipsoid near the center of the trap can be lost due to spin flips to a non-magnetically trapped state. The lifetime associated with this loss rate is proportional to the square of the size of the cloud and is given by

$$\tau = \frac{1}{4} \alpha \sigma_{FWHM}^2, \quad (7)$$

where α is determined experimentally for ^{87}Rb to be

$$\alpha = 3.7(7) \times 10^4 \frac{\text{s}}{\text{cm}^2}, \quad (8)$$

and σ_{FWHM} is the radial full width half maximum of the cloud.³⁵ The loss rate due to spin flips is much smaller than the loss rate from background gas collisions for the typical cloud temperatures (200–400 μK) we have in the quadrupole trap. If we evaporate in the quadrupole trap the size of the cloud will rapidly become small and thus the spin flip rate will become large. We must therefore evaporate in a magnetic trap without a zero of magnetic field, such as a IP trap.

6.4. Ioffe–Pritchard Magnetic Trap

We use a hybrid Ioffe–Pritchard trap, which contains both permanent magnets and electromagnetic coils. Permanent magnets are useful because they produce large magnetic field gradients with no power consumption. On the other hand, permanent magnets are sensitive to temperature fluctuations and thus can lead to instabilities if used to produce a bias field for a magnetic trap. The bias field is the trap parameter most sensitive to drift because it determines the depth of the final evaporative cut and thus the temperature. In our trap the two permanent magnets produce a quadrupole field in the radial direction but no field along the axial (or bias field) direction.

The bias field and axial confinement are created by four electromagnetic coils. The outer two coils produce essentially all of the axial curvature, and the inner two coils control the value of the bias field. Each pair of coils is run in series and controlled independently by a bipolar power supply. The power supplies internally servo the current to better than 1 part in 10^4 using an analog voltage set point supplied by a computer-controlled digital-to-analog converter. We increase the long term stability of our trap by continually running the operating current through our coils except for the 4 s period when the atoms are first transported into the IP trapping region. The IP coils are on even during the loading of the MOT. The trap is therefore always at the same temperature even if our experimental timing changes. If we run operating current through the coils continually, the coils reach a steady state temperature of 75°C . Although this temperature does not affect the operation of the magnetic trap, it does increase the temperature of the glass cell, which it surrounds. Raising the temperature of the glass cell causes an undesirable increase in background pressure. We use forced air cooling to reduce the temperature of the coils

from 75°C to 35°C. We have two air cooling ports fed by filtered compressed air as shown in Fig. 13. We do not use a fan to cool the trap because a fan's motor can generate magnetic field noise. Water cooling is another option, however water tubing takes up considerable space and flowing water can cause vibrations.

Our trap has the advantages of tight radial confinement from permanent magnets and also a stable bias field from well-servoed axial coils. Our trapping frequencies are (230, 230, 7) Hz with a 3 G bias field. The radial frequencies can be adjusted by changing the bias field. The radial frequency is

$$\nu = \frac{1}{2\pi} \sqrt{\frac{\mu_B m_f g_f}{m}} \frac{B'_x}{\sqrt{B_0}} \quad (9)$$

where m_f is the projection of the total angular momentum, m is the mass of Rb, B'_x is the field gradient, and B_0 is the bias field.




Fig. 13 Science cell region showing the magnetic trap holder. Two stainless-steel end caps on the ends of the Boron nitride form (white) are attached to a support structure behind the coil form. The microwave waveguide is shown on the left side of the picture directed towards the trapping region. Not shown is the objective lens on the back of the coil form.




Fig. 14. Ioffe–Pritchard magnetic trap (end on view). The permanent magnets are at a 45° angle with respect to the horizontal axis so as to provide a magnetic field in the same direction as the magnetic field from the quadrupole trap used to transport the atoms. The magnetic trap can be rotated by 45° and still confine the atoms as they are brought into the IP trap region by the quadrupole coils. Four permanent magnets, magnetized through the thin dimension, will also work to provide a two-dimensional quadrupole field with no field along the axial direction of the trap. However, using just two magnets magnetized through the thin direction would create a significant gradient along the axial direction of the trap, which would interfere with the transport of atoms via the moving quadrupole trap.

The trap configuration is shown in Figs. 14 and 15. The two permanent magnets are $5.05\text{ cm} \times 1.91\text{ cm} \times 0.64\text{ cm}$ grade 35 Nd/Fe/B, which combined produce a gradient of 450 G/cm . In a preliminary version of the apparatus, we used permanent magnets that produced a quadrupole gradient of 1200 G/cm . This gradient gave us 600 Hz radial trap frequencies at a bias field of 3 G . We found that having such tight confinement led to pronounced density dependent losses (presumably due to inelastic collisions), which were so large that the final evaporation stage was not efficient, and thus we produced smaller condensates with shorter lifetimes. A valuable lesson in designing evaporative cooling apparatuses is that provided one has a large initial load of atoms in a MOT, a larger transverse quadrupole gradient in the magnetic trap is not always better.

In the axial electromagnets (Fig. 15), the outer(inner) coils are each $20(10)$ turns of 18 gauge magnet wire held in place with thermally conductive epoxy (see Appendix B). In the normal configuration we run 13 A through the outer coils and 6.5 A through the inner coils producing an axial field curvature $\partial^2 B_{ax} / \partial x^2 = 60.6\text{ G/cm}^2$.




Fig. 15. Ioffe–Pritchard magnetic trap (side view) showing the permanent magnet and axial coil positions.

All of the trap components are mounted on a form made from boron nitride. This ceramic has a high thermal conductivity, similar to aluminum, so that the heat generated from the coils can be removed. It also has a low coefficient of thermal expansion, smaller than stainless steel, which ensures that the axial geometry, and thus the trapping field, remains constant as the trap holder changes temperature. Boron nitride also allows microwaves through without attenuation for frequencies below 10 GHz. Being transparent to microwaves is important for our imaging procedure, in which we make transitions between hyperfine ground states, and for our scientific goals.^{36–38} Boron nitride has the consistency of a hard chalk but can be machined into simple shapes (Fig. 16).

We choose to use a hybrid IP trap in the experiment because of its great stability, but it is obviously not the only solution. A fully electromagnetic trap would be necessary if an experiment required the magnetic field to be zero. A quadrupole with Ioffe configuration (QUIC) trap or the time orbiting potential (TOP) trap would work for this purpose.³⁵ The main requirement for a trap is for it to have around 450 G/cm quadrupole gradient, which is not hard to achieve with electromagnetic coils close to the 1.4 cm diameter cell. The quadrupole gradient must be large to have an acceptably high initial collision rate, of at least a few of Hz, to begin evaporation. Initially the cloud is not in the harmonic region of the trap and is mostly confined by the quadrupole gradient. Therefore the quadrupole gradient determines the initial collision rate. A cloud is in the harmonic region of the trap when the mean thermal energy is less than one “bias field worth of energy,” $\mu_B g_f B_0$; for a 3 G bias field a cloud is in the harmonic region when its temperature is below about 30 μ K.




Fig. 16. Machine drawing of the Boron-nitride hybrid IP-trap form (all units in mm).

6.5. Transfer between Magnetic Traps

Transferring atoms between a quadrupole trap and an Ioffe–Pritchard trap can not be done completely adiabatically due to the relative directions of the magnetic fields in each trap. If the transfer is done correctly however, one can limit loss in phase space density to a factor of 2 to 4 during transfer. After the sliding quadrupole coils have come to rest with the center of the quadrupole trap aligned with what will be the center of the IP trap, we start the transfer by slowly (~ 500 ms) ramping down the quadrupole gradient to a point where the cloud is approximately mode-matched in the axial direction; for our axial coils and initial temperature this corresponds to a vertical gradient of 100 G/cm. Next we discontinuously turn on the IP axial coils and turn off the quadrupole coils. We find that the timing of the traps turning on and off is not critical at the 5 ms level.

We optimize the transfer parameters by maximizing the phase space density after the transfer. It is difficult to image and determine properties of hot clouds in the IP trap for a few reasons. First there is a large magnetic-field-induced detuning across the radial direction of the cloud from the permanent magnetic field. Second the cloud's optical depth is large in the magnetic trap, which leads to systematics in determining the number. We overcome these problems by moving the atoms back to the MOT cell and

imaging them with fluorescence. We optimize the transfer between traps by transferring the atoms from the quadrupole trap to the IP trap and back again. The cloud's temperature and number measured in the MOT cell after being brought back from the IP trap are not a completely accurate representation of the parameters that existed in the IP trap, but the comparisons are at least monotonic, which is good enough to allow for optimization.

7. rf EVAPORATION

Now that we have atoms in the IP trap we can evaporatively cool them to degeneracy. The basic idea of evaporation is to remove atoms with more than the average energy of the cloud and allow the ensemble to equilibrate to a lower temperature through collisions.¹⁶ We need an adequate elastic collision rate to have the sample reequilibrate before there is a large loss of atom number or a large increase in energy of the sample from inelastic collisions.

There are three types of inelastic collisions we have to be concerned with during evaporation: one-, two-, and three-body processes. One-body loss from collisions with background gas atoms will cause essentially only number loss and does not induce heating, because all atoms in the trap have about the same probability of removal. During the initial stages of evaporation, one-body loss is the dominant factor because the density is low, inhibiting density-dependent collisions. As the density increases two- and three-body process become important. Two-body processes are significantly suppressed with a spin-polarized gas in the maximum angular momentum state of a ground hyperfine state. An upper bound on the rate constant has been determined experimentally to be 1.6×10^{-16} cm³/s for atoms in the $|F = 1, m_f = -1\rangle$ state.³⁹ Two-body loss is seldom an issue for ⁸⁷Rb in the lower hyperfine state. Three-body loss happens when three atoms collide, two forming a molecule, and the other taking away the residual energy. The per atom decay rate is proportional to density squared; the three-body rate constant has been measured to be $4.3(1.8) \times 10^{-29}$ cm⁶/s for noncondensed ⁸⁷Rb atoms in the $|F = 1, m_f = -1\rangle$ state.³⁹ The three-body process not only causes atom loss but also heating because atoms are preferentially lost from the highest density region of the cloud, which corresponds to the atoms with the least energy in a magnetic trap. When the density and spatial extent of the cloud are such that the products of three-body decay can no longer pass freely out of the cloud but instead multiply scatter, the total the depletion of atoms due to three-body collisions can occur much faster than that suggested by the simple rate constant, and heating can become significant. See Ref. 18 for a discussion of the threshold collision rate for “runaway” evaporation, but a reasonable rule of thumb is that the elastic collision rate should be at least 100 times larger

than the loss rate, except at the very end of evaporation, when larger losses may be tolerated.

We remove or evaporate the higher energy atoms by exploiting the fact that higher energy atoms tend to travel on trajectories that stray farther from the center of the magnetic trap into regimes of larger magnetic fields.⁴⁰ The trap volume is bathed in a spatially uniform, radio frequency magnetic field. There is an ellipsoidal surface of constant dc magnetic field at which the spin flip frequency of an atom is resonant with the rf. Atoms whose trajectories pierce this surface are transferred from $|F = 1, m_f = -1\rangle$ trapped state to the $|F = 1, m_f = 0\rangle$ untrapped or $|F = 1, m_f = 1\rangle$ antitrapped state and are permanently ejected from the trapping region. By ramping down the frequency, we shrink the ellipsoidal surface, forcing evaporation to continue even as the temperature and the mean cloud radius decrease. The goal is to maintain the cloud in approximate thermal equilibrium with $k_B T$ about six times less than the atom's potential energy at the ellipsoidal surface of resonance. If the collision rate is constant, we want to remove the same fraction of energy from the cloud per unit time. This condition corresponds to an exponentially decreasing frequency ramp. As the collision rate changes so will the optimum exponential time constant. The functional form we use is

$$\nu(t) = (\nu_{\text{start}} - \nu_0) e^{-t/\tau} + \nu_0, \quad (10)$$

where ν_{start} is the frequency where we begin evaporating, ν_0 is the frequency corresponding to the bottom of the trap, and τ is the exponential time constant of the ramp. The optimum time constant depends on the elastic collision rate and loss rate.

7.1. rf Coil

We use a simple single-loop coil to deliver rf to the atoms for evaporation. We typically evaporate from 40 → 2 MHz. The large range of frequencies we use prevents us from impedance matching the coil to gain better coupling. Because we are in the near-field limit of the radiation for all evaporation frequencies, one can think of the rf as just an oscillating magnetic field. Only the component of the oscillating magnetic field perpendicular to the local quantization field will cause transitions between the Zeeman states. We place the coil directly outside the glass cell in the narrow space between the outer diameter of the glass cell and the inner diameter of the Boron-nitride trap form (Fig. 16). with the axis of the loop perpendicular to the bias field to maximize the coupling to the atoms when they are cold. Hot clouds will have atoms in the quadrupole field with

quantization axes in every direction in space, so that there will be small regions in the cloud that are not affected by the rf. This does not appear to pose a serious problem. Avoid placing the coil closer than one radius to any electrically conductive object; the conductive object will reduce the flux return path and thus the magnetic field produced. The size of our loop, made from 18 gauge magnet wire, is about 1 cm in diameter. The loop is soldered directly to a RG 175 cable leading to a rf amplifier.

7.2. Evaporation Optimization

We need several stages of evaporation, each with different parameters. Throughout evaporation both the elastic and inelastic collision rates change as well as the rf coupling to the atoms, thus we must adjust the evaporation time constant and the rf power for each stage. As the atoms cool, we decrease both the time constant, due to the increased collision rate, and the rf power delivered to the atoms, to avoid power broadening effects.

Power broadening of the rf “knife” will cause the evaporation process to lose energy selectivity as the width of the knife becomes comparable to the temperature of the cloud. Because the atoms initially have a larger velocity, for the early stages of evaporation we need more rf power to insure that atoms piercing the ellipsoid of resonance will undergo a spin flip. Later, as the cloud approaches zero temperature one must be very careful not to apply too much rf power. Another potential problem with setting the rf power is coil or amplifier resonances. The rf coil may have a self-resonant frequency in the frequency range spanned by the evaporation. An easy way to check for resonances is to measure the rf power delivered to the atoms using a small pick-up coil placed near the evaporation coil. The rf power from the synthesizer may have to be drastically reduced near a resonance to avoid power broadening.

We break up the evaporation into enough stages so that we decrease the temperature by a factor of 2 or 3 with each stage; this criterion sets the start and stop frequency for each stage, points A–E in Fig. 17. We start with a time constant of 10 s. In general the time constant of the evaporation ramp should be about a factor 10 to 20 greater than $1/(\text{collision rate})$. We do not actually continuously ramp the frequency of the rf for the initial stages but instead send discrete steps to the synthesizer through the General Purpose Interface Bus (GPIB). Typically a single GPIB command will take between 30 to 50 ms to be received and executed; therefore we send a new frequency command every 50 ms. The discrete nature of the frequency ramp is not a problem for the initial stages when each step is small compared to the temperature of the cloud (i.e., when the frequency ramp time constant $\tau \gg 50$ ms), but it is a problem in the last stage of evaporation.

For the last stage we sometimes use a programmable frequency synthesizer that can phase continuously ramp the evaporation frequency. The extra synthesizer is not necessary but will produce larger condensates.

We want to optimize the collision rate for each stage of the evaporation. If we image the cloud right after the stage we are optimizing, the cloud may not be in equilibrium due to a too-rapid cut. Imaging a cloud out of equilibrium can systematically misrepresent the collision rate. However if we add an additional evaporation stage before imaging, we can circumvent this problem. The additional stage will not be as efficient if the cloud is out of equilibrium, and thus the cloud will have fewer atoms after the additional stage.

Optimizing the initial stages of evaporation in our trap is difficult because of our inability to image hot clouds. As stated before, we can not obtain an accurate temperature or number of atoms in our cloud when the temperature is above $1 \mu\text{K}$ because of the magnetic field gradients. However, we have created an optimization procedure for the first stages of evaporation that works well enough. We start with two stages (A–C in Fig. 17). The parameters of the first stage, segment A–B, are varied, while the second stage, segment B–C, parameters are kept constant. We image the cloud at point C and maximize the peak optical depth (OD) by changing the parameters for segment A–B. Even with imperfect imaging, the peak OD measured after ramp B–C is monotonic in the true equilibrated collision rate produced by ramp A–B. Next we add a stage C–D and optimize segment B–C and so on. It is important to iteratively adjust the time constant and rf power because they are coupled. The initial evaporation is not very sensitive to the parameters of the cut so this procedure works well.

The final stages of evaporation are more critical than the first stages. Fortunately, for the last stages we can accurately determine the temperature and density of the cloud. Except when optimizing the very last stage, we characterize a given stage by optimizing number in the cloud after an additional stage. It is easy to walk the parameters in the wrong direction, especially with the final stage; one tends to have too short a time constant and too much rf power. We reduce the rf power 19 db from the first to the




Fig. 17. Sample evaporation trajectory with four segments shown. We typically use eight segments, each providing a factor of 2 to 3 decrease in $(v - v_0)$.

TABLE III

Experimental Evaporation Parameters for a Trap With a 3.2 G Bias Field, where the rf Power is the Amplitude of the Signal out of rf Amplifier. Because we change the Frequency over an Order of Magnitude, Coupling into the coil varies considerably and there is no fixed relationship between rf Power and Actual Applied Field. At 3 MHz, the rf Power of 14 dBm corresponds to a rf Magnetic Field Magnitude at the Atoms of approximately 20 mG. The Value of ν_0 Is 2.26 MHz.

Stage	ν_{start} (MHz)	ν_{stop} (MHz)	τ (s)	rf power (dBm)
1	40	20	10	25
2	20	10	5	20
3	10	5	4	20
4	5	3	4	18
5	3	2.60	2	14
6	2.60	2.44	2	14
7	2.44	2.40	1.5	10
8	2.40	2.28	1	6

last stage (Table III). Generally we can change the rf power by plus or minus 5 db in the upper stages and 3 db in the later stages without observing a significant change in evaporation efficiency. When optimizing the last few stages it is also important to remeasure the trap bottom ν_0 , the frequency at which the last of the atoms disappear, which can be more accurately determined now that one has a cold cloud. An example of the rf evaporation parameters is given in Table III.

8. ABSORPTION IMAGING

We image clouds in the IP trap using laser absorption. We illuminate the cloud with resonant light, atoms scatter photons out of the beam, and we focus the shadow cast by the atoms onto a charge coupled device (CCD) array. The amount of light absorbed gives the column optical density (OD) along a particular ray through cloud. Optical density is defined by Beer's law and is given by

$$I = I_0 e^{-OD}, \quad (11)$$

where I_0 and I are respectively the intensities entering and emerging from the atom cloud. In essence, everything that is experimentally known about ultra-cold atoms has come from the analysis of images of optical density structures.

8.1. Optical Setup

The imaging optics are shown in Fig. 18. We use a probe beam that has been filtered spatially by a single-mode fiber. The probe beam is




Fig. 18. Optical system for absorption imaging. Light from a single mode fiber is put through a polarizing beamsplitting cube (PBS) and a half-wave plate to adjust the polarization. After the light passes through the cloud, the image is focused onto a CCD camera by two lenses.

expanded to a diameter of 1 cm so that the intensity across a $100\ \mu\text{m}$ condensate is nearly constant. The light first passes through a polarizer and then through a $\lambda/2$ plate so that we can adjust the angle of the linear polarization. The incoming probe beam passes through a $4\ \text{mm}^2$ aperture on the magnetic trap form to reduce excess light that could scatter into the camera from defects in the glass cell (Fig. 13). The shadow of the atoms is focused onto the camera with two lenses. We use, as the objective, a 1 cm diameter gradient-index singlet lens. The objective lens is mounted directly on the trap coil form to collect the largest possible solid angle. We use a 30 mm diameter achromat doublet as the second lens. We use an achromat not for its reduction in chromatic aberrations, but for its low spherical aberrations when oriented correctly. Our CCD camera is a front-illuminated CCD array with pixels $13\ \mu\text{m}$ on a side. The entire array is 1024×1024 pixels and the readout has 16 bit resolution. The quantum efficiency at 780 nm is around 35%, and the readout noise is 6.1 electrons per pixel in the fastest readout mode of 1 MHz.

The intensity of the probe beam is about $0.3\ \text{mW}/\text{cm}^2$. The frequency of the probe beam is set by adjusting the difference frequency of the two AOMs shown in Fig. 2. This gives the offset frequency from the peak to which the laser is locked. For different imaging schemes we lock to different lines, but for the high-field scheme described below we need a frequency about 140 MHz red of the zero-field $F' = 3 \rightarrow F = 2$ transition, and so it is most convenient to lock to the cross-over peak A (Fig. 3), which is 133.5 MHz red of the zero-field transition and offset about 6.5 MHz red.

8.2. Imaging in a Non-Uniform Magnetic Field

Using permanent magnets requires us to image the atoms in a non-uniform magnetic field. We will discuss solutions to the problems of imaging in a spatially varying magnetic field. One must also keep in mind

that a fully electromagnetic trap (i.e., QUIC or TOP trap) could be used and thus eliminate many of these problems.

Imaging atoms in a strong quadrupole magnetic field presents several problems. Our goals are to (i) expand the cloud well above the resolution limit of our imaging system, (ii) extract useful parameters from the image without systematic errors, and (iii) image along a radial direction so we can observe dynamics along the axial direction. These goals are difficult to meet, because first we obviously can not turn off the quadrupole field created by the permanent magnets to allow for the usual expansion. Second, imaging in the radial direction causes the bias (quantization) axis to be perpendicular to the propagation direction of the probe beam, and thus does not allow us to drive purely $\sigma+$ or $\sigma-$ transitions, which is desirable because it would give us a cycling transition and thus a large signal-to-noise ratio. The last obstacle to overcome is the magnetic field gradient, which causes a spatially varying energy shift due to the Zeeman effect. Therefore, we can not apply light which is resonant with the entire cloud. The spatially varying detuning could cause the image to have systematically the wrong width and optical density.

We have found ways to reduce or eliminate all of our imaging problems. We expand the cloud by transferring the atoms to an anti-trapped state and allowing them to fall off of the potential created by the magnetic trap. We use a microwave adiabatic rapid passage (ARP) to transfer coherently the atoms from the $|1, -1\rangle$ to the $|2, -2\rangle$ state.⁴¹

To ARP the atoms from one state to another we turn on a microwave coupling field far off resonance, ramp the frequency slowly, compared to the Rabi frequency, through resonance, and then turn off the field. This coherently transfers the atoms between the two states. The microwaves, generated by a commercial microwave synthesizer, are transmitted to the atoms by a sawed-off waveguide, which is placed near the IP trap, directed along the axis of the trap (Fig. 13).

Second, we reduce the Zeeman detuning across the cloud by increasing our bias (axial) field to 100 G; this is easily accomplished by reversing the current in the inner coils. The transverse gradient adds in quadrature with the large bias field and thus reduces the spatial variation of the magnetic field from 2.4 G to 0.1 G for a typical expanded radial cloud radius of 100 μm . The residual variation in magnetic field across the cloud corresponds to a spatial inhomogeneity in the resonant frequency of only 140 kHz, which, being much less than a natural linewidth, has no effect on the image.

The last problem to solve is the incorrect imaging polarization. We would like to drive a cycling transition from the $|2, -2\rangle$ to the $|3', -3\rangle$ state. We choose our probe beam polarization linear and perpendicular to the bias field and thus drive in principle not only the desired transition but also the $|2, -2\rangle \rightarrow |3', -1\rangle$ transition, which is obviously not a cycling

Propagation direction wrt quantization axis	Polarization	Transitions wrt quantization axis
Parallel	Right circular	$\sigma+$
Parallel	Left circular	$\sigma-$
Parallel	Linear	$\sigma+, \sigma-$
Perpendicular	Circular	$\sigma+, \sigma-, \pi$
Perpendicular	Linear along quantization field	π
Perpendicular	Linear perpendicular to quantization field	$\sigma+, \sigma-$

Fig. 19. Possible transitions with different probe beam polarizations.

transition. However we image in 100 G bias field, which breaks the degeneracy of these two transitions by 31 linewidths and allows us to have an effective cycling transition for hundreds of photon scattering events. Figure 19 lists the possible transitions for the different probe beam propagation directions and polarizations.

Naively one might expect that at maximum only half of the light could be absorbed, because one thinks of linear light as an equal amount of $\sigma+$ and $\sigma-$ light. However in the atoms' frame these two polarizations are not the correct basis and are actually coupled. Absorption of this type is typically known as the Voigt effect. In fact all of the light can be absorbed by the atoms, and the only effect of the direction of the polarization is to reduce the line strength by a factor of 2.

8.3. Imaging Procedure

We start the imaging procedure with the atoms in the $|1, -1\rangle$ state. We ARP the atoms to the $|2, -2\rangle$ using microwaves. The microwave field must be swept phase continuously for atoms to be efficiently transferred between states. Alternately, one may hold the microwave frequency constant and ramp the atoms resonance by ramping the bias magnetic field. We typically start about 1 MHz (i.e., 1.4 G) away from resonance and sweep through in 0.3 ms; our Rabi frequency is around 100 kHz. Next we

TABLE IV
Image Timing

Time (ms)	Event
0	Camera triggered, probe beam shutter open
9	Centering coil on
10	Microwave on, bias field ramp for ARP started
10.3	Bias ramp stopped, microwave off
10.3	Bias field jumped to 48 G for expansion
10.3 + Expansion time	Bias field jumped to 100 G for imaging
10.4 + Expansion time	Probe beam AOM on
10.402 + Expansion time	Probe beam AOM off
20	Probe beam shutter closed

jump the bias field to 48 G, and then wait for the anti-trapped atoms to expand. If we expanded in our normal 3 G bias field trap, the atoms would expand too rapidly into the anharmonic region of the trap, thus making it difficult to calculate the effect of this expansion. On the other hand, if we jump directly to a 100 G bias field, the atoms would expand too slowly and fall under gravity, once again into the anharmonic region of the trap. The intermediate field keeps the atoms in the harmonic region of the trap during the entire expansion. The atoms are also slightly sagged in the trap due to gravity, so when they begin to expand they are sitting on the side of the potential, which induces some asymmetry to the expansion. We correct for this sag by applying a small magnetic field (~ 0.3 G) to shift the center of the trap just below the center of the cloud just before the expansion. We find the correct magnitude of the centering field by imaging the cloud after long expansion times and adjusting the added field until the cloud remains fixed in the vertical direction during expansion.

After the cloud has expanded the desired amount, we jump the bias field to 100 G and flash the probe beam for 20 μ s. We use a short 20 μ s pulse for two reasons. First we do not want the atoms to be excited to the $|3', -1\rangle$ state and fall back to a dark state. We are only 200 MHz detuned from the $|2, -2\rangle \rightarrow |3', -1\rangle$ transition and therefore will drive transitions to this state, although with a very low probability. Second if the atoms in the cloud scatter too many photons they will pick up enough momentum to move along the direction of the probe beam; this motion could blur the image or cause the atoms' transition frequency to change as they move into regions of larger magnetic field.

After we acquire our data image, I_{atoms} , we take two additional pictures for normalization purposes. One normalization image, I_{light} , is taken with the probe beam on but with no atoms present; this gives our light

image which we use to calculate percent absorption. The other normalization image, I_{dark} , is taken with the probe beam off and the camera shutter open. This image will give a calibration of the camera dark current as well as any stray light that does not come from the probe beam. The images are taken 800 ms apart, which is limited by the readout from our camera. We calculate the OD of each pixel, which is given by

$$\text{OD}_{\text{meas}} = \ln \left(\frac{I_{\text{light}} - I_{\text{dark}}}{I_{\text{atoms}} - I_{\text{dark}}} \right). \quad (12)$$

There are two common systematics that should be addressed with any absorption imaging system. One is that in practice the maximum observable optical density saturates. Any probe beam light collected by the camera that can not be absorbed by the atoms will reduce the observed OD. Two usual culprits are off-resonant light and scattered probe-beam light. A good way to check how much of the probe beam is far off resonant is to send the probe beam through a heated Rb vapor cell and measure the percent transmitted. Some diodes have a broad pedestal of light that is not in the main frequency mode of the laser and which, being far from the atomic resonance, can cause a reduction in the observed OD. The second reason for a low maximum observable OD is probe light which is indirectly scattered onto the CCD. We place a small aperture in front of the cell to reduce scattering light from the cell onto the CCD. In practice we observe a maximum OD (OD_{sat}) around 2.8, even for clouds for which the actual OD is much greater. We must correct for the effect of the OD saturation during the image analysis. The modified OD, taking out the effect of OD saturation, is

$$\text{OD}_{\text{mod}} = \ln \frac{1 - e^{-\text{OD}_{\text{sat}}}}{e^{-\text{OD}_{\text{meas}}} - e^{-\text{OD}_{\text{sat}}}}. \quad (13)$$

We measure OD_{sat} by creating a dense cloud and expanding it for 1 ms. The center of the cloud will have a flat top where the OD is saturated at the maximum value. If the correction factor between OD_{meas} and OD_{mod} is too large, the potential for error increases. We increase the expansion of the cloud until $\text{OD}_{\text{meas}} < \text{OD}_{\text{sat}}/2$.

The other systematic with absorption imaging is the effect of probe beam intensity saturation. The actual OD is

$$\text{OD}_{\text{actual}} = \text{OD}_{\text{mod}} + (1 - e^{-\text{OD}_{\text{mod}}}) \frac{I}{I_s}, \quad (14)$$




Fig. 20. Energy level diagram showing imaging transitions in a 100 G magnetic field. The microwave ARP transition and optical imaging transition are shown with solid arrows.

where I is the intensity of the probe at the position of the cloud and I_s is 3.2 mW/cm^2 for Rb on a cycling transition with our imaging polarization. We like to minimize the correction factor, so we work at $I < I_s/10$.

The resonant frequency for the imaging transition can be calculated easily because both the initial and final states are maximum angular momentum states. F and m_f are therefore good quantum numbers even in a magnetic field of 100 G, and the frequency splitting between the two states is $\Delta\nu = \mu_B B_0$. We confirm we are on resonance by taking a transition line shape, which involves producing a series of identical clouds, and probing them with different frequencies. We change the frequency of the first probe AOM (Fig. 2) and measure the peak optical depth. The resulting curve should be a Lorentzian with the natural linewidth, Γ . Measuring the natural linewidth with the expected center implies that many parameters are correct in the imaging system, including narrow laser linewidth, accurate calibration of magnetic fields, probe beam well below saturation, and correct control of probe frequency.

The line shape can also be a useful diagnostic for probe laser frequency noise. Often the probe laser frequency may be affected by shutter-induced vibrations or current transients right before imaging. Therefore it is important to measure the noise on the laser during the imaging pulse. One can find the shot-to-shot standard deviation of the measured atom number while the probe is tuned on resonance and contrast while the probe beam is half linewidth off resonance. Comparing the two measurements rejects uncorrelated atom number fluctuations. A significant increase in shot-to-shot noise when the laser is tuned a half linewidth off resonance indicates probe laser frequency or magnetic field noise.

8.4. Focusing the Image

We focus the image onto the CCD camera by imaging a small (few times our resolution limit), low density, low OD ($OD < 1$) cloud that has not expanded much. Before we focus the image we first take a line shape to ensure we are on resonance. Above and below the optical resonance frequency the real part of the index of refraction of a gas differs from one, and the ellipsoidal cloud of gas will not only absorb light but also refract or “lens” it. Once we have tuned the probe laser to the resonant frequency of the atomic transition, we adjust the position of the camera along imaging axis. The focus of the image will be at the minimum cloud width. We focus the image by measuring the width in the radial direction. This position is not necessarily the focus in the axial direction of the cloud because of the astigmatism induced by the cylindrical cell. If the cloud is exactly in the center of the glass cell the image will not be astigmatic because all the rays of light hit perpendicular to the glass and therefore are not refracted.

8.5. Measuring the Image Magnification

We measure the magnification of our imaging system by watching a cloud fall under the influence of gravity. We start with atoms in the $|1, -1\rangle$ state and perform an ARP to place them in the $|2, 0\rangle$ state, which is affected only slightly by the magnetic field gradient. We allow the atoms to fall for a varying time and measure their resulting position. A cloud’s position as a function of time including the small acceleration due to the second order Zeeman shift is

$$z(t) = -\left(\frac{a}{p}\right) (\cos \sqrt{p} t - 1) + z_0, \quad (15)$$

where a is the acceleration due to gravity in pixels/ms², z_0 is the cloud's initial position, and p is

$$p = \left(\frac{4\pi\hbar}{m} \right) f B'_x, \quad (16)$$

where f is the second order Zeeman shift of 287 Hz/G², and B'_x is the radial magnetic field gradient. Fitting the position versus drop time data will give a value for a , which can be used to find the magnification, which is

$$\text{Magnification} = \frac{9.81 \mu\text{m}/\text{ms}^2}{\sqrt{2} a}. \quad (17)$$

The factor of $\sqrt{2}$ is included because our imaging axis is at 45° with respect to gravity.

9. IMAGE ANALYSIS

9.1. Image Processing

After the three images have been downloaded to the computer we apply some image processing before fitting the images. We first calculate a measured OD for each pixel using Eq. (12). Occasionally we will get pixels that have anomalously high or low values due to noise or readout error. We remove these pixels by systematically going through the image array comparing nearest neighbor pixels. If there is a difference of 6 or greater, in units of OD, the pixel is replaced by the average of the eight adjacent pixels. We perform the same procedure a second time, this time using a difference threshold of 0.8. After the spikes are removed, we apply corrections for saturation effects using Eqs. (13) and (14). Finally, depending on the size of the cloud, we bin the array. For most condensate images we use 2×2 binning, making a modest sacrifice in resolution in order to reduce calculation time during the fitting routine. The column density of atoms at each point in the image is just OD/A, where the absorption cross-section A is just

$$A = \left(\frac{\text{branching ratio}}{2} \right) \left(\frac{3\lambda^2}{2\pi} \right) \times \frac{1}{1 + 4 \frac{\Delta^2}{\Gamma^2}}, \quad (18)$$

where λ is the wavelength of the transition, and the factor of two in the denominator is due to our particular imaging polarization. The branching ratio for the $|F=2, m_f=-2\rangle$ to $|F'=3, m_f=-3\rangle$ transition, read from Fig. 24 below, is $15/15=1$.

9.2. Image Fitting

We use three different fitting routines depending on the degeneracy of the cloud.²¹ For clouds above the condensation temperature we fit the image to a 2-D Gaussian. Clouds at finite temperature but with a condensate present we fit with two separate functions. The condensate portion of the image can be fit to a Thomas–Fermi profile, which is a paraboloid integrated along the line of sight. The thermal cloud is no longer an ordinary Gaussian when it is degenerate but is modified by Bose statistics and must be fit with the appropriate function.²¹ For instance, using an ordinary Gaussian to fit the normal cloud in the second image in Fig. 21

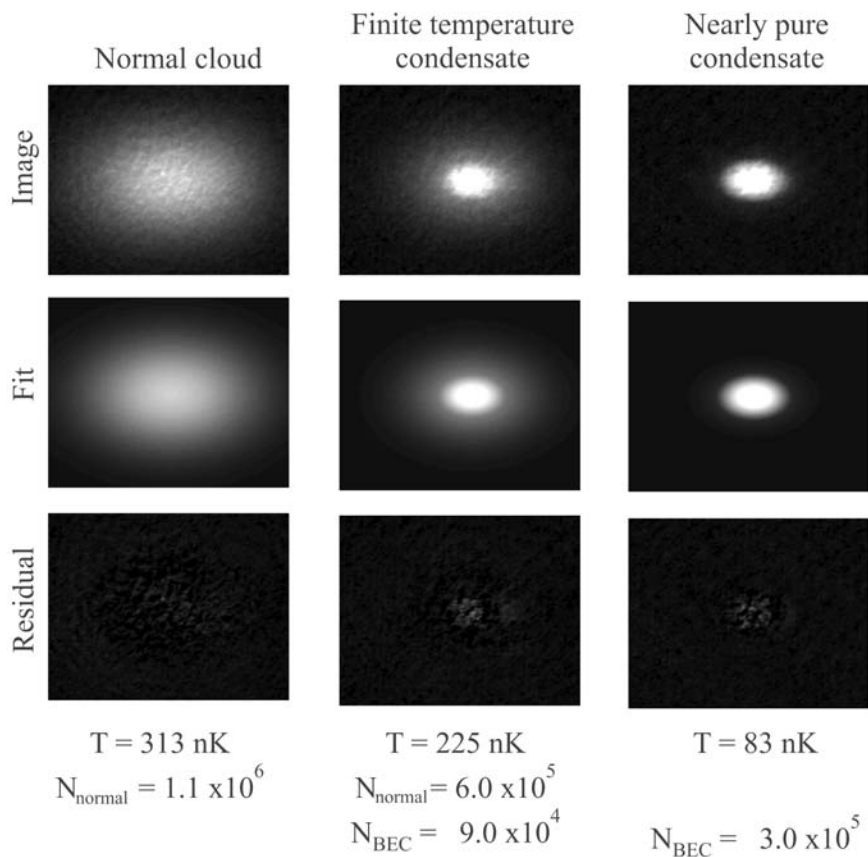


Fig. 21. Examples of images, fits, and residuals of clouds above and below the BEC transition temperature. The normal cloud image was fit using a Gaussian profile. The two images below the transition temperature were fit using a modified Gaussian plus a Thomas–Fermi profile.

underestimates the temperature by 11%. In some cases, where the cloud has no detectable thermal fraction, we just use a Thomas–Fermi distribution. The fitting is done using a Matlab script called from inside LabVIEW. Examples of images, fits, and residuals are shown in Fig. 21.

9.3. Calculating Cloud Parameters

Once we fit the image and extract the fitting parameters we can calculate the properties of the cloud. The first step is to calculate the size of the cloud in the magnetic trap based on our anti-trapped expansion. The Boltzmann equation gives us the functional form of the expansion of a normal cloud in an antitrapped state. The in-trap cloud size is given by

$$\sigma(t=0) = \frac{\sigma(t) \omega}{\sqrt{\omega^2 + (\omega^2 + \omega_0^2) \sinh^2(\omega t)}}, \quad (19)$$

where $\sigma(t)$ is the cloud size after expansion, $-i\omega$ is the harmonic trap frequency during the expansion, ω_0 is the original trapping frequency, and t is the expansion time. This treatment assumes that the initial position and velocity are uncorrelated and that the mean-field does not contribute significantly to the expansion. The effects of being in the hydrodynamic regime, which do affect ballistic expansion, are insignificant for anti-trapped expansion. The condensate expansion is similar to that of the normal cloud except it does not have an initial velocity spread. The axial in-trap condensate size is given by

$$\sigma(t=0) = \frac{\sigma(t)}{\cosh(\omega t)}. \quad (20)$$

(Note Eq. 20 is the same as Eq. 19 with ω_0 set to zero.) The mean-field contribution to the expansion is negligible in the axial direction, from which we calculate all size dependent parameters. The radial expansion of the condensate, however, is significantly affected by the mean field. We calculate the temperature and density of the cloud from the measured axial width and the known in-trap aspect ratio.

10. COMPUTER CONTROL

We use two computers to run our experiment; one computer controls all of the timing, digital, analog, and GPIB commands, while the other is dedicated to running the camera. To initiate an image acquisition, the control computer externally triggers the camera control board, which in

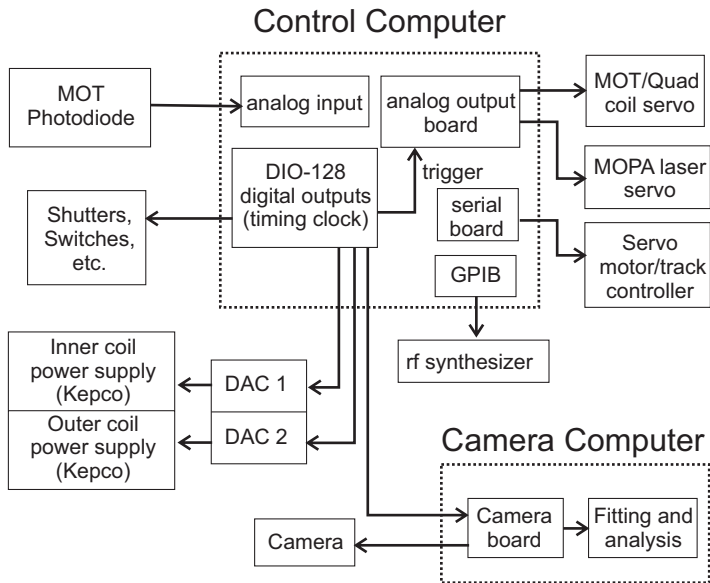


Fig. 22. Computer control diagram showing the different computer boards and what they control.

turn triggers the camera computer. It is useful to allocate control and acquisition tasks to different computers so that the camera computer can analyze the data from the previous shot while the control MOT computer moves on to the next shot. The preliminary analysis of each shot, which includes calculating cloud parameters (such as temperature, density, collision rate, and number in the normal cloud and density, chemical potential, and atom number in the condensate), is completed in real time, greatly increasing the amount of data that can be compiled and digested in a day.

A Bose–Einstein condensation experiment requires precise temporal control of a variety of components. Most functions require timing resolution on the millisecond scale, but for certain key tasks, such as imaging, expansion, and microwave spectroscopy, we need timing on the microsecond scale. There are several basic types of outputs and inputs our control system needs to handle. We need digital, analog, serial, and GPIB outputs. Most of the experiment is controlled by digital outputs, which control items such as shutters, rf switches, and AOMs. Our magnetic coils, both quadrupole and IP trap, have servos that require analog voltage set points. The servo motor driving the track is controlled via a serial connection. We also have several instruments including rf and microwave synthesizers that

use GPIB as the main mode of communication. Our only input port, excluding the camera, is an analog voltage from a photodiode that monitors the fluorescence from the atoms in the MOT. This input is fed into a multipurpose analog input board produced by National Instruments.

We have come up with a complete timing system that includes all the different I/Os. Four computer boards and two external digital-to-analog converters (DACs) make up the control hardware. The programming software we chose is LabVIEW, which is easy to use but has some limitations because unfortunately it runs in a Windows environment. LabVIEW timing can vary by up to 10 ms shot-to-shot, because the operating system can interrupt the program at any time. Therefore we need another source to handle our precise timing. LabVIEW handles our imprecise events, such as GPIB commands, quadrupole coils current ramps, and track motion.

We use a digital input/output board (DIO-128) as the main clock in our system. It has 64 digital inputs and outputs and an internal oscillator that has 500 ns resolution.²⁵ LabVIEW loads the DIO-128 board with an array of time stamps and port levels, which specify the state (hi/low) of each digital port at each time stamp, into the first-in-first-out buffer on the board. The buffer can hold up to 16000 words. When we want the timing sequence to start we send a trigger to the DIO-128 from within LabVIEW. From that point on until the buffer is cleared, the board no longer communicates with the computer and thus is not susceptible to operating system interrupts.

We have two different devices to produce the analog output voltages required in the experiment. The first device that produces an analog voltage is a National Instruments analog output board, which resides inside the noisy environment of the computer. We use this board to control items that are not very sensitive to voltage noise, such as the quadrupole coil current servos. The analog output board can store an array of voltage values and output them when triggered by the DIO-128. The second device, for more noise-sensitive applications, is a pair of 16 bit DACs, located external to the computer case and digitally controlled by the DIO-128. The power supplies which drive the IP trap coils are controlled with the low noise DACs. More detail on the electronic control is available on request.²⁴

11. ROBUSTNESS OF DESIGN

We claim our experimental system is robust and can produce condensates with the system in a less than optimum configuration. We tested this claim by deliberately misadjusting several parameters in the experiment until we saw a reduction of resulting condensate number by a factor of 2 from the fully optimized configuration. These tests gave in some cases an unduly pessimistic view of the vulnerability of the experiment to the

degrading of any one particular performance specification, because we did make any compensating adjustments in the other operating parameters. For instance, the deliberate reduction in MOT trapping beam power caused there to be fewer atoms collected in the MOT, and that in turn led to less efficient evaporation and ultimately smaller condensates. We know from experience however that a smaller MOT yield can be partly compensated for by revisiting the detuning of the CMOT and the time constants of the evaporative sweeps. To simplify the procedures of the tests described below, we did not do this sort of reoptimization, and the results represent therefore a sort of worst-case limit on our sensitivity to a particular parameter.

We first examined what was the maximum background pressure we could have in the science cell and still make a condensate. We found that we needed at least a 25 s magnetic trap lifetime, limited only by background gas collisions, to create a condensate. Our vacuum system routinely produces a lifetime of 170 s or greater.

We concentrated our sensitivity tests on two stages of the experiment: MOT/CMOT and moving coil transfer. For the MOT/CMOT stage we adjusted the power and size of the trapping beams as well as the power in the repumping beam. We found reducing the diameter of an aperture in the trapping beam from 50 mm to 22 mm reduced our condensate number by a factor of 2. This aperturing corresponded to a power reduction of $\sim 25\%$. As a separate test we uniformly reduced the power in the unapertured trapping beam from 160 mW to 125 mW before seeing the factor of 2 condensate loss. We also found we have more than the required amount of repump power. We had to reduce the power in the repump beam by a factor of 5 to give us a factor of 2 reduction in condensate number.

We also looked at how sensitive the system was to the positioning of the quadrupole coils at both ends of travel. We found that the servo linear track's reproducibility of $5\ \mu\text{m}$ was much better than was required. It took a displacement of 3 mm at either end of the travel to decrease the number in the condensate by a factor of 2. These simple tests give an sense of the robustness of our design. Also our experiments in microwave Ramsey spectroscopy, which are not discussed in this text, have produced spectroscopic measurements with precision greater than 1 part in 10^{11} , which attests to the stability of our design.³⁸

12. CONCLUSION

In conclusion we have successfully designed and constructed a simpler system to create a Bose–Einstein condensate. We hope this paper will encourage scientists outside of the trapping and cooling community⁴ to find innovative new uses for Bose–Einstein condensates.

APPENDIX A

TABLE V
Properties of ^{87}Rb

Property	Symbol	Value
Cooling transition		$5S_{1/2}, F = 2 \rightarrow 5P_{3/2}, F = 3$
Nuclear spin		$3/2$
Wavelength in vacuum	λ	780.23 nm
Mass	m	1.44×10^{-25} kg
Lifetime of upper state	τ_n	27 ns
Natural linewidth	Γ	5.9 MHz
Saturation intensity(stretched transition)	I_s	1.6 mW/cm ²
Recoil temperature	T_{rec}	180 nK
Recoil velocity	v_{rec}	0.59 cm/s
Ground hyperfine splitting	ω_{hf}	6834.68261090434(3) MHz
$ 1, -1\rangle$ s-wave scattering length	a_{11}	100.44 a_0
$ 1, -1\rangle/ 2, 1\rangle$ s-wave scattering length	a_{12}	98.09 a_0
$ 2, 1\rangle$ s-wave scattering length	a_{22}	95.47 a_0




Fig. 23. Energy level diagram for ^{87}Rb showing Landé g factors for each state.

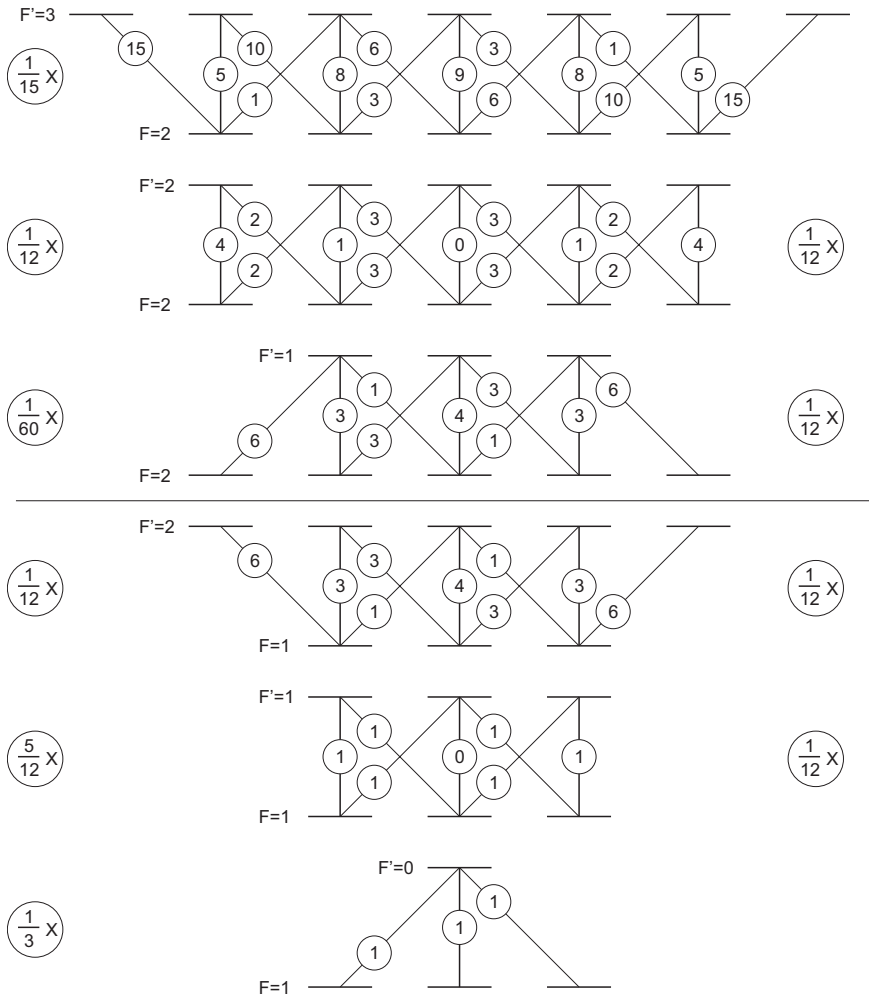
D2-line (780 nm) 87Rb D1-line (795 nm)

Fig. 24. Branching ratios for ^{87}Rb . Multiply by the circled number in the left (right) column to get the branching ratio for the D2 (D1) line.

APPENDIX B

The main components of the BEC apparatus are listed in this appendix. Cables, some power supplies and other common items are not listed but are used in the experiment.²⁶

Item	Quantity	Company	Part number
MOPA laser	1	Toptica	TA100
External cavity diode lasers	2	New focus	Vortex
Dielectric mirrors (1")	25		
Polarizing beamsplitting cubes (1")	5		
Waveplates A.R. coated ($\lambda/2$, 1")	7		
Dielectric mirrors (2")	10		
Polarizing beamsplitting cubes (2")	5		
Waveplates A.R. coated ($\lambda/2$, 2")	5		
Waveplates A.R. coated ($\lambda/4$, 2")	6		
Lens (2")	2		
Lens kit (1")	1		
Mirror mounts (1")	30		
Rotation mounts for waveplates (1")	7		
Mounts for PBS (1")	5		
Lens mounts (1")	15		
Mirror mounts (2")	10		
Rotation mounts for waveplates (2")	11		
Mounts for PBS (2")	5		
Lens mounts (2")	2		
Standard posts (4")	70		
Standard post holders (4")	70		
Post holder bases	70		
Posts (1" diameter)			
Single-mode fiber	1	Tempo	C2C2-1P8-02
Fiber launchers (FC conected)	2	Thorlabs	F2230FC-B
Shutters	4	Uniblitz	LS3T2-105
Optical isolators	3		
Rb vapor cells	3	Technical Glass Inc.	
Photodiode boxes for sat. spec.	3	Home built	
Acoustic optic modulators	3	NEOS	
Voltage controlled oscillators	3	Varil	
Photodiode for MOT monitor	1		
Laser lock boxes	3	Home built	
Lock in amplifiers	3	Home built	
Digital scopes	3	Tektronics	TDS210
Function generators	1	Tektronics	CFG253
Rf synthesizer	1	Agilent	HP8656B
Rf switches	4	Mini-circuits	ZFSW-2-46
Rf amplifier (evaporation, 5 W)	1	Mini-circuits	ZHL-1-2W-S
Rf amplifier (AOMs, 2 W)	3	Mini-circuits	ZHL-5W-1
Microwave synthesizer	1	Agilent	HP8673E
Microwave switch	1	General microwave	F9114A
Microwave amplifier (6.4 W)	1	Microwave power	L0408-38
Microwave circulator	1	Narda microwave	4914
Microwave directional coupler	1	Narda microwave	40146-30
Microwave square law detector	1		
Microwave waveguide	1	Pacific wave systems	D-268, D-200-5
IP trap power supplies	2	Keppco	BOP 20-10M, BOP 20-20M
Quadrupole coil power supplies	1	Agilent	HP6681A
Getter current supply (0-6 A)	1	Topward	6306D

Item	Quantity	Company	Part number
Hall current sensor	1	F.W. Bell	CLN-300
Centering coil power supply	1		
Power MOSFETs	3	Advanced power tec.	APT 10M07JVR
CCD camera	1	Andor	DV434
Security camera	1		
Security camera monitor	1		
Timing/control board	1	Viewpoint USA	DIO-128 or DIO-64
Analog output board	1	National instruments	10 channel analog output
GPIB board	1	National instruments	PCI-GPIB
Multipurpose board	1	National instruments	PC-LPM-16PnP
Linear track	1	Daedel	404 series
Servo motor	1	Parker-compumotor	CM231AR-01015
Servo motor controller	1	Parker-compumotor	APEX6151
Thermally conductive epoxy		Tra-con	2151
Square hollow tubing		Small tube products	
Coating for square tubing		Essex express	Dupont kapton
Rb sources	2	SAES getters	Rb/NF/3.4/12 FT10+10
Ion pump	1	Varian	Starcell VacIon plus 40
Ion pump controller	1	Varian	Midivac
Turbo pump	1	Varian	V70LP
Oil-free diaphragm pump	1	Varian	MDP12

ACKNOWLEDGMENTS

We acknowledge useful conversations with the other members of the JILA BEC collaboration. This work is supported by the NSF and by NIST.

REFERENCES

1. M. Anderson, J. Ensher, M. Matthews, C. Wieman, and E. Cornell, *Science* **269**, 198 (1995).
2. K. Davis, M. O. Mewes, M. Andrews, N. van Druten, D. Durfee, D. Kurn, and W. Ketterle, *Phys. Rev. Lett.* **75**, 3969 (1995).
3. C. Bradley, C. Sackett, J. Tollett, and R. Hulet, *Phys. Rev. Lett.* **75**, 1687 (1995); *ibid.* **79**, 1170 (1997).
4. We thank two anonymous referees who have pointed out that the hypothetical “experimental physicist, regardless of discipline” referred to in our introduction may find some sections of this paper obscure. The operating principles and basic lab techniques of wave plates, for example, are not covered in our paper. Ideally, this paper should be read in conjunction with a handbook on modern laboratory techniques in optical spectroscopy. Unfortunately, we are not aware of any one book that really fits the bill. In the catalogs or on the websites of many of the companies that sell optical components one may find helpful user application notes on AOMs, waveplates, etc.
5. E. Raab, M. Prentiss, A. Cable, S. Chu, and D. Pritchard, *Phys. Rev. Lett.* **59**, 2631 (1987).
6. A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, *Phys. Rev. Lett.* **54**, 2596 (1985).

7. N. Masuhara, J. M. Doyle, J. C. Sandberg, D. Kleppner, T. J. Greytak, H. Hess, and G. P. Kochanski, *Phys. Rev. Lett.* **61**, 935 (1988).
8. C. Myatt, N. Newbury, R. Ghrist, S. Loutzenhiser, and C. Wieman, *Opt. Lett.* **21**, 290 (1996).
9. W. Phillips and H. Metcalf, *Phys. Rev. Lett.* **48**, 596 (1982).
10. M. Greiner, I. Bloch, and T. Esslinger, *Phys. Rev. A* **63**, 031401(R) (2001).
11. W. Hansel, P. Hommelhoff, T. W. Hansch, and J. Reichel, *Nature* **413**, 498 (2001).
12. M. D. Barrett, J. A. Sauer, and M. S. Chapman, *Phys. Rev. Lett.* **87**, 010404 (2001).
13. S. R. Granade, M. E. Gehm, K. M. O'Hara, and J. E. Thomas, *Phys. Rev. Lett.* **88**, 120405 (2002).
14. This manuscript is a considerably revised version of a chapter from H. J. Lewandowski, *Coherences and Correlations in an Ultracold Bose Gas*, Ph.D. thesis, University of Colorado (2002).
15. E. A. Cornell, J. R. Ensher, and C. E. Wieman in *Proceedings of the International School of Physics—Enrico Fermi*, IOS Press (1999), p. 15; cond-mat/9903109.
16. H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Springer-Verlag (1999).
17. C. Wieman and L. Hollberg, *Rev. Sci. Instrum.* **62**, 1 (1991); **113**, 151 (1998).
18. K. MacAdam, A. Steinbach, and C. Wieman, *Amer. J. Phys.* **60**, 1098 (1992).
19. C. Wieman, G. Flowers, and S. Gilbert, *Amer. J. Phys.* **63**, 317 (1995).
20. W. Ketterle and N. van Druten, *Adv. At. Mol. Opt. Phys.* **37**, 181 (1996).
21. W. Ketterle, D. Durfee, and D. Stamper-Kurn, in *Proceedings of the International School of Physics—Enrico Fermi*, IOS Press (1999), p. 67.
22. B. DeMarco, *Quantum Behavior of an Atomic Fermi Gas*, Ph.D. thesis, University of Colorado (2001).
23. Acousto-optic modulators, or AOMs, are small nonlinear devices that allow one in essence to mix a phonon together with another photon, and to produce an outgoing photon with modified energy and k -vector. A transducer is mounted on the side of a transparent crystal, and launches a large-amplitude ultra-sound wave across the crystal, usually in the frequency range of 40 400 MHz. The sound-wave looks like a Bragg grating to the incoming light, and the light undergoes Bragg diffraction from the grating. The $+1$ order diffraction peak corresponds to the light deflecting away from the transducer (picking up the energy and the momentum of the phonon) and the -1 order diffraction peak corresponds to the light deflecting towards the transducer (giving up the energy and momentum into the phonon field). By inserting the crystal into the light beam and tilting it slightly from side to side while observing the transmitted intensity pattern on a IR disclosing card, one can readily identify the different diffraction orders (the 0 order mode is the one that persists when the rf power driving the transducer is disconnected. $+2$ and -2 modes are sometimes observed as well) and optimize the intensity diffracted into the desired mode. Because the transducer is driven by radio-frequency power, physicists familiar with rf technology may find it easiest to think of the AOM as a mixer that takes as its inputs one electromagnetic wave in the 100 MHz range, and one in the 400 THz range, and generates sum and difference frequencies. The outgoing frequencies are diffracted in slightly different directions, so one can readily put up opaque blocks to absorb all but the desired mode, say the sum frequency. In this mode the AOM-mixer can act as a fast optical switch.
24. On request, one of the authors (Eric Cornell ecornell@jilau1.colorado.edu) could provide various supporting materials such as image processing software and circuit diagrams to parties seriously considering building an apparatus similar to the one described here.
25. Viewpoint USA has recently begun to produce a similar board (DIO-64) with 64 outputs and an improved resolution of 50 ns. This new board is PCI rather than ISA and should be compatible with the existing software drivers written for the DIO-128.
26. Trade names are used here for identification purposes only and do not constitute an endorsement by the authors or their institutions.
27. J. H. Moore and M. A. Coplan, *Building Scientific Apparatus*, 2nd ed., Perseus Books (1991).
28. K. E. Gibble, S. Kasapi, and S. Chu, *Opt. Lett.* **17**, 526 (1992).

29. M. Stephens and C. Wieman, *Phys. Rev. Lett.* **72**, 3787 (1994).
30. P. D. Lett, R. N. Watts, C. E. Tanner, S. L. Rolston, W. D. Phillips, and C. I. Westbrook, *J. Opt. Soc. Amer. B* **6**, 2084 (1989).
31. W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, *J. Opt. Soc. Amer. B* **11**, 1332 (1994).
32. C. G. Townsend *et al.*, *Phys. Rev. A* **52**, 1423 (1995).
33. W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard, *Phys. Rev. Lett.* **70**, 2253 (1993).
34. E. Hecht, *Optics*, 2nd ed., Addison–Wesley (1990).
35. W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, *Phys. Rev. Lett.* **74**, 3352 (1995).
36. H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, *Phys. Rev. Lett.* **88**, 070403-1 (2002).
37. J. M. McGuirk, H. J. Lewandowski, D. M. Harber, T. Nikuni, J. E. Williams, and E. A. Cornell, *Phys. Rev. Lett.* **89**, 090402 (2002).
38. D. M. Harber, H. J. Lewandowski, J. M. McGuirk, and E. A. Cornell, *Phys. Rev. A* **66**, 053616 (2002).
39. E. Burt, R. Ghrist, C. Myatt, M. Holland, E. Cornell, and C. Wieman, *Phys. Rev. Lett.* **79**, 337 (1997).
40. D. Pritchard *et al.*, in *Proceedings of the 11th International Conference on Atomic Physics*, S. Haroche, J. C. Gay, and G. Grynberg (eds.), World Scientific, Singapore (1989), pp. 619–621.
41. C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions*, Wiley-Interscience (1992).