Chapter 2

Cooling to the ground state of axial motion

2.1 In situ cooling: motivation and background

While measurements in Ref. [13] established a lifetime of 2–3 seconds for atoms trapped in the FORT, these were atoms "in the dark," i.e., only interrogated once after a variable time t to determine if they were still present. Subsequent experiments have required the trapped atom to interact repeatedly with fields applied either along the cavity axis or from the side of the cavity. In this case, lifetimes have in practice been limited to hundreds of milliseconds due to heating of the atom by the applied fields.

In order to counterbalance these heating processes, one can imagine some method for cooling the atom in situ after it has been loaded into the trap: cooling intervals could then be interleaved as often as necessary between experimental cycles. Interleaved cooling would allow more cycles to occur before the atom was eventually heated out of the trap; ideally, the interrogation time would be limited only by the intrinsic FORT lifetime. In addition, cooling would localize the atom's center-of-mass motion within a single FORT well. As the atom-cavity coupling g is spatially dependent, an atom moving within a potential well sees a periodically modulated coupling whose amplitude is proportional to temperature. Effective cooling would thus restrict the range of g values that the atom could sample.

One means of cooling a trapped atom is by driving a series of Raman transitions that successively lower the atom's vibrational quantum number n, a process known as sideband cooling. Resolved sideband cooling — in which the vibrational trap frequency ω_a is much greater than the Raman transition linewidth Ω, so that n-changing transitions can be independently addressed — was first demonstrated for trapped ions [33, 34] and subsequently extended to neutral atoms trapped in optical lattices [35, 36, 37]. Over the past several years, David Boozer has developed a powerful set of Raman techniques for use in the lab 11 cavity QED experiment. By introducing a second 935.6 nm laser along the cavity axis, phase-locked and polarized perpendicular to the FORT but detuned from it by the cesium ground state hyperfine splitting $\Delta_{HF}/2\pi = 9.2$ GHz, we can drive Raman transitions between the $F = 3$ and $F = 4$ cesium ground states. Thus, the FORT provides both an optical trap for atoms and one arm of a Raman pair. The FORT-Raman pair can be used to prepare an atom in a desired internal state, to measure the magnetic field at the atom's location, and finally, to cool its center-of-mass motion [19]. Because the spatial structure of both beams is defined by the cavity mode, the two standing waves are automatically superimposed upon one another, and the Rabi frequency (proportional to the product of the two field intensities) is the same at the bottom of each FORT well. (The second beam in this configuration will be referred to hereafter as the "Raman beam.")

The theory of Raman cooling as well as measurements in which cooling is shown to extend the mean lifetime can be found in Dave's thesis [19]. This cooling procedure was also successfully incorporated into the vacuum-Rabi splitting experiment [17], and Fourier-transformed data of $g^{(2)}(\tau)$ from this experiment demonstrated a shift in the distribution of axial vibrational frequencies in the presence of cooling [29]; the shift was due to the anharmonic character of the FORT well, which means that the trap frequency depends upon the atom's kinetic energy. However, the fact that the Raman and FORT beams are spatially superimposed, though convenient for other applications, makes it difficult to achieve effective cooling. If the relative phase between a FORT and a Raman antinode is given by α, then in the harmonic approximation at the bottom of a FORT well, the Rabi frequencies for transitions

between vibrational levels $n, n \pm 1, n \pm 2$ are given by

$$\Omega_{n \to n}/\Omega_\rho = 1/2 + (1/2 - \eta^2(2n+1)) \cos 2\alpha,$$

$$\Omega_{n \to n \pm 1}/\Omega_\rho = -\eta\sqrt{n \pm 1} \sin 2\alpha,$$

$$\Omega_{n \to n \pm 2}/\Omega_\rho = -\eta^2\sqrt{n \pm 1}\sqrt{n \pm 2} \cos 2\alpha,$$

where η is the Lamb-Dicke parameter and Ω_ρ is the (radially dependent) Rabi frequency at the FORT antinode [38]. As $\alpha = 0$ for the FORT-Raman configuration, we cannot drive couplings between vibrational states n and $n - 1$. Cooling must take place on $n \to n - 2$ transitions, which are suppressed by $\eta^2 \approx 0.0025$ in our experiments.

In order to access $n \to n-1$ transitions, we introduced a new Raman configuration in the cavity in early 2006. The FORT still provided a trapping potential at $\lambda_F = 936$ nm (a standing wave with $n_F = 90$ half-wavelengths inside the cavity) but no longer contributed to the Raman process; Raman coupling now took place 10 nm away at $\lambda_R = 945.6$ nm, the next longitudinal mode of the cavity ($n_R = 89$ half-wavelengths), using a pair of phase-stable lasers tuned $\pm\Delta_{HF}/2$ above and below the cavity resonance. As $\lambda_R \neq \lambda_F$ and n_R and n_F have no common denominator, α takes a different (nonzero) value for every well. Transitions from $n \to n - 1$ are possible in each well, though cooling is more effective in certain wells (i.e., one-quarter and three-quarters of the way along the cavity, where $\alpha \approx \pi/4, 3\pi/4$).

2.2 Experimental setup

Phase stability between the FORT and Raman lasers in previous experiments was provided by an electronic phase lock: an optical beat note between the two lasers was used to generate an error signal, which was then fed back to the Raman laser current and the piezoelectric transducer controlling the length of the external laser cavity [19]. However, ~ 500 kHz noise observable on the phase lock beat note at Δ_{HF} presented serious problems, especially as the $n \pm 1$ vibrational sidebands of the FORT

were 530 kHz on either side of the carrier and could not be separated from this noise. In setting up the pair of 946 nm diode lasers for Raman cooling, we have instead chosen a master-slave configuration, with one laser injection-locked to a frequency sideband of the other; frequency sidebands at Δ_{HF} are applied with a fiber-coupled electro-optic modulator (EOM) from EOSPACE Inc. In his thesis, Russ Miller will describe the specifics of this injection lock, which has since been transferred to the FORT-Raman 936 nm pair.

The lab 11 experimental apparatus is described in detail in Jason McKeever's thesis [39]; here I summarize the essential components and enumerate recent developments. At the heart of the experiment is the high-finesse optical cavity, constructed by Jun Ye and David Vernooy and enclosed in a two-chamber vacuum system [40]. As we are now two generations of graduate students removed from the original assembly of this system, we treat the cavity and chamber with a great deal of caution and a bit of superstition. A valve controls the flow of cesium into the upper chamber and is opened for a few hours before we attempt to load atoms into the cavity. Surrounding the vacuum chamber are ion pumps, cameras, assorted optics for coupling light in and out of the chamber, and coils to generate magnetic gradient and bias fields for the atoms.

2.2.1 Lasers

We use a collection of home-built diode lasers to cool, trap, and probe cesium atoms and to transfer population between their internal states. The central laser in this scheme is the "master laser" on the southwest corner of the optical table. For the purpose of frequency stabilization, this laser is locked to a ~ 30 cm transfer cavity that is subsequently locked to cesium, and specifically, to the crossover between the $F = 4 \rightarrow F' = 4$ and $F = 4 \rightarrow F' = 5$ $D2$ resonances in a vapor cell via modulation transfer spectroscopy. (For reference, a diagram of the cesium $D2$ spectrum can be found in Figure 4.2.) A pair of acousto-optic modulators (AOMs) is used to shift the master laser into resonance with either the $F = 4 \rightarrow F' = 4$ or $F = 4 \rightarrow F' = 5$

transition and to shutter the beam on and off. After a series of optical filters, light from the laser is then coupled into the physics cavity; we often refer to this laser as the "probe laser" because it is used to interrogate the atom-cavity system. It has never been necessary to mode-match the probe efficiently into the cavity, and the most recent measurement (in March 2006) found the probe input coupling to be 12.5%.

Light from the master laser is also sent through optical fiber to injection-lock a "slave laser," an 852 nm diode laser without an external grating. This slave laser provides the red-detuned $F = 4 \rightarrow F' = 5$ trapping light for MOTs in the upper and lower chamber, which can be shifted to a different frequency and intensity setting for polarization-gradient cooling [29]. Along a separate path, the slave laser also provides blue-detuned $F = 4 \rightarrow F' = 4$ light that is then circularly polarized and enters the cavity from the side to generate orthogonal standing waves. In tandem with the $F = 3$ light described below, these "lattice beams" are used to cool falling atoms into the FORT and to prepare the internal states of trapped atoms.

On the north side of the table, a "repump laser" is locked to the $F = 3$ cesium resonance. Its primary function is to generate beams for both MOTs on the $F = 3 \rightarrow F' = 4$ transition that can recycle atoms to the $F = 4$ manifold. In addition, it is the source of an $F = 3 \rightarrow F' = 3$ lattice beam collinear with the one described above (Ω_3 and Ω_4 in Figure 2.3(b)). The current-modulated repump laser was locked to the saturated absorption spectrum of a cesium vapor cell through 2006. At that time, we replaced the cesium lock with a phase lock to the master laser, for purposes described in Chapter 3.

For the experiment described in this chapter, we introduced a frequency-shifted beam from the probe laser path that was resonant with the $F = 4 \rightarrow F' = 4$ transition. This beam was coupled into optical fiber that carried it to the side of the cavity (Ω'_4, Figure 2.3(b)). Focused into the cavity with a cylindrical lens, it can now be used in combination with the Ω_4 lattice beam to pump a trapped atom into the $F = 3$ manifold; in the absence of a magnetic field, alternating the two beams is necessary to avoid a dark state in the $F = 4$ manifold [4, 19].

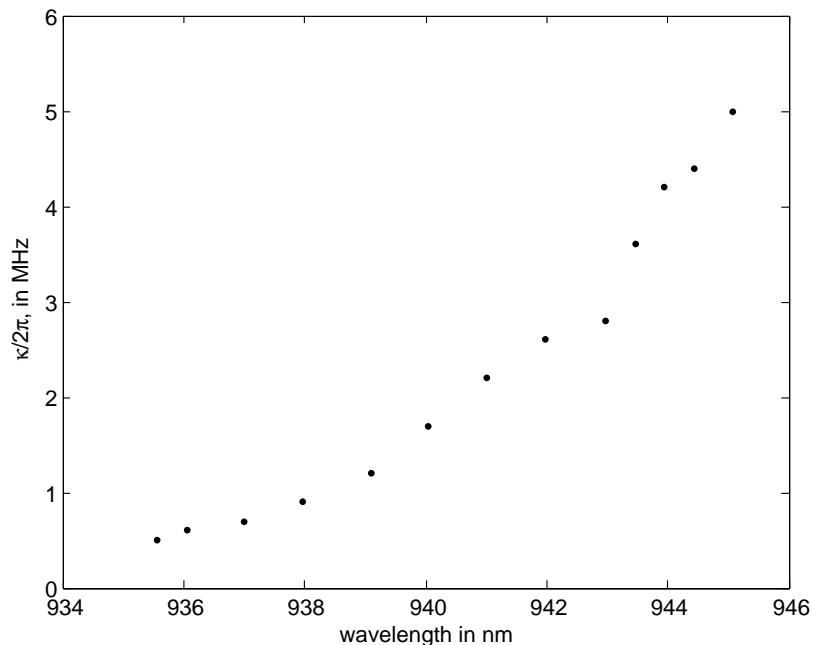At 836 nm, our one store-bought diode laser (thus referred to as the "New Focus

Figure 2.1: Measured physics cavity linewidths between the two Raman wavelengths of 935.6 nm and 945.6 nm

laser") is resonant with the cavity mode two free spectral ranges away from the cesium resonance. It is used to provide active feedback to maintain the cavity length via a Pound-Drever-Hall lock in transmission [41]. The New Focus laser frequency is stabilized by a reflection lock to the transfer cavity; the physics cavity is then locked to a frequency sideband of the New Focus laser. The sideband lock allows us to tune the sideband frequency manually, so that we can adjust the physics cavity frequency ω_c with respect to the cesium resonance ω_a. The sidebands, applied with a highly tunable traveling-wave modulator, are usually between 300 and 800 MHz in frequency, and the transmission lock is derived from a 4 MHz modulation of this signal.

The probe laser and New Focus laser are combined at the cavity input with the FORT laser at 936 nm. While Jason's thesis mentions that the FORT was frequency-modulated and locked to the physics cavity, this lock has been since discontinued and the EOM removed, as it was found to be unnecessary: since the FORT cavity linewidth is so wide ($\kappa_F/2\pi = 0.8$ GHz), once the laser is tuned into resonance, its

slow transmission drifts can be corrected manually on the timescale of hours. In fact, changes in intracavity power due to amplitude drift at the cavity input can be more substantial than those due to frequency drift and would not be corrected by a frequency lock. More recently, we have implemented a feedback loop to directly stabilize the FORT power in transmission using a liquid-crystal variable oscillator (LCVO).

At the side of the lab, a second $F = 3$ diode laser with its own cesium lock sits on a honeycomb breadboard. In 2005, we rebuilt this laser so that it provided separate beams on the $F = 3 \rightarrow F' = 3$ and $F = 3 \rightarrow F' = 4$ transitions, each with switching capabilities. The beams were recombined and coupled onto orthogonal axes of a polarization-maintaining fiber. As with the $F = 4 \rightarrow F' = 4$ beam above, they were brought in from the side of the cavity and used in several optical pumping attempts (Chapter 4). Despite our best efforts to stabilize the diode and isolate the assembly from air currents, this laser has always been difficult to work with, probably because it does not benefit from the vibration isolation of an optical table.

The final lasers in our experiment, the 946 nm Raman pair, have already been discussed above. It was not originally known whether the cavity would support a mode at the new Raman wavelength, as the coating curve for the cavity mirrors falls off steeply in this vicinity [42]. In order to measure the cavity linewidth as a function of wavelength, we increased the wavelength of the old Raman laser in 1 nm steps from its original value of λ_F. We used the sidebands of the New Focus locking laser, set to 1 GHz, as a meterstick in order to obtain the cavity linewidth κ from the scope trace at each wavelength. The linewidth data are plotted in Figure 2.1; at our wavelength of interest, $\kappa_R/\kappa_F = 6$ GHz/0.8 GHz, and the finesse at λ_R decreases proportionally from $\mathcal{F}_F = 2200$ to $\mathcal{F}_R = 300$.[1]

---

[1] The values given here for κ_R and κ_F are drawn from Section 2.3 rather than the data in Figure 2.1. In particular, a subsequent measurement determined that $\kappa_F = 0.8$ rather than 0.5 as shown above.

2.2.2   Output path and detection

At the cavity output, a half-waveplate at 852 nm and a polarizing beamsplitter cube (PBS) are used to select output light along one of the cavity's two birefringent axes.[2] At the rejected port of the PBS, a New Focus DC-125 MHz detector monitors the FORT transmission. Unfortunately, this scheme cuts our detection efficiency in half for unpolarized photon generation and prevents us from examining orthogonal output polarizations simultaneously. Because we have some well-founded hesitation about dismantling the output path of a working experiment, we have left it in place until now but have plans for a rebuild in the near future.

After the PBS, the cavity output is collimated, and an angle-tuned mirror separates the 836 nm light for the cavity lock. A series of four dichroic mirrors and two interference filters suppresses the remaining 836 nm and 936 nm light in the path before the light is finally coupled into optical fiber. In December 2005, we measured the propagation efficiency p_{table} to be 66% before fiber coupling, with losses attributed to individual elements listed in Table 2.1. This measurement is roughly consistent with the values of $p_{path} = p_{couple}p_{table} = 32\%$ and 40% given in Refs. [16, 29], where p_{couple} includes fiber coupling losses; in November 2006, we improved the fiber alignment, which boosted p_{path} from 30% to 50%. The dichroic elements are relatively ineffective at suppressing FORT light — each dichroic removes less than half of incident light at 936 nm — and should probably be removed from future beam paths, as they are also not perfect at 852 nm. Each interference filter, on the other hand, has a measured suppression of 10^4 at 936 nm. It is interesting to note that given the measured FORT propagation losses at each step of the path, we would still expect the FORT power to saturate our detectors, while in fact we see fewer than 5 photon counts/second due to 936 nm light. It must be true that optimizing the probe coupling into fiber at 852 nm results in very poor FORT fiber coupling. We believe that this is due to refraction in the PBS cube at the cavity output, which causes light at the two wavelengths to

---

[2]The cavity supports two orthogonal linear modes, with a frequency splitting between the modes roughly equal to κ. The most recent measurements of this frequency splitting can be found in Ref. [29]. Chapter 6 provides a discussion of birefringence in the context of cavity design.

Path element(s)	Propagation losses, in percent
first mirror	0
angle-tuned splitter	3
second mirror	2
third mirror + lens	1
iris + half-waveplate + lens + dichroic 1	1
dichroics 2 + dichroic 3	10
first interference filter	10
second interference filter	17
dichroic 4	3

Table 2.1: Measured propagation losses at 852 nm in the lab 11 cavity output path. Some path elements are grouped together because of the practical difficulty of inserting a detector between them.

follow slightly different paths.

Once the cavity light is coupled into single-mode fiber, a fiber beamsplitter carries it to two single-photon-counting modules (Perkin-Elmer SPCM−AQR). We record TTL logic pulses from the SPCMs using a P7888 four-channel data acquisition card; these cards are also currently in use in lab 2 and in the microtoroid experiment. A heterodyne detection path, while still in place on the optical table, has been used in recent years only for alignment purposes, in cases where we were concerned about the possibility of too much incident light on the SPCMs.

2.2.3   Analog and digital control

Analog and digital timing control signals for the experiment are generated by an ADwin-Gold system (Jäger Computergesteuerte Messtechnik GmbH). The ADwin has an internal 40 MHz CPU; we map a series of instructions into its RAM from a PC via USB connection, and it then runs continuously until the next program is loaded. Dave has written a straightforward user interface that compiles our timing programs into the (less intuitive) ADbasic language that the system requires. In Chapter 4, I describe recent modifications to our timing programs that allow us to condition the experiment on the presence of at most one atom in the cavity.

Most of the experiment timing is digital, but the analog outputs of the ADwin are
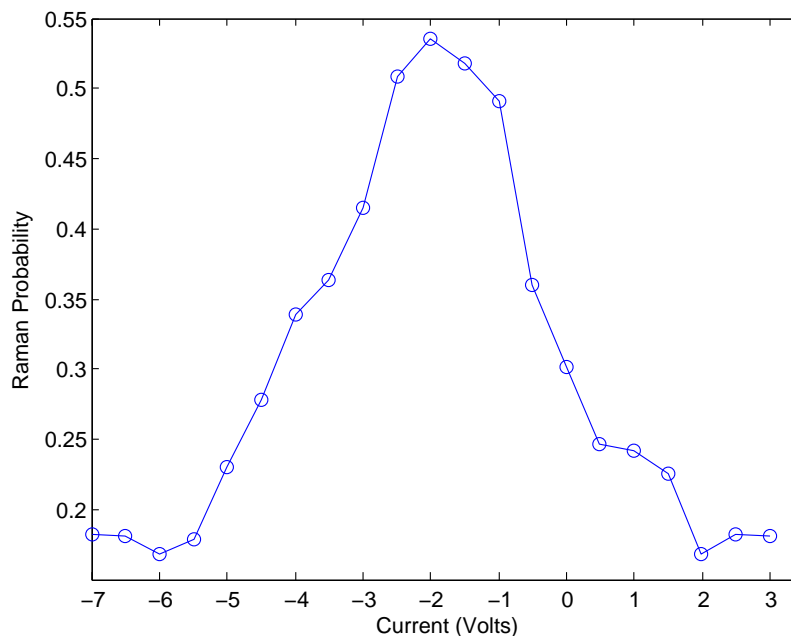
Figure 2.2: Probability to drive Raman transitions as a function of control voltage for the magnetic bias field along the "fed from above" axis, where the conversion factor is 0.97 Gauss/V. The peak indicates a local field minimum.

used to provide control voltages for the bias coil current supplies, which determine the magnetic fields at the cavity. In order to have control over the magnetic field that an atom experiences during an experiment, we first need to adjust the bias coil currents in order to null the ambient fields seen by the atom. Here we rely on the fact that when the fields are nulled, the Raman transition frequencies between all Zeeman sublevels are degenerate. Thus, for a given atom in an unknown Zeeman state, the probability to drive a Raman transition is greatest in a nulled field [19]. If we record the Raman transition probability as a function of bias coil current (that is, as a function of ADwin control voltage) along each axis, we can then associate the maximum transition probability with a local field minimum.

An example of a Raman field nulling scan along one axis is shown in Figure 2.2. During the measurement, fields along the other two axes are fixed at their nominal null values. As these initial values may not be accurate, it is sometimes necessary to iterate the nulling procedure a few times in order to find the field minimum along all

three axes. This was at first a rather time-consuming process, as we would increment the bias coil current by hand. Once we switched to an ADwin control program that automatically scanned the current and stamped the data records accordingly, the situation improved greatly, and we are now able to null the magnetic fields within approximately half an hour. We find that for experiments requiring sensitive control of magnetic fields, it is best to null the fields within a few days of data collection to avoid any surprising drifts. The axial bias field value seems to be the most stable, while the current through the "fed from below" coil is the most likely to need adjustment.

2.3   Demonstration of ground-state cooling

The following section has been adapted from Ref. [30].

Localization to the ground state of axial motion is demonstrated for a single, trapped atom strongly coupled to the field of a high-finesse optical resonator. The axial atomic motion is cooled by way of coherent Raman transitions on the red vibrational sideband. An efficient state detection scheme enabled by strong coupling in cavity QED is used to record the Raman spectrum, from which the state of atomic motion is inferred. We find that the lowest vibrational level of the axial potential with zero-point energy $\hbar\omega_{\mathrm{a}}/2k_B = 13 \ \mu K$ is occupied with probability $P_0 \simeq 0.95$.

Single atoms strongly coupled to the fields of high-quality optical resonators are of fundamental importance in Quantum Optics and, more generally, can be used for many tasks in quantum information science, including the implementation of scalable quantum computation [43, 44] and the realization of distributed quantum networks [2, 45]. In recent years, significant experimental progress to develop tools suitable for these tasks has been made by employing optical forces to localize individual atoms within optical cavities in a regime of strong coupling [12, 13, 17, 46, 47, 48, 49], as well as by combining trapped ions with optical cavities [50, 51]. Scientific advances thereby enabled include the observation of the vacuum-Rabi spectrum for an individual atom [17] and vacuum-stimulated cooling [48].

Although great strides are being made with atoms localized and strongly coupled

to the fields of optical cavities, it has not previously been possible to access the quantum regime for the atomic center-of-mass motion in cavity QED. Qualitatively new phenomena have been predicted in this regime for which a quantized treatment is required for both the internal (i.e., the atomic dipole and cavity field) and external (i.e., atomic motion) degrees of freedom, as was first recognized in the seminal work of Refs. [52, 53, 54] and in the years since [55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66]. Examples include the transfer of quantized states of atomic motion to quantum states of light, and conversely [64, 65, 66], as well as for measurements that surpass the standard quantum limit for sensing atomic position [55, 56].

Our effort towards quantum control of atomic motion in cavity QED follows the remarkable set of achievements for trapped ions [34] and atoms in optical lattices [35, 36, 37], for which such control has led to the creation of manifestly quantum states of motion and to the manipulation of quantum information. A first step in many of these investigations has been the capability to cool to the ground state of motion for single trapped atoms or ions.

Here we report localization to the ground state of motion for one atom trapped in an optical cavity in a regime of strong coupling [49]. Resolved sideband cooling to the ground state is accomplished with a coherent pair of intracavity Raman fields. To deduce the resulting state of atomic motion, we introduce a scheme for recording Raman spectra by way of the interaction of the atom with a resonant cavity probe. Our scheme is the cavity QED equivalent of state detection in free space by quantum-jump spectroscopy [34] and achieves a confidence level for state discrimination $> 98\%$ in 100 μs. From the Raman spectra, we infer that the lowest vibrational level $n = 0$ of the axial potential is occupied with probability $P_0 \simeq 0.95$ for one trapped atom.

A schematic of the experiment is given in Figure 2.3. At the heart of the system is the Fabry-Perot cavity formed by mirrors (M_1, M_2). The cavity length is stabilized to $l_0 = 42.2$ μm using an independent locking laser, such that a TEM_{00} mode is resonant with the $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F' = 5$ transition of the $D2$ line in Cs. The resulting atom-cavity coupling gives a maximum single-photon Rabi frequency of $2g_0/2\pi = 68$ MHz for $(F = 4, m_F = \pm 4) \rightarrow (F' = 5, m'_F = \pm 5)$. The decay
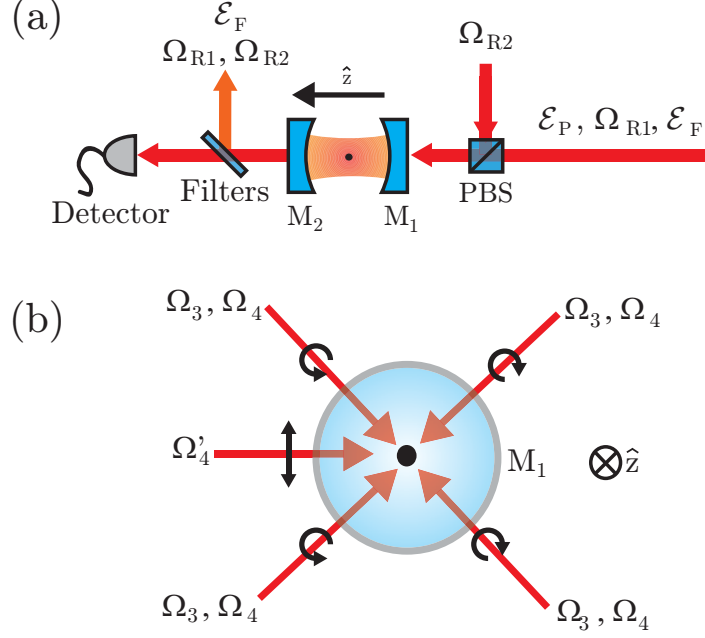
Figure 2.3: Schematic of the experiment. The cavity is represented (a) from the side and (b) along its axis, with a trapped atom indicated as a dot at the cavity center. Shown are the various beams used in the experiment: linearly polarized probe $\mathcal{E}_p$, FORT $\mathcal{E}_F$, pumping Ω_4', and Raman beams $\Omega_{R1,R2}$, as well as the circularly polarized lattice beams $\Omega_{3,4}$.

rates are $\gamma/2\pi = 2.6$ MHz for the $6P_{3/2}$ excited states, and $\kappa/2\pi = 4.1$ MHz for the cavity field.[3] Because $g_0 \gg (\gamma, \kappa)$, our system is in the strong coupling regime of cavity QED [49], with critical photon and atom numbers $n_0 \equiv \gamma^2/(2g_0^2) \approx 0.0029$ and $N_0 \equiv 2\kappa\gamma/g_0^2 \approx 0.018$.

Atoms are trapped by an intracavity far-off-resonant trap (FORT) at $\lambda_F = 935.6$ nm, which is driven by a linearly polarized input field $\mathcal{E}_F$ and is resonant with a TEM_{00} mode of the cavity with linewidth $\kappa_F/2\pi = 0.8$ GHz. At λ_F, states in the ground $F = 3, 4$ and excited state $F' = 5$ manifolds experience nearly equal trapping potentials. For states in the $F = 3, 4$ manifolds, this potential is independent of m_F and has a peak value of $U_F/h = -41$ MHz, while for states in the $F' = 5$

---

[3]This value for $\kappa/2\pi$ differs from the 3.8 MHz rate found elsewhere in this thesis, which is based on a more recent measurement. As discussed in detail in Ref. [29], we have seen as much as 20% variation in measured cavity linewidths using different techniques. Some of this variation may be due to input light which is not polarized along one of the cavity's birefringent axes and thus couples to both cavity modes.

manifold it has a weak dependence on m'_F [13, 17]. The standing-wave structure of the FORT forms independent wells where atoms may be trapped. Near the bottom of a FORT well the potential is approximately harmonic, with axial and radial vibrational frequencies $\omega_a/2\pi = 530$ kHz, $\omega_r/2\pi = 4.5$ kHz.

To load atoms into the FORT, we release a cloud of cold atoms located ~ 3 mm above the cavity [13]. As the atoms fall through the cavity, we apply 5 ms of Ω_3, Ω_4 light by way of two pairs of counterpropagating $\sigma_+ - \sigma_-$ polarized beams; see Figure 4.1(b). These beams are blue detuned by $+10$ MHz from the $F = 3 \rightarrow F' = 3$ and $F = 4 \rightarrow F' = 4$ transitions respectively, and cool the falling atoms via polarization gradient cooling [67]. We adjust the powers[4] of these beams so that the probability of loading at least one atom is ~ 0.3 per cloud drop.

Raman coupling between the $F = 3$ and $F = 4$ manifolds is generated by driving a cavity mode at $\lambda_R = 945.6$ nm with a pair of beams $\Omega_{R1,R2}$ that are phase-locked, lin $\perp$ lin polarized, and have a relative detuning $\omega_{R1} - \omega_{R2} = \Delta'_{HF} + \delta_R$, where $\Delta'_{HF}/2\pi = 9.19261$ GHz is the Cs hyperfine splitting[5] and δ_R is the Raman detuning. This cavity mode has a linewidth $\kappa_R/2\pi = 6$ GHz , and the Raman beams are tuned such that $\Omega_{R1(R2)}$ lies $\Delta'_{HF}/2$ above (below) cavity resonance. Since $\Omega_{R1,R2}$ drive a different mode of the cavity than $\mathcal{E}_F$, atoms trapped in different FORT wells see different Raman powers. If the relative spatial phase along the cavity axis between the FORT and the Raman pair at the bottom of a given well is α, then an atom at this potential minimum sees a Raman power proportional to $\cos^2 \alpha$.

We set the optical power transmitted on resonance through the cavity for $\Omega_{R1,R2}$ to $P_{R1} = P_{R2} = 140$ μW, which gives a Rabi frequency[6] $\Omega_0/2\pi = 200$ kHz for atoms with $\alpha = 0$. The ac-Stark shift due to these beams adds a correction to the FORT potential of $U_R/2\pi = 0.84$ MHz. To avoid heating the atom by switching U_R, we

---

[4]For loading, the total power in the four lattice beams is $50I_4^{\text{sat}}$ for Ω_3 and Ω_4, where $I_4^{\text{sat}} \sim 3.8$ mW/cm^2. For detection, the Ω_3 intensity is $5I_4^{\text{sat}}$.

[5]Δ'_{HF} for a trapped atom is slightly reduced compared to its free space value Δ_{HF}, because the FORT potential for the $F = 3$ manifold is slightly weaker than for $F = 4$. For our FORT, the correction is $\Delta'_{HF} - \Delta_{HF} = -21$ kHz.

[6]Here Ω_0 is the Rabi frequency for the transition $F = 3, m = 0, n = 0 \rightarrow F = 4, m = 0, n = 0$, where the quantization axis is along the cavity axis.
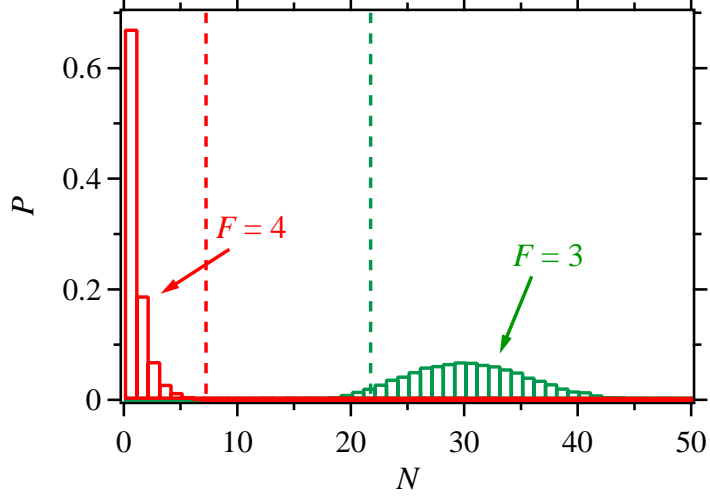
Figure 2.4: Histogram of counts recorded during 100 μs probing intervals for $N_e = 30$. The left (right) curve is the probability P of detecting a number of counts N for an atom prepared in the $F = 4$ ($F = 3$) state. The dashed lines indicate detection thresholds at $N = 0.25N_e$ and $N = 0.75N_e$.

leave $\Omega_{R1,R2}$ on all the time, but usually keep them far-detuned ($\delta_R/2\pi = 85$ MHz) from Raman resonance. To drive Raman transitions, we change ω_{R2} to bring the pair into Raman resonance, whereas to fine-tune δ_R we vary ω_{R1}.

Because the intensity of the Raman pair is spatially varying, the Raman coupling can connect states with different vibrational quantum numbers. The form of this motional coupling is simple for atoms near the bottom of a FORT well where the axial and radial motions decouple, allowing us to consider the effect of the Raman coupling on the axial motion alone. In this harmonic limit, we can define a set of Fock states $\{|n\rangle\}$ for the axial motion. For transitions from $F = 3, m = 0$ to $F = 4, m = 0$ and to first order in η, the Rabi frequency for an $n \to n$ transition is $\Omega_{n \to n} = (1/2)(1 + \cos 2\alpha)\,\Omega_0$, while for an $n \to n - 1$ transition, $\Omega_{n \to n-1} = \eta\sqrt{n}\sin 2\alpha\,\Omega_0$, where $\eta = (2\pi/\lambda_R)\sqrt{\hbar/2m\omega_a} = 0.056$ is the Lamb-Dicke parameter for axial motion. Note that the $n \to n - 1$ transition is strongest for atoms with $\alpha = \pi/4$.

The spatial dependence of the Raman coupling, together with the fact that the the axial motion of the atom is in the Lamb-Dicke limit, allows us to implement Raman sideband cooling [34]. We tune the Raman pair to the red axial sideband

($\delta_R = -525$ kHz $\simeq -\omega_a$) and apply the Ω_4 lattice beams. An atom that starts in $F = 3$ is coherently transferred by $\Omega_{R1,R2}$ to $F = 4$, where it is incoherently repumped to $F = 3$ by Ω_4. The coherent transfer lowers the axial vibrational quantum number n by one, while the incoherent repumping usually leaves n unchanged since n-changing transitions are Lamb-Dicke suppressed. Thus, the beams continually lower n, cooling the atom to the axial ground state. Also, the Ω_4 light provides Sisyphus cooling [67] in the radial direction.

Strong atom-cavity coupling enables a versatile detection scheme for determining if an atom is present in the cavity, and if so, if it is in the $F = 3$ or $F = 4$ manifold. In 100 μs we measure the atomic hyperfine state with a confidence level of $> 98\%$. The scheme involves driving the cavity with a 100 μs pulse of resonant $F = 4 \rightarrow F' = 5$ probe light $\mathcal{E}_P$. If an $F = 3$ atom is present (or if the cavity is empty), then the light is transmitted, while if an $F = 4$ atom is present, then the light is blocked because of strong coupling [17]. We set the $\mathcal{E}_P$ intensity such that on average N_e photons are detected[7] per probing interval with no atom in the cavity. As shown in Figure 2.4, if the number N of detected photons is such that $N < 0.25N_e$, we assume an $F = 4$ atom is present; if $N > 0.75N_e$, we assume an $F = 4$ atom is not present; otherwise, the measurement is inconclusive ($< 2\%$ of the time) and we ignore the result. Whenever we detect the atomic state, we perform two such measurements: the first with $\mathcal{E}_P$ to find out if an $F = 4$ atom is present, the second with $\mathcal{E}_P$ together with Ω_3 as a repumper, to detect an atom, regardless of its internal state.

We measure the Raman transfer probability P_4 for a given δ_R by preparing an atom in $F = 3$, applying a Raman pulse, and then detecting the atomic state using the above scheme (with $N_e \sim 22$). For each measurement cycle (or trial), we first Raman-sideband cool the atom for an interval Δt_c. Next, we pump it into $F = 3$ by alternating 1 μs pulses of Ω_4 lattice light with 1 μs pulses of Ω'_4 linearly polarized resonant $F = 4 \rightarrow F' = 4$ light from the side of the cavity (10 pulses of each). After the atom is pumped to $F = 3$, we apply a $\Delta t_R = 500$ μs Raman pulse, which

---

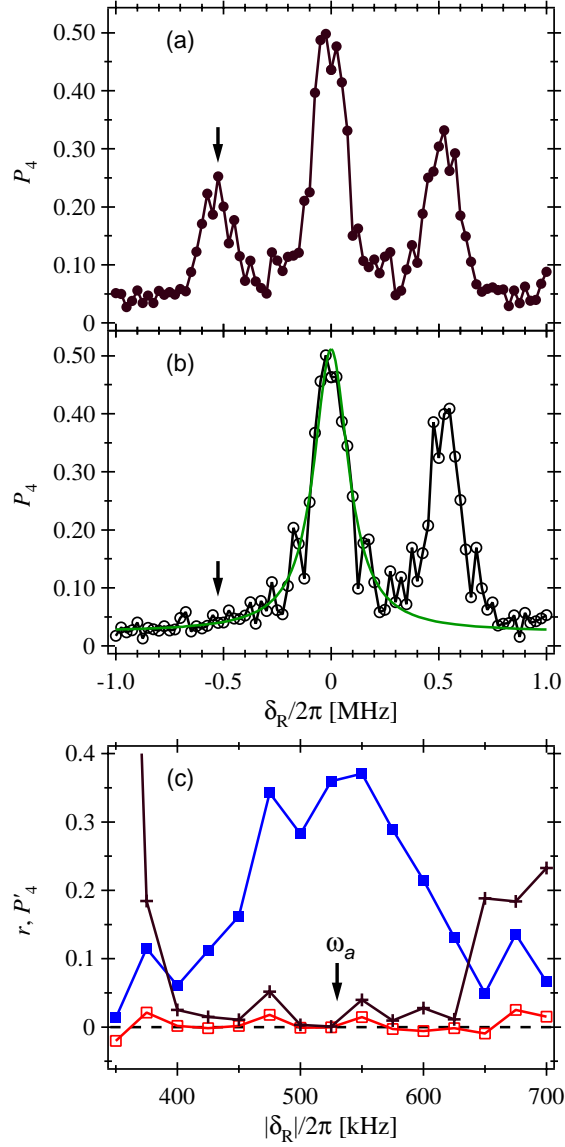[7]The probability for a photon starting out within the cavity to be recorded by our detectors is $e \sim 0.06$.

Figure 2.5: Population P_4 in the $F = 4$ state vs. Raman detuning $\delta_R/2\pi$. The data in (a) are taken with $\Delta t_c = 250$ μs of cooling, and an Ω_4 total 4-beam intensity $I_4 = 5I_4^{\text{sat}}$; those in (b) with $\Delta t_c = 5$ ms of cooling, and $I_4 = 0.5I_4^{\text{sat}}$ (on average, ~ 33 atoms per data point). The arrow marks the detuning used for sideband cooling. (c) Zoom-in on the two sideband regions for the data in (b), with detuning axis folded around $\delta_R = 0$. The red ($\square$) and blue ($\blacksquare$) sidebands, and their ratio r (+), are shown after subtracting a Lorentzian fit to the carrier [superimposed in (b)].

sometimes transfers it to $F = 4$. Finally, we measure the atomic state and check if the atom is still present. For each atom we fix the absolute value of the Raman

detuning $|\delta_R|$, and alternate trials at $+|\delta_R|$ with trials at $-|\delta_R|$ (299 trials each). By combining data from atoms with different values of $|\delta_R|$, we map out a Raman spectrum. Note that because the initial Zeeman state of the atom is random, all allowed $F = 3 \rightarrow F = 4$ Zeeman transitions contribute to these spectra.

Two example Raman spectra are plotted in Figure 2.5. For the curve in Figure 2.5(a), we cool for $\Delta t_c = 250$ μs; for the curve in Figure 2.5(b), for $\Delta t_c = 5$ ms. These scans are performed after nulling the magnetic field to within ~ 40 mG; the widths of the peaks are set by the splitting of different Zeeman levels due to the residual magnetic field. For the curve in panel (a), we see peaks at the carrier ($\delta_R = 0$), as well as at the blue/red sidebands ($\delta_R/2\pi \simeq \pm 530$ kHz $= \pm \omega_a$). Already we note a sideband asymmetry, indicating that a significant fraction of the population is in the $n = 0$ vibrational state. For the (b) data, the red sideband at $\delta_R/2\pi \simeq -530$ kHz is suppressed such that it cannot be distinguished from the background and contribution from off-resonant excitation of the carrier.

The ratio r of transfer probabilities for the red and blue sideband gives information about the temperature of the atom. For a two-state atom in a thermal state, this ratio r_0 at $|\delta_R| = \omega_a$ is related to the mean vibrational quantum number $\bar{n}$ by $r_0 = \bar{n}/(\bar{n}+1)$ [34]. In Figure 2.5c, we plot r as a function of $|\delta_R|$ for the $\Delta t_c = 5$ ms data. As shown in Figure 2.5b, we fit a Lorentzian curve to the carrier, then subtract its contribution from both the red and the blue sideband data, as shown in Figure 2.5(c). We find $r_0 \simeq \bar{n} = 0.01 \pm 0.05$, and the ground state population $P_0 = 1/(\bar{n} + 1) = 0.99 \pm 0.05$, where the error bars reflect fluctuations in the data around $|\delta_R| = \omega_a$. If instead we subtract the constant background of $P_4^B = 0.024$ but not the carrier's Lorentzian tail, we find $r_0 \simeq \bar{n} = 0.05 \pm 0.04$, and $P_0 = 0.95 \pm 0.04$. Finally, if we use the raw data from Figure 2.5b with no subtractions, we obtain $r_0 = 0.10 \pm 0.03$, $\bar{n} = 0.12 \pm 0.04$ and $P_0 = 0.89 \pm 0.03$. Because the atom is not a two-state system and the motional state is not known to be thermal, these estimates are approximate.

The axial cooling rate and asymptotic value of $\bar{n}$ depend on δ_R, on the $\Omega_{R1,R2}$ Rabi frequencies, and on the power and detuning of Ω_4. We have performed computer simulations to help us choose optimal values for these parameters. A common feature
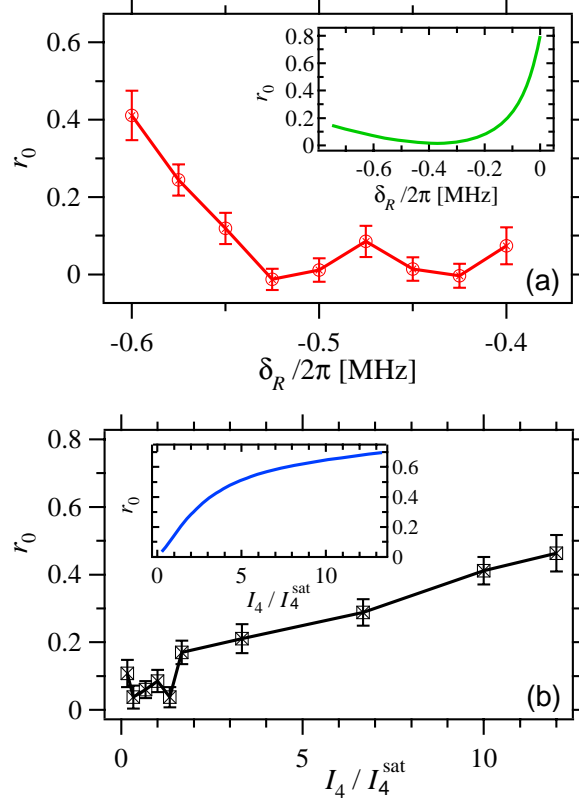
Figure 2.6: Varying cooling parameters. The sideband ratio r_0 is shown as a function of (a) the Raman detuning δ_R employed for cooling and (b) the $4 \to 3$ repumping intensity I_4 . Insets show the results from a simple calculation for a 2-state atom trapped in a FORT well with $\alpha = \pi/4$.

of both our theoretical and experimental investigations is the robustness of $\bar{n}$ under variations of the cooling parameters. As an example[8] , in Figure 2.6 we plot the measured sideband ratio r_0 at $\delta_R/2\pi = -500$ kHz $\simeq -\omega_a$ as a function of (a) the detuning δ_R used for sideband cooling, and (b) the recycling intensity I_4. The sideband asymmetry is maintained over a range of at least 200 kHz in detuning, and of two orders of magnitude in the intensity I_4 of the Ω_4 beams. The insets give results from a 2-state calculation of r_0, displaying similar insensitivity to the exact values of δ_R and I_4.

We use two methods for estimating the mean energy E_r for radial motion. The

---

[8]Unless otherwise noted, the settings for this measurement are: $I_4 = 0.3I_4^{\text{sat}}$, $\Delta t_c = 5$ ms, $\delta_R = -500$ kHz, $\Delta t_R = 50$ μs; $P_4^B = 0.065$ was subtracted before computing r_0.

first method involves adiabatically lowering the FORT depth to zero, so that only the U_R trapping potential remains, and measuring the probability that the atom survives the process [68]. If we assume a thermal state for the radial motion, this method limits $E_r \lesssim 200 \; \mu$K. The second method relies on the fact that radial motion would shift and/or broaden the axial sidebands in our Raman spectra. By applying a small axial bias field, we can resolve motional sidebands of transitions between specific Zeeman levels; the positions and widths of these sidebands also limit $E_r \lesssim 200\mu$K. However, the Sisyphus cooling we use radially has been previously shown to reach temperatures of $\sim 1 \; \mu$K [67], which corresponds to $n_{rad} \simeq 4$ for an atom in our FORT. A possibility for improved determination of the radial "temperature" is to monitor the transmission of the cavity probe field, as in Ref. [17]. Note that the ratio of radial to axial trapping frequencies ($\simeq 0.01$) is such that any modulation of the axial frequency due to radial motion would be adiabatic. As a result, we expect that an atom cooled to the axial ground state does not change its state due to radial motion.

In conclusion, we have demonstrated cooling to the ground state of axial motion for single cesium atoms strongly coupled to the field of a small optical resonator. Together with existing capabilities for strong coupling of the internal degrees of freedom, control over the external center-of-mass motion in cavity QED enables a new set of phenomena to be explored at the light-matter interface. For example, arbitrary states of atomic motion can be prepared from the ground state by coherent Raman transitions [34], then mapped to the electromagnetic field by way of the strong atom-field coupling [64, 65, 66]. Investigations of sensing atomic motion at the standard quantum limit and feedback control now become feasible [55, 56, 69].

2.4   Resolved second-order sideband cooling

As discussed in Section 2.1, it is possible to cool trapped atoms on the second-order red sideband using the FORT-Raman pair. Although this method is less efficient, there are many experiments in which we rely on the FORT-Raman pair to drive Raman transitions whose Rabi frequency is well-independent, such as the Rabi flopping
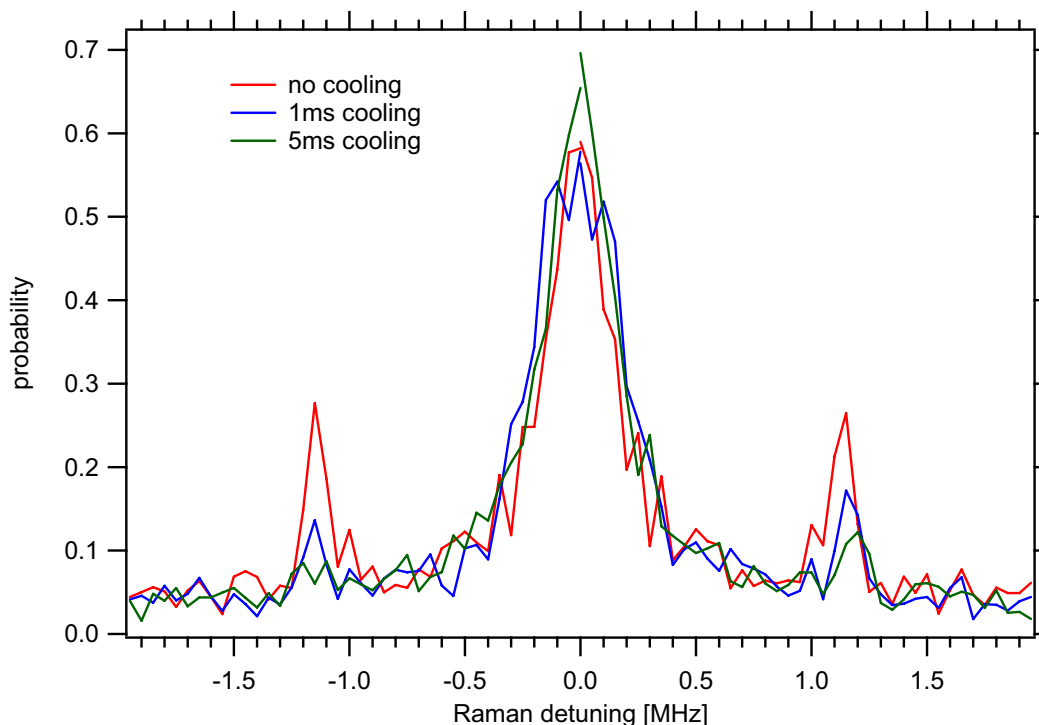
Figure 2.7: Population P_4 in the $F = 4$ state vs. Raman detuning $\delta_R/2\pi$, after second-order sideband cooling. Cooling intervals of (a) 0 ms, (b) 1 ms, and (c) 5 ms are applied using the FORT-Raman pair.

measurements in Chapter 5. It is useful in these situations to be able to cool the atoms, even if we cannot achieve the ground-state cooling described in Section 2.3. (Ideally, our lab would have a pair of Raman lasers at 946 nm for ground-state cooling and a separate FORT-Raman pair for well-independent transitions. At the moment, however, we have only three lasers: dedicated FORT and Raman lasers, and a third injection-locked laser whose wavelength must be adjusted manually in order to move between the two configurations.)

In the past, we were unable to observe resolved-sideband cooling using the FORT-Raman pair because of ~ 500 kHz noise inherent to the phase lock between the lasers. With the replacement of the phase lock by an injection lock during the ground state cooling experiment, this noise was no longer an issue. In the spring of 2007, we measured Raman spectra similar to those presented in Figure 2.5, but in this case with cooling at 936 nm; these data are shown in Figure 2.7. After initial preparation

in $F = 3$, atoms are cooled with an applied pulse of the FORT-Raman pair, tuned 1.150 MHz red of resonance to the $n \to n - 2$ vibrational sideband, in conjunction with a $F = 4 \to F' = 4$ repump. In order to measure the Raman transfer probability for a given detuning, we use the same technique as in the previous section; that is, for each atom at detuning δ, we apply a 500 μs Raman pulse, followed by state detection to determine whether that pulse has transferred an atom to $F = 4$. For these measurements, the bias fields have been nulled, so all $|3, m\rangle \to |4, m\rangle$ transitions are degenerate.

We emphasize two features of the data in Figure 2.7. First, while in the absence of cooling, the red and blue sidebands are roughly symmetric, there is a clear relative suppression of the red sideband after 5 ms of cooling (although the red sideband can still be discerned from the carrier). Second, the widths of both sidebands become narrower, and the center frequencies of the sidebands shift toward higher detunings. This also provides strong evidence for cooling: recall that the FORT potential is sinusoidal and only harmonic in the low-temperature limit. The vibrational frequency spacing at the top of each well is about half the spacing at the bottom, and so the sideband narrowing reflects the fact that we are driving $n \to n \pm 2$ transitions for increasingly smaller n.

Second-order sideband cooling is a promising technique that we hope to incorporate in future experiments, both to localize atoms in the FORT wells and to extend their lifetimes in the trap. However, one caveat is that bias fields must be nulled in order for the cooling to address all Zeeman sublevels, while many of our recent experiments use an axial bias field to split out these sublevels. In principle, one could use an array of RF frequency generators to address the second-order sidebands of all seven $|3, m\rangle \to |4, m\rangle$ transitions independently, but that would constitute a new level of complexity for our experiment.