

Chapter 3

Experimental Apparatus I



3.1 Overview

Now we turn our attention to the experimental apparatus. In this chapter, we will confine our discussion to the basic parts of the apparatus that were necessary to carry out the experiments in Chapter 4. For the later experiments on chaos-assisted tunneling (Chapter 6), several major improvements were made to the apparatus, and these changes are discussed in Chapter 5.

The apparatus that we discuss here includes several parts. We begin by discussing the various laser systems involved in the experiment. Two diode lasers provided the light that trapped, cooled, and performed the measurements on the cesium atoms, and they are discussed separately in Sections 3.2 and 3.3. We will not discuss the general operation of a magneto-optic trap (MOT), as these traps have become quite commonplace (a nice discussion can be found in [Wieman95]); rather, we will give the specifics of our implementation and discuss how they impact the analysis of the experimental results. The optical lattice, which was the heart of the experiment, was provided by a home-built Ti:sapphire laser system, which is described in Section 3.4. For these experiments to work, the atoms had to be extremely well isolated from their surroundings, so they were trapped in an ultrahigh-vacuum system, detailed in 3.5. Finally, we discuss the experimental measurement procedure in Section 3.7 and the electronics needed to control the experiment in Section 3.8.

We will concentrate mostly on the salient features of the experiment, since much of the information in this chapter has already been covered in the dissertation of Bruce Klappauf [Klappauf98c].

3.2 DBR Laser

The basic operation of laser cooling in a MOT requires stabilized laser light tuned slightly to the red of an atomic transition. In the experiments here, we used the light from a distributed Bragg reflector (DBR) laser diode (model SDL-5712-H1, manufactured by SDL, Inc., which is no longer available as of the time of writing), which drove the $F = 4 \rightarrow F' = 5$ hyperfine component of the cesium D_2 transition. The advantage of a DBR laser is that it incorporated

a frequency-selective element that narrows its line width and causes it to lase near the desired wavelength (852 nm, in our case). The frequency of the laser is coarsely tuned by changing the diode temperature, and fine-tuning is achieved by adjusting the injection current. This particular laser nominally produced 100 mW of power with 180 mA of current, and incorporated a thermoelectric cooler and thermistor within the TO-3 package for convenient temperature stabilization.

3.2.1 Construction and Operation

The current and temperature controllers for the DBR laser were nuclear instrument module (NIM) style units that were designed and produced by Leo Hollberg's group at NIST-Boulder. These units worked extremely well, providing quiet current sources with reliable protection for the laser from electrical transients. The diode laser was also protected from electrical transients with several "protection diodes" in the cable connected to the laser diode, as suggested by [MacAdam92]. The protection-diode circuit, placed only a few inches from the laser diode, comprised a fast 1N5711 diode connected across the laser diode, but in the opposite direction, to protect against reverse voltages; and four 1N914 diodes in series connected across and in the same direction as the laser diode, to protect against overvoltages. The DBR diode was mounted on a commercial, air-cooled heatsink mount (model SDL-800 heatsink, with a model SDL-800-H adaptor plate, both by SDL, Inc.). This mount was bolted to a stainless steel post, which was in turn clamped to the optical table. The laser light was collimated by an antireflection-coated lens, Rodenstock model 1403-.108-.020, which had a focal length of 5 mm and a numerical aperture of 0.5. The lens was originally mounted on a three-dimensional translation stage, for precise and flexible positioning. However, thermal creep in this setup necessitated a nearly daily adjustment of this lens to keep the beam line correctly aligned. To solve this problem, we implemented a method used in the ultracold atom group at the Laboratoire Kastler Brossel in Paris, illustrated in Fig. 3.1 (a photograph of the DBR laser assembly is also shown in Fig. 3.2). The lens, which came mounted in a cylindrical aluminum case, was bonded to the end of a rod using a small drop of five-minute epoxy. The rod was in turn attached to a rotation stage and then an x-y-z translation stage for precise positioning. The lens was then adjusted into the desired position

over a V-grooved block that was attached directly to the laser diode heatsink mount. Two glass rods were then dropped gently onto the lens and V-groove, effectively filling in the space between them. This assembly was then bonded together by running beads of 24-hour epoxy using a syringe along the contact lines of the glass rods. After the epoxy was set, the bond for the positioning setup was broken. The setup (after the collimating lens) required minor adjustments for a few weeks as the bonds settled, but later the beam line became very stable, often remaining aligned for months without adjustment.

After collimation, the light was reshaped by an anamorphic prism pair (by Melles Griot) to be nearly circular, although the astigmatism of the beam required a cylindrical-lens telescope later in the beam line to maintain circularity. The beam then passed through two optical isolators (Conoptics, Inc. model 713), for 77 dB measured total isolation, because the DBR laser was especially sensitive to back reflections (compared to the grating-stabilized laser). This isolation came at a cost of 20% of the optical intensity. Another 10% of the light was split off for the saturated absorption setup used to frequency-lock the laser, as described below, leaving about 73 mW of power in the main beam. The beam was then double-passed through a tunable 60-100 MHz acousto-optic modulator (AOM), an IntraAction Corp. model ATD-801AL2. Since the laser was locked 195 MHz to the red of the $F = 4 \rightarrow F' = 5$ “cycling” transition, the light after the double-pass setup could be tuned from -75 MHz to $+5$ MHz relative to the cycling transition. We used a detuning of -15 MHz for normal (loading) trap operation. This AOM also controlled the intensity of the trapping light. The beam was focused through a $25 \mu\text{m}$ pinhole to clean its spatial profile, leaving about 27 mW of power. The diverging beam after the pinhole was collimated, with a beam waist parameter $w_0 = 11$ mm, to ensure a uniform illumination of the atoms. This light was split with two 2” diameter beamsplitters (66% and 50%) into three beams of equal power. Each of these beams passed through a 1.5” diameter quarter-wave plate before entering the vacuum chamber that contained the MOT, to give the beams the appropriate circular polarization for proper MOT operation. The beams were retroreflected through another set of quarter-wave plates after they exited the chamber. This setup provided the usual six-beam $\sigma^+ - \sigma^-$ molasses configuration, and together with the anti-Helmholtz magnetic field

coils (described below), provided the usual six-beam MOT configuration. A diagram of this beam line, along with the other two laser setups, is shown in Fig. 3.3.

3.2.2 Saturated Absorption Spectroscopy

The 10% of the beam that was picked off before the AOM was used to perform saturated absorption spectroscopy [Pappas80; MacAdam92] on the light for frequency stabilization. The thermal motions of cesium atoms in a vapor cell smear out the optical resonances, and the resulting Doppler line width [Verdeyen95] of cesium at room temperature is around 0.4 GHz. This width is much larger than the 5 MHz natural width of the cesium hyperfine resonances, and is even larger than the splittings between the excited hyperfine levels, so that these states cannot be resolved by simple absorption spectroscopy. The idea behind saturated absorption spectroscopy is that the atoms are illuminated by two counterpropagating beams of the same frequency (a “pump” and a “probe”), where the intensity of one of the beams (the probe) is monitored. The Doppler-broadened line will be apparent as usual, but at the resonance frequencies, the two beams will pick out the zero-velocity atoms, because only these atoms will be simultaneously




Figure 3.1: Diagram of DBR laser diode and collimation lens hardware. Only the front plate of the mounting heatsink is shown. A custom aluminum block with a V-groove was bolted to the front of the mounting plate. The collimation lens, which was mounted in a cylindrical aluminum housing, was held in place above the groove by two glass rods. The collimation lens, glass rods, and V-groove block were bonded together with 24-hour epoxy, resulting in a very stable configuration.

pumped by both beams (for moving atoms, the two beams will be Doppler-shifted to two different frequencies). The spectral manifestation of this effect is that there are narrow “Lamb dips” in the Doppler-broadened background, because these atoms are more saturated at certain frequencies and thus are more transparent to the probe beam. The Lamb dips occur at the atomic hyperfine resonances and also at “crossover resonances,” which occur halfway between pairs of hyperfine resonances (in these cases, the two beams pick out atoms moving such that the two Doppler-shifted beams excite two different hyperfine transitions, and thus still oversaturate the atoms). The widths of the Lamb dips are affected by many factors, including beam alignment, beam intensity, and cell pressure, but the width of the dips can be made comparable to the natural line width.

The saturated absorption spectrum thus provided high-resolution, frequency-dependent spectral features by which the DBR laser frequency could be determined. In the sense of




Figure 3.2: Photograph of the DBR laser, showing the laser itself inside the gold-toned TO-3 package, the mounting heatsink, and the collimation lens assembly.

actively stabilizing the frequency of the laser, though, these dips are not so convenient, because it is necessary to lock to a sloping part of a spectral feature. It is possible to lock to the side of a Lamb dip, but then the laser frequency would be sensitive to the width of the dip, which could drift in time. A more convenient situation arises for a dispersive (as opposed to absorptive) line shape, where the maximum slope occurs at the center of the resonance. One way to obtain such a shape is through frequency-modulation (FM) spectroscopy [Bjorklund83]. In this technique the frequency of the laser light is modulated, and phase-sensitive detection produces the dispersive line shape. In our setup, the probe beam was modulated at 11 MHz using a Conoptics model 350-40 electro-optic phase modulator (EOM) before passing through a 75 mm long cesium vapor cell (made by Environmental Optical Sensors, Inc.), where it overlapped the counterpropagating pump beam. The probe was monitored with a New Focus model 1801 low-noise photodiode, and the photodiode signal was mixed with the phase-shifted signal that drove the EOM. Additionally,

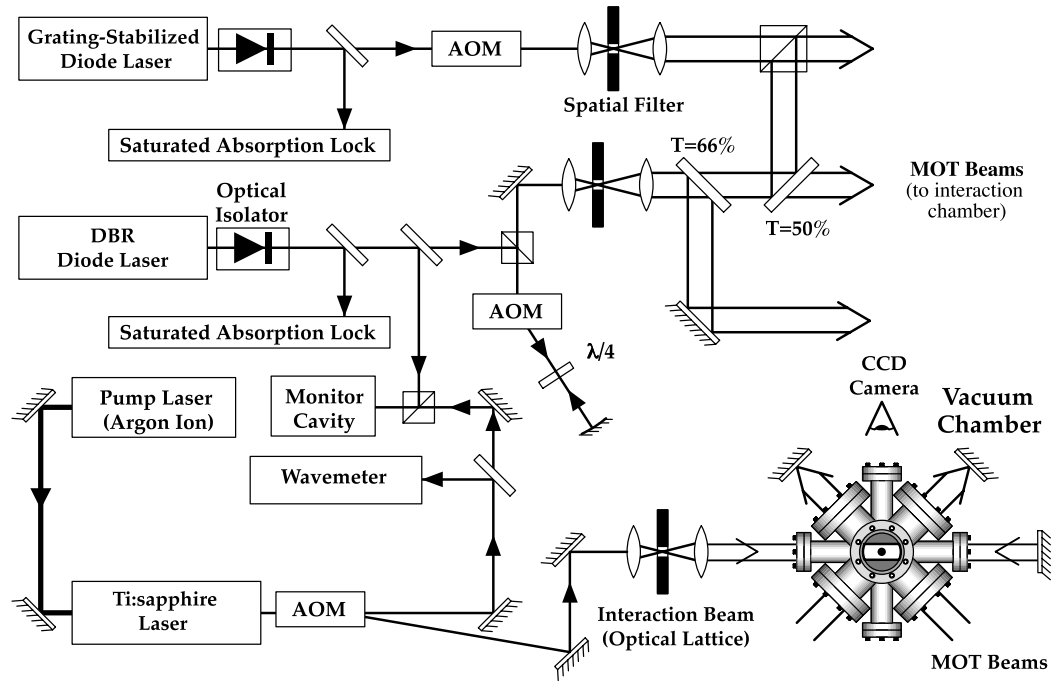


Figure 3.3: Layout of optical table, showing the schematic beam paths of the DBR diode laser (which provided the MOT trapping light), the external cavity diode laser (which provided the repumping light for the MOT), and the Ti:sapphire laser (which produced the optical lattice for the time-dependent interaction).

the pump beam was double-passed through a 70 MHz AOM, which was chopped on and off at 50 kHz. The mixed-down photodiode signal was then processed by a Stanford SR510 lock-in amplifier, with this chopping signal as the reference. The consequences (beyond the usual noise immunity of lock-in detection) of this pump-beam modulation were twofold: first, the spectrum was shifted to the red by 70 MHz, because the pump and probe beams were 140 MHz apart, and second, the Doppler contribution to the spectrum was suppressed, because only the Lamb dips (and not the Doppler-broadened absorption shape) were affected by the presence or absence of the pump beam. A spatial filter with a $75\ \mu\text{m}$ pinhole placed in the input beam of this setup greatly enhanced the quality and reproducibility of the spectrum by reducing fringes on the DBR laser beam, which had a relatively poor spatial mode quality. The measured spectrum for our setup is shown in Fig. 3.4. Because of the tunability requirements of the MOT trapping light, and the fact that the $F = 4 \rightarrow F' = 4, 5$ crossover transition was the largest and cleanest feature in the spectrum, we locked the DBR laser to this crossover and shifted the

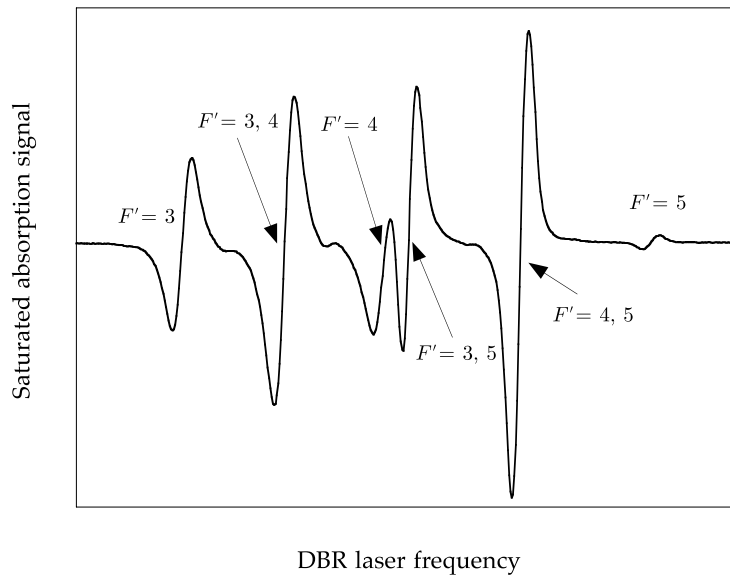


Figure 3.4: Saturated-absorption spectrum for the cesium D_2 line $F = 4 \rightarrow F'$ hyperfine manifold, as measured by the DBR laser setup described in the text. The three excited-state hyperfine levels coupled to the $F = 4$ ground state are visible along with the three crossover transitions (labeled by pairs of quantum numbers) as dispersive resonances. The DBR laser was actively locked to the $F' = 4, 5$ crossover resonance during normal MOT operation.

laser frequency closer to the MOT cycling transition (before the beam entered the chamber) as described above. The laser was locked to the crossover resonance using a standard P-I feedback loop to the laser injection current. This detection and locking method was considerably more complex than the method implemented for the repumping laser diode, but the advantage of this setup was that the main beam that went to the chamber was not frequency-modulated. After active locking, the line width of the DBR laser was on the order of 1 MHz, which is not nearly as narrow as is possible with other (e.g., grating feedback) stabilization techniques. It may have been possible to improve this line width electronically by using a much faster feedback loop, but this setup was certainly sufficient for the cooling and trapping of atoms. More details, as well as schematic diagrams, of this saturated absorption setup can be found in [Klappauf98c].

3.3 Grating-Stabilized Diode Laser

Although the $F = 4 \rightarrow F' = 5$ trapping transition is in principle a closed cycling transition (because the $F' = 5$ state decays only to $F = 4$), the trapping light can off-resonantly excite one of the other excited hyperfine levels, which could then decay to the $F = 3$ “dark state,” in which case the trapping laser would no longer cool the atom. The DBR laser alone was therefore insufficient on its own to trap cesium atoms. To address this problem, a “repumping” laser, tuned to the $F = 3 \rightarrow F' = 4$ transition, was used to return atoms in the dark state to the cycling transition. For this repumping light, we employed a home-built, grating-stabilized (pseudo-external cavity) diode laser system [Wieman91; MacAdam92]. This laser system was constructed in the Littrow configuration, where a grating was placed in front of the laser such that the first-order diffracted beam was reflected back into the diode, while the zeroth-order light was taken as the laser output. The grating thus served as a wavelength-selective component, which narrowed and controlled the lasing frequency.

3.3.1 Construction and Operation

A diagram of the Littrow laser is shown in Fig. 3.5, and a photograph of the assembled laser is shown in Fig. 3.6. The diode laser was a model 5421-G1 laser by SDL, Inc., which is housed in a

standard 9 mm package, and nominally produced 150 mW with 200 mA of injection current. The diode was mounted in a Thorlabs model LT230P-B collimation tube, which included a collimating lens with focal length 4.5 mm and numerical aperture 0.55. The collimation-tube assembly was mounted in a structure made of 954 aluminum bronze, which also supported the diffraction grating. This material had high strength and resistance to fatigue, which was important because the adjustability of the system relied on the flexing of thin sections of the structure, and the material also had reasonably good thermal conductivity, which facilitated temperature stabilization.

The diffraction grating was a $1/2'' \times 1/2''$ section of an inexpensive Edmund Scientific E43,005 grating, which was $3/8''$ thick and blazed for 500 nm (17°). We had a 500 \AA gold coating (over a 60 \AA chromium coating to facilitate adhesion) evaporated onto the stock aluminum coating to improve the reflection efficiency. At the Littrow angle (31°), the coated grating diffracted 21% of the input power into the first (retroreflected) order and 67% into the zeroth-order beam. The grating angle was set by two New Focus 9300-series fine-adjustment screws, whose ball-ends contacted sapphire windows for precise and stable alignment. Additionally, the horizontal grating angle was controlled by a piezo stack, which provided electronic tuning capability. The piezo stack was assembled from three American Piezo Ceramics piezo discs, which were 8 mm in diameter and 2.54 mm thick. The stack had a maximum voltage of 1500 V, a capacitance of 39 pF, and a displacement of $1.35 \mu\text{m}$ at 1 kV (the laser tuning rate was $12 \text{ GHz}/\mu\text{m}$ of displacement), and the stack was driven by a Trek model 601B-4 high-voltage (1 kV) amplifier, which had the upsetting habit of periodically self-destructing. Although we designed the pivot point of the horizontal adjustment to be near the optimal location [McNicholl85], the laser could only be scanned for 1 GHz frequency intervals without mode hops. This restricted tunability was likely due to a competition between the pseudo-external cavity modes and the free-running laser diode modes. Some possible solutions would include using an antireflection-coated diode or scanning the injection current along with the grating angle to match the corresponding modes. However, this setup was certainly sufficient to scan continuously over the $F = 3 \rightarrow F' = 4$ excited-state manifold, and has operated for years without mechanical intervention.




Figure 3.5: Diagram of the grating-stabilized diode laser for driving the repumping transition for the MOT. The top view shows the main elements involved in the operation and tuning of the laser, while the bottom view shows how the laser assembly is mounted and thermally stabilized.

We used a second set of NIST current and temperature modules (as with the DBR laser) to control the laser, with the same protection diodes as in the DBR laser. The temperature of the aluminum bronze structure was monitored with a Fenwal Electronics 50 k Ω glass-bead thermistor, and a separate Analog Devices AD590 temperature sensor provided the panel reading on the NIST temperature controller. The bronze structure was attached to an aluminum base plate with nylon screws, and a Melcor CP1.0-127-05L thermoelectric cooler was sandwiched between them to provide active temperature control over the bronze structure. Because the diode laser was operated near room temperature, the aluminum plate was an adequate thermal reservoir. A Lucite cover provided thermal and also some acoustical isolation from the surroundings, and a microscope slide at Brewster's angle allowed the light to escape from the housing. The base plate was mounted rigidly to the optical table with four 1" diameter stainless steel posts, so that mechanical resonances would occur only at high frequencies. A second aluminum plate was




Figure 3.6: Photograph of the grating-stabilized laser assembly, used for the MOT repumping light.

bolted directly to the bottom of the base plate with a layer of 1/8" thick Sorbothane viscoelastic damping material between the two, in order to damp any vibrations of the base plate.

As in the case of the DBR laser, the Littrow laser output beam was passed through an anamorphic prism pair to make the beam circular before it passed through a single Conoptics isolator. The beam then made a single pass through a 100.5 MHz AOM, which shifted the frequency to the blue (because the laser was again locked to a crossover transition rather than the desired line itself) and provided intensity control for the repumping light. The beam was then spatially filtered and expanded to nearly the same size as the MOT beams, and was combined with only the vertical MOT beam (i.e., the beam along the symmetry axis of the anti-Helmholtz coils, described below) by a 1.5" polarizing cube beamsplitter before passing through the large quarter-wave plate mentioned above. With a typical operating current of 93 mA, this setup provided about 40 mW of power after the grating, and about 16 mW of repumping light after the spatial filter, which was far above the necessary intensity for proper MOT operation.

3.3.2 Frequency Control

The frequency of this laser was also locked via a saturated-absorption setup, which was considerably simplified compared to the setup for the DBR laser. The optical setup we used here is essentially the same as the one described in [MacAdam92; Monroe92]. About 10% of the repumping beam was split off of the main beam before the AOM, which then passed through a 3/8" thick uncoated window. The reflections from the window provided a pair of probe beams with about 1% of the main-beam intensity, and passed through a cesium vapor cell (identical to the vapor cell used in the DBR setup). The remainder of the picked-off beam formed the pump beam, overlapping one of the probe beams in the vapor cell in a counterpropagating fashion. The intensities of the probes were detected and subtracted by a differential photodetector. The Lamb dips only appeared on the probe beam that overlapped the pump beam, while the Doppler absorption profile appeared as a common mode signal on the two beams, and was thus removed by the subtraction. Again, from the standpoint of active frequency locking, it was convenient to have a dispersive signal. To achieve this, we dithered the laser frequency by applying a small 12

kHz signal to the piezo stack, and the subtracted photodiode signal was analyzed by an EG&G model 5204 lock-in amplifier. The dither had the disadvantage that the main repump beam was modulated, but the frequency excursion was not more than a few MHz, and the MOT operation was relatively insensitive to the repump characteristics, because the atoms spent a relatively small fraction of their time out of the cycling transition. The measured saturated-absorption spectrum for this setup is shown in Figs. 3.7 and 3.8 (see also [Klappauf98c] for other details of this setup, including a diagram of the optics and electronics). This signal was then fed back to the piezo control amplifier, using another P-I lock circuit based on the design in [Monroe92]. The laser was locked to the $F = 3 \rightarrow F' = 3, 4$ crossover transition, which was the strongest feature in the spectrum, and the AOM shift of the main beam brought the repump light to the center of the $F = 3 \rightarrow F' = 4$ repumping transition resonance. With this fairly simple setup, the diode laser remained locked to the proper frequency for hours, provided the optical table was not bumped too severely, and only required occasional adjustment of the injection current to move the laser mode-hop points away from the hyperfine resonances.

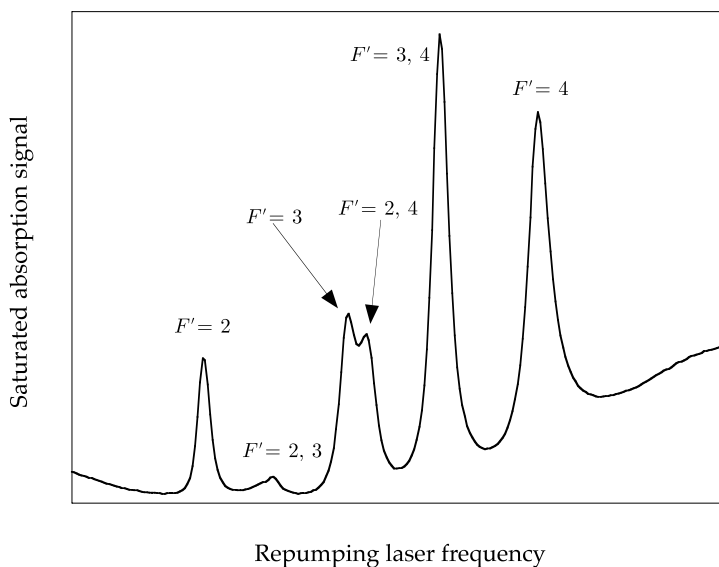


Figure 3.7: Saturated absorption spectrum for the cesium D_2 line $F = 3 \rightarrow F'$ hyperfine manifold, as measured by the grating-stabilized (repumping) laser setup described in the text. The signal plotted here is the output of the differential photodiode, so the resonances appear as absorptive line shapes. The spectrum after additional processing by the lock-in amplifier to extract the dither information is shown in Fig. 3.8.

3.4 Ti:sapphire Laser

The two diode lasers worked quite well for cooling and trapping atoms, but to realize a clean optical lattice, we required a high-intensity, far-detuned source of light. These requirements were met by a home-built Ti:sapphire laser. This laser, when pumped by a Coherent Innova 90 argon-ion laser (producing about 8 W of power in multi-line mode) produced 0.5 W of single-mode light, tunable several nm around 852 nm. This laser (along with the associated control electronics) was originally a dye laser, built by Patrick Morrow, and was similar to a dye laser used in the sodium-based apparatus in our laboratory [Fischer93; Robinson95a]. The dye laser and the subsequent conversion to a Ti:sapphire laser were based on designs by Jim Bergquist at NIST-Boulder.




Figure 3.8: Saturated absorption spectrum for the cesium D₂ line $F = 3 \rightarrow F'$ hyperfine manifold, as measured by the grating-stabilized (repumping) laser setup described in the text. The signal plotted here is the output of the lock-in amplifier, which extracts the frequency-dither information from the signal, so that the resonances are dispersive. This spectrum is otherwise the same as that shown in Fig. 3.7. During normal operation, the repump laser is actively locked to the $F = 3 \rightarrow F' = 3, 4$ crossover resonance.

3.4.1 Laser Design and Construction

The basic layout of the Ti:sapphire laser is shown in Fig. 3.9, and a photograph of the laser is shown in Fig. 3.10. The overall construction of the laser was a series of 1" thick aluminum plates mounted on four long 1" diameter Invar rods, providing a stable frame for mounting the laser components. The mirrors were mounted on plates that were in turn mounted on the main plates by extension springs and Bill Lees fine-adjustment screws.

The crystal itself was produced by Union Carbide. The crystal was doped with 0.05% titanium, with a guaranteed figure of merit ≥ 450 , and was 6 mm in diameter and 20 mm long, with Brewster-cut ends. The crystal was mounted in a split copper disk, which was in turn water-cooled. We found that the water cooling was necessary for efficient laser operation.

The cavity was a four-mirror, folded-ring design. The two mirrors on either side of the crystal (in the beam path) were the "pump mirrors," which allowed the argon-ion light in and out of the cavity. These mirrors were coated by VLOC (now a part of II-VI, Inc.), and were specified to transmit at least 90% at 488 and 514 nm, while maintaining at least 99.9% reflectivity at 852 nm. These two mirrors also had a 20 cm radius of curvature, while the other two mirrors were flat. The arrangement of the mirrors and crystal was such that the aberrations introduced by the Brewster surfaces and the off-axis curved mirrors cancelled [Dunn77], leading to clean Gaussian intensity profiles both inside the crystal and at the output coupler. The output coupler had a reflection coefficient of 97.3% (several other output couplers in the 90-95% range failed to give better output powers). The fourth mirror was small and mounted on a piezo stack to give rapid control over the cavity length. The cavity mode was focused to its tightest point inside the crystal (to a waist w_0 on the order of 40 μm), and the beam at the output coupler had a waist around 0.7 mm. The argon-ion beam was focused by a lens (175 mm focal length) to match the cavity mode inside the crystal, and the lens was tilted to compensate for the aberrations introduced by the Brewster-cut crystal end and the propagation through the slightly curved mirror (the axis of the tilt was perpendicular to the plane of the cavity mode).

The cavity also included several optical devices for stable, single-mode operation, some of which are visible in Fig. 3.11. An optical diode (optical rotator/Faraday rotator pair, manufac-

tured by Coherent Laser Group for their model 899 Ti:sapphire laser) forced unidirectional lasing by rotating the polarization of light propagating in the wrong direction. A birefringent filter, also manufactured by Coherent, provided coarse frequency tuning, making use of the frequency-dependent polarization shift of three birefringent plates. Single-mode operation was enforced by a Coherent intracavity assembly (ICA), which housed two etalons, a thick etalon and a thin etalon. The thick etalon, which had a free spectral range (FSR) of 10 GHz, was dithered and actively locked to maximize the laser power. The thin etalon, with a 225 GHz FSR, was used to select the longitudinal cavity mode of operation, in roughly 10 GHz steps. The laser frequency could be tuned continuously by a pair of Brewster plates, mounted on galvanometer drives, to vary the cavity length; the two windows were arranged to cancel deflections caused by their synchronous rotation. The short-term line width of the laser was around 10 MHz, as measured by a Fabry-Perot cavity with a similar resolution, with fluctuations at the 100 MHz level during a data run.

3.4.2 Laser Operation and Control

The laser could be actively locked to arbitrary frequencies by using a polarization-based spectroscopic technique [Hänsch80] to derive a dispersive lock signal from a 1.5 GHz FSR stable Invar Fabry-Perot cavity. The cavity resonance could be tuned continuously by adjusting the angle of a Brewster plate inside the cavity. In practice, though, the Ti:sapphire laser was stable enough that active frequency locking was unnecessary.

The frequency of the laser was monitored on a NIST LM-10 wavemeter (scanning Michelson interferometer), which measured the absolute laser wavelength by direct comparison to a known, stabilized HeNe reference laser. The resolution of the wavemeter was limited to about 500 MHz, so the Ti:sapphire beam was measured simultaneously with the DBR laser beam by a scanning “monitor” Fabry-Perot cavity with a 1.5 GHz FSR. The Ti:sapphire laser was tuned so that the resonances of the two lasers were aligned, and thus the frequency could be known to much better than 1.5 GHz. The Ti:sapphire beam intensity was controlled by an 80 MHz fast AOM (IntraAction model ATM-801A2-2), which shifted the frequency 80 MHz to the blue.




Figure 3.9: Diagram of the Ti:sapphire laser layout and beam paths. The structure of the laser consisted of thick aluminum plates mounted on four long Invar rods. Only the most basic optical elements are shown, including the four mirrors (shown in blue) and the Ti:sapphire crystal (shown in pink), where the beam path “jogs” in the right half of the diagram. The two mirrors adjacent to the crystal (in the sense of the beam path) transmitted the argon-ion laser pump light, which enters in the diagram from the upper right. The output coupler is the leftmost mirror. The jog in the uppermost leg of the beam path represents the location of the optical diode, which forced the lasing mode to oscillate in only one direction, indicated by the arrows. For a sense of scale, the uppermost leg of the beam path is about 64 cm long.

Since the DBR laser beam (before the AOM) was 195 MHz to the red of the $F = 4 \rightarrow F' = 5$ transition, and it was the zeroth-order beam of the Ti:sapphire AOM that was measured by the cavity, the frequency of the optical lattice seen by the cesium atoms was $n \times 1.5 \text{ GHz} - 115 \text{ MHz}$ relative to the MOT cycling transition, when the two resonances were overlapped on the monitor cavity. For the experiments in Chapter 4, we used $n = -4$, for a detuning of -6.1 GHz .

The first-order AOM beam was spatially filtered by focusing through a $50 \mu\text{m}$ pinhole. The beam was then collimated with a waist parameter of 1.5 mm and sent to the vacuum chamber. This beam size was significantly larger than the size of the MOT atom cloud (which was Gaussian, with $\sigma_x = 0.15 \text{ mm}$), to ensure uniform illumination of the atoms. The beam was retroreflected to form the optical lattice, using a mirror mounted directly to a vacuum-chamber window flange. This mounting method made the mirror position very rigid in the direction




Figure 3.10: Overall photograph of the Ti:sapphire laser. The view is from the pump input end of the laser, and the pump light input periscope is visible here. The argon-ion laser is also visible behind and to the right of the Ti:sapphire laser.

of the lattice, which was important because the retroreflector position determined the spatial phase (and hence the positional stability) of the lattice. We measured the phase stability of the optical lattice with a Michelson interferometer setup, and found that the phase fluctuations over the time scales of the experiments (2 ms) were 8% of the lattice period, with typical ringing frequencies of around 600 Hz. This number is a generous upper bound, since the noise in the measurement was probably dominated by the motion of the beamsplitter and reference mirror. These optics were mounted on a Newport model 45 damped post and were hence not nearly as rigidly fixed in the direction of the standing wave. The dominant source of the phase instability was vibration from the cooling water flowing through the argon-ion laser, which was transmitted to the chamber and retroreflector through the optical table. The short-term (i.e., over a single 2 ms interaction with the atoms) noise on the intensity of the laser was measured to be 0.2% (rms) of the average intensity, with longer-term (i.e., on the order of seconds) fluctuations at the same




Figure 3.11: Overhead view of the interior of the Ti:sapphire laser, showing a clear view of the input light as it enters the crystal. Several of the frequency-stabilization components are also visible here.

level. (Note that the intrinsic noise of the New Focus model 1801 low-noise photodetector used for this measurement was of the same order as the laser intensity noise, and was thus subtracted to obtain the reported numbers.) The short-term fluctuations were dominated by an oscillation at 2.4 kHz, due to the dither of the thick etalon at the same frequency. The laser power also drifted by several to many percent over the course of hours as the room temperature drifted. For the experiments in Chapter 4, where the data were acquired over a few days, we manually kept the power at approximately the same level. For the experiments in Chapter 6, where the data were acquired in runs that took place continuously over more than a week at a time, a photodiode after the spatial filter monitored the laser intensity, and the computer scaled the control signals to compensate for laser drifts at the beginning of each run (i.e., about every 15-30 minutes). Also, if the power drifted by more than 1% from the initial power during a single data run, the entire run was discarded. This procedure was important, as the laser intensity was significantly different from the day to the night because of the difference in ambient temperature.

3.4.3 Intensity Calibration

As the local laser intensity and detuning determined the optical potential experienced by the atoms, it was important to have an accurate method for calibrating the laser intensity. Here we will discuss the calibration method used for the experiments in Chapter 4. The beam power was measured using a Newport model 1825C power meter with a model 818-SL semiconductor detector (and a model 883-SL attenuator). The power meter had a NIST-traceable calibration, with a nominal absolute uncertainty of 2%. The intensity used for the well-depth calculation was reduced by 8% from this measured value to account for losses due to the vacuum chamber windows and other optics in the beam line after the point where the power was measured.

The beam diameter was measured using both a knife-edge method and an imaging method using a CCD camera. Both methods measured the beam profile approximately 2 m before the beam entered the chamber. The knife-edge method measured only the beam profile as the beam propagated to the chamber, while the CCD camera method measured the beam propagating both to and from the chamber by using beam splitters on kinematic mounts to pick

off a sample of the beam for the camera. The knife-edge method involved scanning a knife edge across the beam using a micrometer-driven translation stage. The detector placed behind the knife edge thus measured the total intensity integrated over a half-plane. As the beam was quite Gaussian after the spatial filter, the detected signal could be well fit to an error-function model to yield the beam-waist parameter w_0 in the horizontal and vertical directions. In the CCD camera method, the picked-off beam passed through several attenuators and then into a Sony model XC-77 CCD video camera. The camera output was digitized using a ComputerEyes/RT SCSI frame grabber by Digital Vision, Inc., and the beam intensity profile was fit directly to a Gaussian model to obtain the beam-waist parameters. Some interference fringes due to reflections between the CCD surface and the protective window were visible on the beam, but did not significantly affect the measurement. The CCD images were valuable in ensuring that the beam spots were circular and nearly the same size before and after the chamber. Because the fitted spot sizes depended slightly on the beam intensity in the CCD method, this method was used only to obtain the relative spot sizes of the beam before and after the chamber, and the knife-edge measurement was used to set the absolute scale of the measurements. The local intensity used for well-depth calculations was also corrected for the specific widths of the MOT and laser beam size. The correction factor is

$$\eta_{\text{MOT}} = \frac{w_0^2}{4\sigma_x^2 + w_0^2} , \quad (3.1)$$

where σ_x is the spatial standard deviation of the Gaussian MOT profile. This correction represents the average intensity experienced by the atoms relative to the intensity experienced by an atom at the center of the beam. For typical values of $\sigma_x = 0.15$ mm and $w_0 = 1.5$ mm, this correction factor is about 96%. Combining these measurements yielded an absolute potential amplitude that was correct to better than 10%. The oscillations in the diffusion rate $D(K)$ as a function of the laser intensity, as described in Chapter 4, provided an independent check of this calibration.

3.5 Vacuum System

The ultrahigh vacuum system, shown in Figs. 3.12 and 3.13, had as its main feature a large 10-way stainless steel cross, custom-fabricated by HPS (a division of MKS Instruments). The cross had six 4.25" ConFlat (CF) flanges, arranged along the three major axes to accept the large MOT trapping beams. Along the equatorial plane of the chamber, there were four 2.75" CF flanges between the four 4.25" flanges. Two of these smaller flanges provided access for the optical lattice, a third provided access for the main imaging system, and the fourth provided a path to the vacuum pump and cesium source. The flanges providing optical access were covered with zero-length (Kovar-sealed) glass viewports, which were antireflection (AR) coated by VLOC to have $\leq 0.25\%$ reflection per surface. The glass cells used in previous experiments [Robinson95a] have the advantages of compact size and rapid magnetic field switching. The steel chamber




Figure 3.12: Photograph of the vacuum chamber. The rear section of the chamber is visible here, including the ion pump and cesium ampule. The main (trapping) chamber section is also visible, along with the anti-Helmholtz (AH) and Helmholtz (HH) coils.

used here, on the other hand, has the advantages of AR-coated windows, which largely suppress interference fringes on the beams, and optical access for large beams, like the MOT trapping and repumping beams. The chamber was mounted on a 5/8" thick aluminum plate (later changed to a G10 glass phenolic plate), which was in turn attached to the table via four 6" long stainless steel posts. All the viewport flanges had additional 1/4-20 tapped holes for mounting optics directly to the chamber.

A 20 l/s ion pump (Varian model 919-0236) maintained the vacuum in the chamber, and an HPS model 10000 5836 nude Bayard-Alpert ion gauge provided a means to monitor the chamber pressure (in addition to the ion pump current). From the ion-pump current, the vacuum pressure was around 8×10^{-8} , although it is likely that the actual pressure was substantially lower. With the proper magnetic bias field adjustment, the atoms could be released from the MOT by turning off the anti-Helmholtz coils (but not the optical molasses), and a substantial




Figure 3.13: Photograph of the main (trapping) section of the vacuum chamber. This view looks into one of the large MOT trapping beam windows and the smaller optical lattice input window. The anti-Helmholtz and Helmholtz coils are more visible in this picture.

fraction could still be recaptured after 4 s by turning on the anti-Helmholtz fields again. With the measurements in [Monroe92], a lifetime of 4 s implies a pressure of 2×10^{-9} torr (assuming that cesium vapor dominated the background pressure). The readings from the ion gauge and ion pump were probably anomalously high due to cesium contamination.

We introduced cesium vapor into the chamber from a 1 g sample in a glass ampule, which was attached via a small CF flange (and a glass-metal seal) to a series pair of all-metal bakeable UHV 1.5" diameter vacuum valves (Varian model 951-5027). The inner valve was always fully open, and we installed it so that we could remove the cesium ampule without breaking vacuum in the main chamber by sealing both valves and breaking the CF seal between them. The second valve was usually closed hand-tight, and opened by a few turns to increase the cesium pressure in the main chamber whenever necessary (typically only once per week). At first, though, we heated the cesium ampule to 100°C for several days in order to build up enough vapor pressure to see trapped atoms in the MOT. This procedure coated the surfaces of the main chamber with cesium until it no longer acted as an effective pump for the cesium vapor. In an earlier attempt, we used a UHV variable leak valve (Varian model 951-5106) to introduce the cesium, but this valve did not have sufficient conductance to maintain a reasonable quantity of cesium in the chamber, even with substantial heating of the ampule. By contrast, the large-valve system required no temperature control of the cesium ampule.

To achieve good vacuum, we baked the chamber for three days at a maximum temperature of 220°C, with an additional day each for the gradual ramping up and down of the temperature. During the bake, the ion pump was off, but the chamber was pumped by a 70 l/s turbo pump through one of the all-metal bakeable valves. After this main bake, the second valve and cesium ampule were installed, and this subsection was separately baked for several days at a similar temperature while being pumped by the turbo pump through a pinch-off tube. After the second bake, we broke the ampule's inner glass seal using the included breaker and a strong magnet. Then the seal to the turbo pump was pinched off and we opened the valve to the main chamber.

3.5.1 Magnetic Field Control

The trapping action of a MOT relies on the presence of magnetic-field gradients in the trapping region. In our experiment these fields were generated by a pair of coils in the anti-Helmholtz configuration (where the coils form a mirror-image pair on either side of the atoms, but the currents in the two coils flow in opposite directions). The circular coils were 6.2" in diameter and each was wound from 202 turns of 24 gauge, Kapton-coated copper magnet wire onto a water-cooled aluminum form. The layers of the coil windings were fixed in place and further electrically insulated with clear fingernail polish. After winding, the coils measured 8.7Ω and 2 mH each. The aluminum forms were attached to the chamber by Lucite clamps that gripped the two vertical flanges, such that the separation between the coils was about half their diameter. At a typical operating current of 3 A, the coils yielded a gradient of 11 G/cm along the symmetry axis of the coils (and thus half this value perpendicular to this axis, which follows from cylindrical symmetry and $\nabla \cdot \mathbf{B} = 0$). Although the coils generated 80 W of heat, they remained cold to the touch with 10°C water flowing through the forms at 10 gallons/hour.

The current through these coils was provided by two push-pull pairs of LM12 (National Semiconductor) high-current operational amplifiers, which were powered by a massive Sorenson DCR40-35A single-ended power supply. This circuit was capable of switching off the 3 A of current in the coils in 100 μ s, but due to eddy currents in the chamber and the coil forms, the major part of the magnetic field damped away exponentially with a much longer time constant of 3 ms. The residual magnetic fields were not of substantial concern for the experiments in Chapter 4. However, the Raman velocity selection technique described in Chapter 5 is very sensitive to magnetic fields, and with this setup we could detect changing magnetic fields even 350 ms after the coil currents were extinguished. This time scale was much longer than expected for eddy currents in the chamber or coil forms, and even persisted when the aluminum chamber mounting plate was replaced with a phenolic plate. Thus, this long field decay was likely due to the presence of ferromagnetic Kovar in the chamber viewports.

The chamber was also enclosed within three pairs of Helmholtz coils, which allowed the magnetic field at the center of the chamber to be nulled out. Two of these coil pairs also

provided the bias field for the optical pumping described in Chapter 5. These coils had a square profile, measured about 15” per side, and were separated by about 8” (with slight variations among the pairs so that they fit together properly and so that the coil spacings were 0.54 times the coil widths [Moore91]). Each of these coils had 44 turns of the same wire used for the anti-Helmholtz coils, with about 6 Ω of resistance per coil. These coil pairs produced about 2 G/A of magnetic field along their respective axes. For typical operating conditions, where around 0.8 G was needed in the vertical direction to compensate for the magnetic field of the Earth, and a combined 1.5 G was pulsed on in the horizontal direction (using two coil pairs) for optical pumping, no special cooling was needed for these coils. These coils were driven by circuits that were similar to the anti-Helmholtz drivers, but were based on the Burr-Brown OPA2544T dual high-current op-amp, with the two sections operating in push-pull mode on a series pair of coils. The push-pull arrangement permitted bidirectional operation of all the coils, even though all the drivers (including the anti-Helmholtz drivers) were driven from a common single-ended power supply.

3.6 Imaging System

All measurements of the atoms in our experiments were performed by a CCD camera (Princeton Instruments TE/CCD-5122TK/1UV). The CCD chip was an array of 512×512 pixels, with each pixel approximately a 20 μm square. For low-noise operation, the CCD was chilled to around -30°C . The light was imaged by a Nikon 105 mm f/2.8 D macro lens, which reduced the image in the plane of the atoms by 1:1.8. The camera viewed the atoms through one of the smaller viewports in the chamber, such that the camera view was orthogonal to the optical lattice.

The action of the camera shutter was quite audible, and we designed the camera mount to minimize the transmission of vibration to the lasers and other components on the optical table. The camera was attached via a short post to an aluminum plate, which was sandwiched between two sheets of 1/8” thick Sorbothane rubber. This stack was clamped between an aluminum base plate and a third aluminum plate, so that the camera was vibrationally isolated from the base. The base plate was then mounted on a Newport model 45 vibration-damping post. The post was

mounted on a large aluminum base plate that rested on the table through another intermediate layer of Sorbothane, and was clamped down with clamps padded with more rubber. With this setup, the shutter action vibrations were visible on the laser lock signals, but did not have a significant effect on the measurements.

The camera images were transferred to the Power Macintosh 7100 controlling computer via a Princeton ST-138 camera controller connected to a NuBus-style computer interface card. The data transfer took several seconds for large images, so for real-time alignment we used a separate and much less expensive Ikegami camera connected to a small CRT monitor. This camera viewed the MOT through the side edge of one of the MOT beam windows, so it also provided an independent view of the MOT.

3.7 Measurement Technique

With these experimental tools in hand, we can proceed with a discussion of the measurement methods for the data presented in the next chapter. The basic experimental sequence for these experiments is shown in Fig. 3.14. The sequence began by preparing a sample of cesium atoms by loading the MOT from ambient vapor. This trapping phase lasted for about 5 s. The atoms were cooled further by reducing the intensity of the trapping light to about 40% of its maximum value and increasing the detuning to 55 MHz (from the 15 MHz used during the initial trapping phase). This procedure resulted in an atomic sample that was approximately Gaussian in position and momentum, with typical sizes of $\sigma_x = 0.15$ mm and $\sigma_p/2\hbar k_L = 4$ (corresponding to a one-dimensional temperature of 12 μ K).

After this initial preparation, the trapping light and anti-Helmholtz fields were extinguished, and the atoms were exposed to the time-dependent optical lattice (in the form of a periodic pulse train, in Chapter 4). It was during this stage that the interactions of physical interest to us occurred. Since we tracked the evolution of the atoms as they were exposed to the optical lattice, this stage lasted from zero to about 2 ms. This time was limited by both the capabilities of the control electronics as well as the requirement that the atoms not fall significantly compared to the 1.5 mm waist of the lattice during this time. During this interaction, the




Figure 3.14: Schematic picture of the experimental sequence for the experiments in Chapter 4. (Graphics rendered by W. H. Oskay.)

optical lattice modified the atomic momentum distribution, typically heating the center-of-mass motion of the atoms through the optical forces.

The next stage in the experimental sequence was the measurement of the atomic momenta. All laser light was extinguished, and the atoms were allowed to drift in the dark for 15 ms (or 20 ms for the experiments in Chapter 6; these quoted drift times include the variable lattice interaction times, so the time spent in the dark depended on the duration of the lattice interaction). During this time, the typical atom moved a distance that was large compared to the initial diameter of the cloud. Thus, assuming that the initial sample was a point source, the displacements of the atoms after the drift times was proportional to their corresponding momenta, and so the spatial distribution of the atoms had the same form as the momentum distribution. (More carefully, the final spatial distribution had the form of the momentum distribution convolved with the initial spatial distribution). To detect the momentum distribution, the atoms were illuminated by the MOT trapping light, but without the anti-Helmholtz magnetic fields. In this configuration, the cooling effect of these optical-molasses beams “froze” the atoms in place, and the atomic fluorescence was recorded on the CCD camera. The imaging time was kept relatively short (10 ms for the experiments of Chapter 4, 20 ms for Chapter 6) so that the atoms did not move significantly during the measurement. To avoid an effectively nonuniform exposure of the CCD as a result of the shutter opening (5 ms) and closing (12 ms) times, we pulsed the freezing molasses light on only when the shutter was fully open.

The two-dimensional intensity distributions from the camera were therefore a direct measurement of the two-dimensional momentum distributions of the atoms. Since we only care about the atomic momenta in the direction of the optical standing wave, we integrated over the transverse direction to obtain one-dimensional distributions. An example of such a distribution is shown in Fig. 3.15. This measurement is the simplest type of experiment, where no optical-lattice interaction is applied to the atoms, and hence constitutes a measurement of the momentum distribution of the MOT (which is the initial condition for the experiments in Chapter 4). The momentum distribution in the trap was mostly Gaussian, with broadened tails. Similar distributions have been observed by other groups, especially for strong magnetic-field

gradients (e.g., [Petrich94; Townsend95]). A good model for this momentum distribution is an incoherent sum of a Gaussian and an exponential distribution:

$$\frac{\eta_g}{\sqrt{2\pi}\sigma_p} \exp\left(-\frac{p^2}{2\sigma_p^2}\right) + \frac{\eta_e}{2\xi_p} \exp\left(-\frac{|p|}{\xi_p}\right) . \quad (3.2)$$

The parameters for this model were determined from a best fit; the widths of the components are $\sigma_p/2\hbar k_L = 3.9$ and $\xi_p/2\hbar k_L = 13.0$, and the relative weights are $\eta_g = 82\%$ and $\eta_e = 18\%$. The measured distribution is plotted in Fig. 3.15 along with the best-fit function (3.2). To characterize the temperature of the atoms, we normally used a simple Gaussian model. When fitted to the same atomic distribution, this model accounts for 96% of the atomic population, with the remaining 4% augmenting the tails of the momentum profile. To minimize the convolution effects of the MOT cloud on the temperature measurements, we also measured the nearly Gaussian momentum distribution at several drift times. We could then extract the corresponding true σ_p by fitting Gaussian lineshapes to measurements at each drift time, and fitting these Gaussian




Figure 3.15: Plot of the momentum distribution in the MOT (solid line), with model distribution (3.2) (dotted line). Most of the atoms are in the main Gaussian component, while a smaller fraction is contained in the broad, non-Gaussian tails. This distribution corresponds to the initial conditions of our kicked-rotor experiments.

widths to the model

$$\sigma_x(t_{\text{drift}}) = \sqrt{(\sigma_{x0})^2 + (\sigma_p t_{\text{drift}}/m)^2} , \quad (3.3)$$

where $\sigma_x(t_{\text{drift}})$ is the measured spatial width parameter of the expanded MOT, σ_{x0} is the spatial width parameter of the initial MOT, and t_{drift} is the drift time. This model is the exact convolution if the MOT spatial and momentum distributions are Gaussian.

3.8 Control Electronics

We give here a brief overview of the control electronics, and a more thorough discussion will be published separately [Oskey01b]. For the experiments in Chapter 4, the entire experiment was controlled by a single Power Macintosh 7100/80 computer. All data acquisition and control software was programmed in LabVIEW. As mentioned above, the computer acquired the camera images via a NuBus interface card. The computer also used a National Instruments NB-MIO-16L-9 multipurpose I/O board. Although this board included many input and output functions, the functionality of this board was severely limited by the fact that its various sections could not be easily synchronized. This board was only used as a master trigger for the experiment and for control of the anti-Helmholtz fields, but precision timing was delegated to several external devices, including a Stanford Research Systems (SRS) DG535 timing/pulse generator and two SRS DS345 arbitrary waveform synthesizers. The computer programmed these timing devices through a National Instruments GPIB interface. The timing of various events was controlled with the DG535, while the optical lattice and MOT trapping beam intensities were controlled by the two DS345 synthesizers. To ensure precise timing, these three devices were slaved to a common frequency reference, and later all timing devices were slaved to a rubidium clock (LPRO model by EFRATOM) for extremely high absolute accuracy and stability.

For the more complex experiments in Chapter 6, the acquisition and control setup was upgraded considerably. Most of the control and data processing was handled by a fast Power Macintosh G4 (“Sawtooth”). However, because Princeton Instruments (now Roper Scientific) does not provide a PCI interface card for our camera controller, the old 7100 computer still acquired the camera images. The two LabVIEW processes on the two computers communicated via TCP.

The more complex optical lattice intensity waveform was moved to an Agilent 33250A waveform synthesizer which has much better temporal resolution and waveform memory (65, 535 samples with as fast as a 5 ns sampling rate, compared to the 16, 299 sample, 25 ns rate of the DS345). The new computer also handled many timing functions through a PCI-DIO-32HS digital interface by National Instruments, which outputs up to 32 channels of TTL data through 50 Ω drivers (at a clock rate of 1.25 MHz, in our setup). Other timing functions that required voltage levels other than TTL were handled by two DG535 units, and more complicated waveforms were generated by five DS345 synthesizers.

