

Chapter 4

Atom preparation: optical pumping and conditional loading

This chapter presents two significant advances in our ability to prepare trapped atoms within an optical cavity. First, we demonstrate that we can use Raman transitions to pump a trapped atom into any desired Zeeman state. Second, we introduce a feedback scheme for conditioning our experiment on the presence of at most one atom in the cavity. When combined with a new Raman-based cavity loading method which allows us to trap multiple atoms consistently, this scheme allows us to perform single-atom experiments after almost every MOT drop.

While there exist well-known methods for optically pumping atoms in free space, we have struggled to implement them in our lab. These conventional methods rely on classical fields which address both hyperfine manifolds of the atom, and we have used the lattice beams as well as linearly polarized light from the side of the cavity for this purpose. One source of our difficulties may be the Zeeman-dependent AC Stark shift which the FORT induces in the excited states of the atom [39]; this potentially leads to mixing of the excited-state Zeeman populations during the optical pumping process. In addition, recent calculations suggest that diffraction of beams focused into the cavity from the side results in significant intensity variation along the cavity axis.

After a series of frustrating optical pumping attempts in the cavity, we realized that it might be better to study our techniques first in the simpler setting of a MOT.

We rehabilitated the old lab 1 upper vacuum chamber for this purpose, and undergraduate Eric Tai is currently working with Dal in lab 9 to explore free-space optical pumping. In the meantime, we have found an alternative method of preparing atoms in a specific Zeeman state, presented in Sections 4.1 and 4.2: by driving incoherent Raman transitions across a specifically chosen frequency range, we can create a tunable dark state for population trapping. Not only does our new scheme avoid the problems associated with side beams and excited-state mixing, but also it no longer restricts us to the Zeeman levels $m_F = 0, +F, -F$ as in conventional optical pumping.

Variation in intracavity atom number is another bugbear of the experiment which in the past has been only indirectly confronted. The number of atoms loaded into the FORT after each MOT drop is believed to obey Poissonian statistics over the course of a data run. Therefore, if we can establish a maximum acceptable level of events with two or more atoms present, we can simply turn down the loading rate (and thus, the average atom number \bar{N} in the Poissonian distribution) until this threshold is met. (In practice, this is accomplished by attenuating the power in the lattice beams which cool atoms into the FORT, so that cooling is less effective.) However, we must rely on an unfortunate trade-off between two-atom contamination and data acquisition time. For example, 90.8% of all trials in the single photon generation experiment were discarded because they had no atoms, and the experiment thus required almost a day of continuous data acquisition; nevertheless, the effects of two-atom events could still be seen in the photon statistics [16, 39]. In order to decouple these two parameters, we introduce a real-time measurement of atom number which allows us to enforce the condition $N < 2$ before the experiment begins (Sections 4.3 and 4.4).

4.1 Optical pumping via incoherent Raman transitions

The following section has been adapted from Ref. [32].

A new optical pumping scheme is presented that uses incoherent Raman tran-

sitions to prepare a trapped cesium atom in a specific Zeeman state within the $6S_{1/2}, F = 3$ hyperfine manifold. An important advantage of this scheme over existing optical pumping schemes is that the atom can be prepared in any of the $F = 3$ Zeeman states. We describe an experimental implementation of the scheme and show that a fraction 0.57 ± 0.02 of the total population can be prepared in the desired state, with the remaining population distributed fairly uniformly among the six other states. We demonstrate the scheme in the context of cavity quantum electrodynamics, but the technique is equally applicable to a wide variety of atomic systems with hyperfine ground-state structure.

4.1.1 Introduction

Many experiments in atomic physics rely on the ability to prepare atoms in specific internal states. For example, spin-polarized alkali atoms can be used to polarize the nuclei of noble gases [84], to act as sensitive magnetometers [85], and to provide frequency standards that exploit magnetic-field-insensitive clock transitions [86]. In the field of quantum information science, internal atomic states can be used to store and process quantum bits [1, 2, 34, 87, 88] with extended coherence times.

A standard method for preparing an atom in a specific internal state is optical pumping [89, 90, 91], which involves driving the atom with light fields that couple to all but one of its internal states; these light fields randomly scatter the atom from one internal state to another until it falls into the uncoupled “dark” state. Various optical pumping schemes have been analyzed and demonstrated for alkali atoms [92, 86, 93] and today are well-established techniques. These schemes rely on dark states that are set by the polarization of the driving field, and this imposes restrictions on the possible Zeeman states in which the atom can be prepared. Specifically, one can prepare the atom in the $m_F = 0$ state by using light that is linearly polarized along the quantization axis, or in one of the edge states ($m_F = \pm F$) by using light that is circularly σ_{\pm} -polarized along the quantization axis.

In contrast, the scheme presented here allows the atom to be prepared in any

of the Zeeman states within the lowest ground state hyperfine manifold of an alkali atom, which in our case is the $6S_{1/2}, F = 3$ manifold of cesium. The key component of the scheme is a pair of optical fields that drive Raman transitions between pairs of Zeeman states $|3, m\rangle \leftrightarrow |4, m\rangle$. We apply a magnetic bias field to split out the individual Zeeman transitions and add broadband noise to one of the optical fields, where the spectrum of the noise is tailored such that all but one of the transitions are driven. The two Zeeman states corresponding to the undriven transition are the dark states of the system, and we exploit these dark states to perform optical pumping. We verify the optical pumping by using coherent Raman transitions to map out a Raman spectrum, which allows us to determine how the atomic population is distributed among the different Zeeman states; these measurements show that a fraction 0.57 ± 0.02 of the total population is prepared in the desired state, with the remaining population distributed fairly uniformly among the six other states. The capability of driving Raman transitions between hyperfine ground states has many additional applications, such as state manipulation [94], ground-state cooling [30, 33, 35, 37], precision measurements [95, 96], and Raman spectroscopy [97]. The scheme described here shows that this versatile tool can also be used for atomic state preparation.

We have demonstrated this scheme in the context of cavity quantum electrodynamics (QED), specifically in a system in which a single atom is strongly coupled to a high-finesse optical cavity. Cavity QED offers a powerful resource for quantum information science, and the ability to prepare the atom in a well-defined initial state is a key requirement for many of the protocols that have been proposed for this system, such as the generation of polarized single photons [4, 82] and the transfer of Zeeman coherence to photons within the cavity mode [76]. Conventional optical pumping to a single Zeeman sublevel has been previously demonstrated within a cavity [83], but we find our new method to be particularly effective given the constraints of our system, in which optical access to the atom is limited and we must address the large multiplicity of cesium sublevels. However, optical pumping via incoherent Raman transitions has much broader applications beyond the cavity QED setting and can be




Figure 4.1: Schematic of experiment. (a) View from the side of the cavity. Shown are the linearly polarized FORT, Raman, and probe beams that drive the cavity, and the circularly polarized 4 – 4' lattice beams and linearly polarized 4 – 4' side beam that drive the atom. (b) View along the cavity axis. Shown are the 4 – 4' lattice beams and the 4 – 4' side beam.

used in a wide variety of atomic systems with hyperfine ground-state structure.

4.1.2 Experimental apparatus

Our system consists of a single cesium atom that is strongly coupled to a high-finesse optical cavity, as shown in Figure 4.1. The cavity supports a set of discrete modes, and its length is tuned so that one pair of modes ¹ is nearly resonant with the atomic transition $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F = 5'$ at $\lambda_{D2} = 852 \text{ nm}$ (see the level diagram shown in Figure 4.2). The atomic dipole associated with this transition couples to the electric field of the resonant mode, allowing the atom and cavity to exchange excitation at a characteristic rate $g = (2\pi)(34 \text{ MHz})$ for the $6S_{1/2}, F = 4, m_F = 4 \rightarrow 6P_{3/2}, F = 5', m_{F'} = 5$ transition, a rate that is much larger than either the cavity decay rate $\kappa = (2\pi)(3.8 \text{ MHz})$ or the atomic decay rate $\gamma = (2\pi)(2.6 \text{ MHz})$; thus, the system is in the strong-coupling regime [49].

We hold the atom inside the cavity via a state-insensitive far off-resonance trap (FORT) [13]. The FORT is produced by resonantly driving a cavity mode at $\lambda_F = 936 \text{ nm}$ with a linearly polarized beam, which creates a red-detuned standing wave inside the cavity. Each antinode of this standing wave forms a potential well in which an atom can be trapped; for the experiments described here, the optical power of the

¹Since there are two polarization degrees of freedom, the cavity modes occur in nearly degenerate pairs.

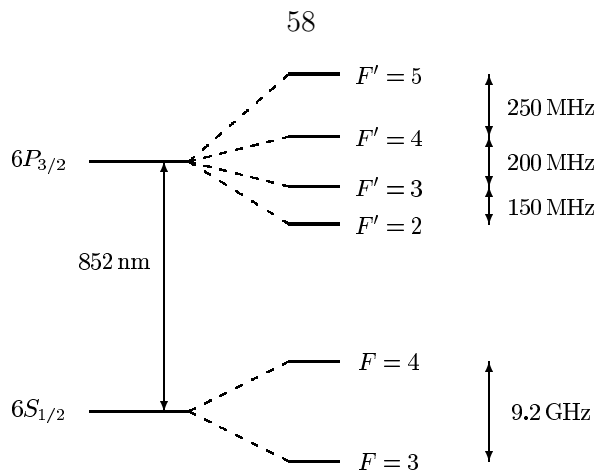


Figure 4.2: Level diagram for the $D2$ line in cesium

FORT beam is chosen such that the depth of these wells is $U_F = (2\pi)(45 \text{ MHz})$.

We drive Raman transitions between the $F = 3$ and $F = 4$ hyperfine ground-state manifolds of the atom by adding a second beam, referred to here as the Raman beam, which drives the same cavity mode as the FORT beam but is detuned from the FORT by the atomic hyperfine splitting $\Delta_{HF} = (2\pi)(9.2 \text{ GHz})$ (this scheme was first proposed in [19] and was used to perform Raman sideband cooling in [17]). The FORT and Raman beams are combined on a polarizing beam splitter (PBS) before entering the cavity, so the Raman beam is linearly polarized in a direction orthogonal to the polarization of the FORT beam. To stabilize the frequency difference between the FORT and Raman beams, the external-cavity diode laser that generates the Raman beam is injection-locked to the red sideband of light that has been picked off from the FORT beam and passed through an electro-optical modulator (EOM), which is driven at Δ_{HF} . The FORT and Raman beams form the two legs of a Raman pair and drive Raman transitions between pairs of Zeeman states $|3, m\rangle \leftrightarrow |4, m\rangle$, where the quantization axis \hat{z} is chosen to lie along the cavity axis². Typically we use a strong FORT beam and a weak Raman beam, so the Raman beam does not significantly alter the FORT trapping potential³.

²The FORT-Raman pair generates a Raman coupling between the hyperfine ground states that is proportional to $\vec{J} \cdot (\hat{\epsilon}_F \times \hat{\epsilon}_R)$, where \vec{J} is the electron angular momentum operator and $\hat{\epsilon}_F, \hat{\epsilon}_R$ are the polarization vectors for the FORT and Raman beams, so in general $\Delta m = \pm 1, 0$ transitions are possible [19]. For our system $\hat{\epsilon}_F \times \hat{\epsilon}_R = \hat{z}$, so only the $\Delta m = 0$ transitions are driven.

³The FORT and Raman beams give level shifts $U_F \sim \Omega_F^2/\Delta$ and $U_R \sim \Omega_R^2/\Delta$, and the effective




Figure 4.3: Ground state spectrum of cesium in the presence of an axial bias field. The Raman coupling drives transitions between pairs of Zeeman states $|3, m\rangle \leftrightarrow |4, m\rangle$, as indicated by the arrows.

In order to address individual Zeeman transitions, we apply a magnetic bias field B_a along the cavity axis. As shown in Figure 4.3, this axial field shifts the $|3, m\rangle \leftrightarrow |4, m\rangle$ transition by

$$\delta(|3, m\rangle \leftrightarrow |4, m\rangle) = \omega_B m, \quad (4.1)$$

where

$$\omega_B \equiv (g_4 - g_3)\mu_B B_a = (2\pi)(700 \text{ kHz/G}) B_a, \quad (4.2)$$

and $g_4 = 1/4$, $g_3 = -1/4$ are the Lande g -factors for the $F = 4$ and $F = 3$ ground-state hyperfine manifolds. For the experiments described here, we typically set the axial bias field such that $\omega_B \simeq (2\pi)(910 \text{ kHz})$.

The strong atom-cavity coupling allows us to determine whether the atom is in the $F = 3$ or $F = 4$ hyperfine manifold by driving the cavity with a $100 \mu\text{s}$ pulse of resonant $4 - 5'$ probe light, as described in [30]. If the atom is in $F = 4$, it couples to the cavity and blocks the transmission of the probe beam, while if the atom is in $F = 3$, it decouples from the cavity, and the probe beam is transmitted. Using this technique, we can determine the hyperfine ground state of the atom with an accuracy

Rabi frequency for the Raman transitions driven by the FORT-Raman pair is $\Omega_E \sim \Omega_F \Omega_R / \Delta$, where $\Omega_{F,R}$ are the Rabi frequencies of the FORT and Raman beams and Δ is the detuning from atomic resonance. Thus, the ratio of the level shifts is $U_R/U_F \sim (\Omega_E/U_F)^2 \sim 10^{-5}$ for the typical values $U_F = (2\pi)(45 \text{ MHz})$, $\Omega_E = (2\pi)(120 \text{ kHz})$.

of $\sim 98\%$ for a single $100\ \mu\text{s}$ measurement interval.

Atoms are delivered to the cavity by releasing a magneto-optical trap located a few millimeters above the cavity, and the falling atoms are loaded into the FORT by cooling them with $4 - 4'$ lattice light. This lattice light consists of two pairs of counter-propagating beams in the $\sigma_+ - \sigma_-$ configuration, which are applied from the sides of the cavity. We ensure that only one atom is trapped in the FORT by applying the Raman beam and driving the cavity with a resonant $4 - 5'$ probe; this combination gives an effect analogous to that in [18], which allows us to determine the number of atoms in the cavity based on the amount of $4 - 5'$ light that is transmitted.

4.1.3 Coherent and incoherent Raman transitions

If the FORT and Raman beams are both monochromatic, then they drive coherent Raman transitions between pairs of Zeeman states $|3, m\rangle \leftrightarrow |4, m\rangle$, and the atomic populations oscillate between the two states in each pair. The effective Rabi frequency for the $|3, m\rangle \leftrightarrow |4, m\rangle$ transition is

$$\Omega_E(|3, m\rangle \leftrightarrow |4, m\rangle) = \Omega_0 (1 - m^2/16)^{1/2}, \quad (4.3)$$

where Ω_0 is set by the power in the FORT and Raman beams [19]. For the experiments described here, the powers in these beams are chosen such that that $\Omega_0 \simeq (2\pi)(120\ \text{kHz})$. The Raman detuning for the FORT-Raman pair is given by $\delta_R = \omega_F - \omega_R - \Delta_{HF}$, where ω_F and ω_R are the optical frequencies of the FORT and Raman beams, which means that the effective detuning for the $|3, m\rangle \leftrightarrow |4, m\rangle$ transition is

$$\delta_E(|3, m\rangle \leftrightarrow |4, m\rangle) = \delta_R - \omega_B m. \quad (4.4)$$

We can also drive incoherent Raman transitions by using a monochromatic FORT beam and a spectrally broad Raman beam, where the spectral width is typically $\sim 10\ \text{MHz}$. In contrast to coherent Raman transitions, in which the atom undergoes

coherent Rabi oscillations, for incoherent Raman transitions the atomic population decays at a constant rate from $|3, m\rangle \rightarrow |4, m\rangle$ and from $|4, m\rangle \rightarrow |3, m\rangle$. In Section 4.1.7, we show that these decay rates are proportional to $S(\Delta_{HF} + \omega_B m)$, where $S(\omega)$ is the power spectrum of a beat note formed between the FORT and Raman beams.

4.1.4 Measuring the population distribution

Given an initial state of the atom in which the entire population lies in the $F = 3$ manifold, we can use coherent Raman transitions to determine how the population is distributed among the various Zeeman states. To perform this measurement we prepare the atom in the desired initial state, apply a coherent Raman pulse of fixed duration, Rabi frequency, and Raman detuning, and then drive the cavity with a resonant $F = 4 \rightarrow F = 5'$ probe beam to determine if the atom was transferred to $F = 4$. By iterating this process, we determine the probability p_4 for the atom to be transferred by the Raman pulse, and by repeating the probability measurement for different Raman detunings δ_R we can map out a Raman spectrum $p_4(\delta_R)$. For the Raman spectra presented here, the Raman pulses have Rabi frequency $\Omega_0 = (2\pi)(120 \text{ kHz})$ and duration $25 \mu\text{s}$. This is long enough that the Rabi oscillations decohere, and the Raman spectrum just records the Lorentzian envelope for each Zeeman transition.⁴ Thus, when the $|3, m\rangle \leftrightarrow |4, m\rangle$ Zeeman transition is resonantly driven by the Raman pulse, roughly half the population that was initially in $|3, m\rangle$ is transferred to $|4, m\rangle$.

As a demonstration of this technique, Figure 4.4 shows a Raman spectrum for an initial state with comparable populations in all of the $F = 3$ Zeeman states. To prepare this state, we optically pump the atom to $F = 3$ by alternating 7 pulses of resonant $F = 4 \rightarrow F = 4'$ lattice light with 7 pulses of resonant $F = 4 \rightarrow F = 4'$ side light, where each pulse is 300 ns long. The beams that deliver the lattice and side light are shown in Figure 4.1.

⁴Subsequent measurements (Section 5.2.2) have shown that decoherence times are longer over a certain frequency range on the red side of each transition. Pulses of $25 \mu\text{s}$ are in fact not quite long enough for decoherence at these special detunings, but the Lorentzian assumption still holds in the context of a broad scan.




Figure 4.4: Raman spectrum for a random initial state. Shown is the transfer probability p_4 vs. Raman detuning δ_R : the points are the experimental data, the curve is a fit of $p_4(\delta_R)$, as given by equation (4.5), and the vertical green lines indicate the predicted frequencies $\delta(|3, m\rangle \leftrightarrow |4, m\rangle)$ for individual Zeeman transitions.

To determine the population $p_{3,m}$ in the Zeeman state $|3, m\rangle$, we fit a sum of Lorentzians, one for each Zeeman transition, to the experimental data:

$$p_4(\delta_R) = p_b + (1/2) \sum_m (1 + (\delta_R - \omega_B m)^2 / (1 - m^2/16) \Omega_0^2)^{-1} p_{3,m}, \quad (4.5)$$

where p_b is a constant background. We fit the Zeeman state populations, the Rabi frequency Ω_0 , and the frequency ω_B that characterizes the strength of the axial bias field, and perform an independent measurement to determine the background probability $p_b = 0.006$. The fitted value of Ω_0 agrees to within 14% with the value we would expect based on the measured optical powers in the FORT and Raman beams, and the fitted value of ω_B agrees to within 5% with the value we would expect based on the known axial coil current and geometry. As a consistency check, we sum the fitted populations and obtain the result 1.10 ± 0.03 , in reasonable agreement with the expected value of 1.

4.1.5 Optical pumping scheme

We can prepare the atom in a specific Zeeman state by using a Raman beam whose spectrum is tailored to drive incoherently all but one of the Zeeman transitions. As




Figure 4.5: (a) Power spectrum of noise used for pumping into $|3,0\rangle$. (b) Power spectrum of coherent signal used for driving coherent Raman transitions with $\Omega_0 = (2\pi)(120 \text{ kHz})$. Both curves are obtained by combining the FORT and Raman beams on a photodetector and measuring the spectrum of the photocurrent, mixed down from $\Delta_{HF} = 9.2 \text{ GHz}$; shown is the RF power in a 3 kHz bandwidth vs. detuning from Δ_{HF} .

an example, Figure 4.5a shows the power spectrum of the noise used for pumping into $|3,0\rangle$. This graph was obtained by measuring the power spectrum of a beat note formed between the FORT and Raman beams by mixing them on a photodetector with a non-polarizing beam splitter. For comparison, Figure 4.5b shows the power spectrum for a monochromatic Raman beam tuned to Raman resonance, as would be used for driving coherent Raman transitions.

Comparing the noise spectrum shown in Figure 4.5a to the Raman spectrum shown in Figure 4.4, we see that the noise drives incoherent Raman transitions from $|3, m\rangle \leftrightarrow |4, m\rangle$ for $m \neq 0$, but because of the notch around zero detuning, the $|3,0\rangle \leftrightarrow |4,0\rangle$ transition is not driven. We optically pump the atom into $|3,0\rangle$ by first driving incoherent Raman transitions for $10 \mu\text{s}$, then pumping the atom to $F = 3$ using the method discussed in Section 4.1.4, and iterating this sequence 40 times. It is straightforward to modify this procedure so as to pump into the $|3, m\rangle$ Zeeman state for any m ; we simply shift the notch in the noise so that it overlaps with the $|3, m\rangle \leftrightarrow |4, m\rangle$ transition.

To characterize the optical pumping, we first pump the atom into a specific Zeeman state and then measure the Raman spectrum as described in the preceding section.

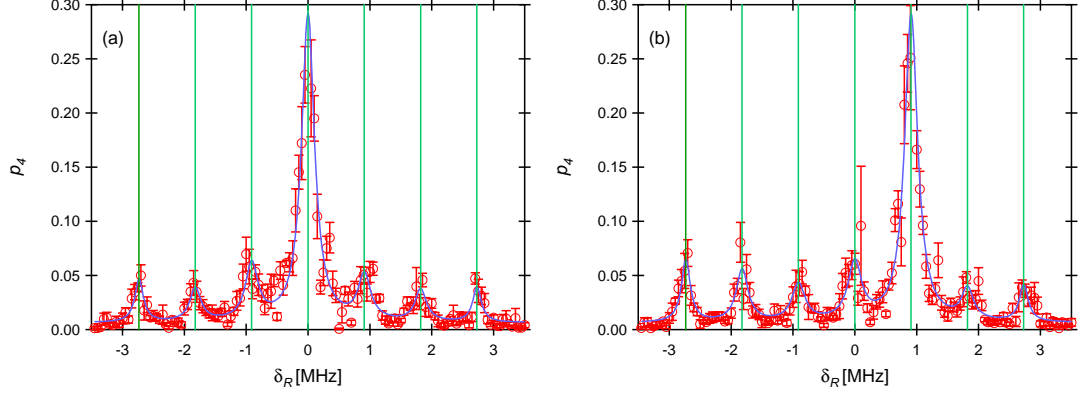


Figure 4.6: (a) Raman spectrum for optical pumping into $|3, 0\rangle$. (b) Raman spectrum for optical pumping into $|3, 1\rangle$. Raman spectrum for a random initial state. Shown is the transfer probability p_4 vs. Raman detuning δ_R : the points are the experimental data, the curve is a fit of $p_4(\delta_R)$, as given by equation (4.5), and the vertical green lines indicate the predicted frequencies $\delta(|3, m\rangle \leftrightarrow |4, m\rangle)$ for individual Zeeman transitions.

Figure 4.6 shows Raman spectra measured after pumping into (a) $|3, 0\rangle$ and (b) $|3, 1\rangle$. We find that the fraction of the atomic population in the desired state is 0.57 ± 0.02 for pumping into $|3, 0\rangle$ and 0.57 ± 0.02 for pumping into $|3, 1\rangle$, where the remaining population is roughly equally distributed among the other Zeeman states (these numbers are obtained by fitting equation (4.5) to the data, as described in Section 4.1.4). Summing the fitted populations in all the Zeeman states, we obtain the value 1.02 ± 0.04 for (a) and 1.08 ± 0.04 for (b), in reasonable agreement with the expected value of 1.

To generate the Raman beam used in Figure 4.5a, we start with an RF noise source, which produces broadband noise that is spectrally flat from DC to ~ 10 MHz. The noise is passed through a high-pass filter at 500 kHz and a low-pass filter at 5 MHz, where both filters roll off at 60 dB per octave. The filtered noise is then mixed against an 85 MHz local oscillator, and the resulting RF signal is used to drive an acousto-optical modulator (AOM) that modulates a coherent beam from the injection-locked Raman laser. The first-order diffracted beam from the AOM forms a Raman beam with the desired optical spectrum. Note that previous work has demonstrated the use of both synthesized incoherent laser fields [98, 99], such as that used here, as well as

the noise intrinsic to free-running diode lasers [100, 101] to resonantly probe atomic spectra.

Although the scheme presented here relies on incoherent Raman transitions, it is also possible to perform optical pumping with coherent Raman transitions. The basic principle is the same: we simultaneously drive all but one of the Zeeman transitions, only instead of using a spectrally broad Raman beam, we use six monochromatic Raman beams, where each beam is tuned so as to resonantly drive a different transition. We have implemented such a scheme, and found that it gives comparable results to the incoherent scheme described above, but there are two advantages to the incoherent scheme. First, it is simpler to generate a Raman beam with the necessary spectral properties for the incoherent scheme. Second, when coherent Raman transitions are used, the six frequency components for the Raman beam must be tuned to resonance with their respective transitions, and hence are sensitive to the value of the axial magnetic field. When incoherent Raman transitions are used, however, the same Raman beam can be used for a broad range of axial field values.

4.1.6 Conclusion

We have proposed a new scheme for optically pumping atoms into a specific Zeeman state and have experimentally implemented the scheme with cesium atoms in a cavity QED setting. An important advantage over existing schemes is that atoms can be prepared in any of the Zeeman states in the lower hyperfine ground state manifold.

We have measured the effectiveness of the optical pumping, and have shown that a fraction ~ 0.57 of the atomic population can be prepared in the desired Zeeman state. Some possible factors that could be limiting the effectiveness of the optical pumping include fluctuating magnetic fields transverse to the cavity axis, misalignment of the cavity axis with the axial bias field, and slow leaking out of the dark state due to scattering from background light. We are currently investigating these factors.

The scheme presented here operates on a fundamentally different principle from existing optical pumping schemes, in that it relies on incoherent Raman transitions

to create an atomic dark state. Raman transitions have many different applications in atomic physics, so there are often independent reasons for incorporating a system for driving Raman transitions into an atomic physics laboratory; our scheme shows that such a system can also be applied to the problem of atomic state preparation. The scheme should serve as a useful tool for experiments in atomic physics, both in a cavity QED setting and beyond.

4.1.7 Transition rate for incoherent Raman transitions

As described in Section 4.1.3, we drive incoherent Raman transitions between pairs of Zeeman states $|3, m\rangle \leftrightarrow |4, m\rangle$ by using a monochromatic FORT beam and a spectrally broad Raman beam. For incoherent Raman transitions the atomic population decays at a constant rate from $|3, m\rangle \rightarrow |4, m\rangle$ and from $|4, m\rangle \rightarrow |3, m\rangle$, and in this section we calculate these decay rates.

We will consider a single Zeeman transition $|3, m\rangle \leftrightarrow |4, m\rangle$, so we can treat the system as an effective two-level atom with ground state $g \equiv |3, m\rangle$ and excited state $e \equiv |4, m\rangle$, where the energy splitting between g and e is $\omega_A \equiv \Delta_{HF} + \omega_B m$. The FORT-Raman pair drives this effective two-level atom with broadband noise, which we can approximate as a comb of classical fields with optical frequencies ω_k and Rabi frequencies Ω_k . Let us assume that we start in the ground state g . If we only consider the coupling of the atom to field k , then the equation of motion for the excited state amplitude c_e is

$$i\dot{c}_e = \frac{\Omega_k}{2} e^{-i\delta_k t} c_g, \quad (4.6)$$

where $\delta_k \equiv \omega_k - \omega_A$ is the detuning of the field from the atom. At small times the population is almost entirely in the ground state, so we can make the approximation $c_g = 1$ and integrate equation (4.6) to obtain

$$c_e(t) = \frac{\Omega_k}{2\delta_k} (e^{-i\delta_k t} - 1). \quad (4.7)$$

Thus, the transition rate from g to e for a single frequency ω_k is

$$\gamma_k = \frac{|c_e(t)|^2}{t} = \frac{\pi}{4} t \Omega_k^2 D(\delta_k t/2), \quad (4.8)$$

where

$$D(x) \equiv \frac{\sin^2 x}{\pi x^2}. \quad (4.9)$$

The total decay rate is obtained by summing the decay rates for all the fields in the comb:

$$\gamma = \sum_k \gamma_k = \frac{\pi}{4} t \sum_k \Omega_k^2 D(\delta_k t/2). \quad (4.10)$$

To evaluate this expression we need to know the distribution of Rabi frequencies Ω_k . This information can be obtained by forming a beat note between the FORT and Raman beams on a photodetector, and measuring the power spectrum $S(\omega)$ of the photocurrent using a spectrum analyzer. Let us first consider this measurement for a monochromatic Raman beam, and then generalize to a spectrally broad Raman beam. If both the FORT and Raman beams are monochromatic, with optical frequencies ω_F and ω_R , then the resulting photocurrent $i(t)$ is given by

$$i(t) = i_F + i_R + 2\eta \cos((\omega_F - \omega_R)t) \sqrt{i_F i_R}, \quad (4.11)$$

where i_F and i_R are the cycle-averaged photocurrents for the FORT and Raman beams taken individually and η is the heterodyne efficiency. Thus, the power spectrum of the photocurrent has a spike at the difference frequency $\Delta \equiv \omega_F - \omega_R$:

$$S_c(\omega) = P_c \delta(\omega - \Delta), \quad (4.12)$$

where the integrated power P_c of the spike is proportional to $i_F i_R$. If the difference frequency Δ is tuned to Raman resonance ($\Delta = \omega_A$), then the FORT-Raman pair

drives coherent Raman transitions with a Rabi frequency Ω_c that is proportional to $\sqrt{i_F i_R}$, so

$$\Omega_c^2 = \alpha P_c, \quad (4.13)$$

where α is a constant that depends on various calibration factors.

Now consider the case of a spectrally broad Raman beam, which results in a photocurrent with power spectrum $S_i(\omega)$. The effective Rabi frequency Ω_k corresponding to comb line k is given by

$$\Omega_k^2 = \alpha S_i(\omega_k) \delta\omega, \quad (4.14)$$

where $\delta\omega$ is the frequency spacing between adjacent comb lines. Substituting this result into equation (4.10), and replacing the sum with an integral, we obtain

$$\gamma = \frac{\pi}{4} \alpha t \int S_i(\omega) D((\omega - \omega_A)t/2) d\omega. \quad (4.15)$$

If the power spectrum near ω_A is flat over a bandwidth $\sim 1/t$, then we can approximate D as a delta function and perform the integral:

$$\gamma = \frac{\pi}{2} \alpha S_i(\omega_A). \quad (4.16)$$

It is convenient to use equation (4.13) to eliminate the calibration factor α :

$$\gamma = \frac{\pi}{2} \frac{S_i(\omega_A)}{P_c} \Omega_c^2. \quad (4.17)$$

The spectrum analyzer trace given in Figure 4.5a displays the power spectrum in terms of the power $P_i(\nu) \simeq 2\pi B S_i(\omega)$ in a bandwidth $B = 3$ kHz, so we can also write this as

$$\gamma = \frac{1}{4} \frac{P_i(\omega_A/2\pi)}{P_c} \frac{\Omega_c^2}{B} = \frac{1}{4} (1 - m^2/16) \frac{\Omega_0^2}{B} \frac{P_i((\Delta_{HF} + \omega_B m)/2\pi)}{P_c}, \quad (4.18)$$

where we have substituted $\Omega_c = (1 - m^2/16)^{1/2} \Omega_0$ and $\omega_A = \Delta_{HF} + \omega_B m$.

We can calculate the time evolution of the atomic populations using rate equations. It is straightforward to show that the decay rate $e \rightarrow g$ is also given by γ , and from the rate equations one can show that the excited state population is

$$p_e(t) = \frac{1}{2}(1 - \exp(-2\gamma t)). \quad (4.19)$$

We can calculate the decay rates for the noise spectrum shown in Figure 4.5. For this noise spectrum the power $P_i(\nu)$ has roughly the same value \bar{P}_i at the frequencies of all the $m \neq 0$ Zeeman transitions, so we can write the decay rates for these transitions as

$$\gamma(|3, m\rangle \rightarrow |4, m\rangle) = \gamma(|4, m\rangle \rightarrow |3, m\rangle) = (1 - m^2/16) \Gamma, \quad (4.20)$$

where

$$\Gamma \equiv (1/4)(\Omega_0^2/B)(\bar{P}_i/P_c). \quad (4.21)$$

From the power spectrum for the noise shown in Figure 4.5a we have that $\bar{P}_i = -63$ dBm, and from the power spectrum for the coherent signal shown in Figure 4.5b we have that $P_c = -36$ dBm, where the corresponding Rabi frequency is $\Omega_0 = (2\pi)(120 \text{ kHz})$. Substituting these values into equation (4.21), we obtain $\Gamma = 0.084 \mu\text{s}^{-1}$.

4.2 Optimizing optical pumping

The discussion in the previous section centers on an example in which a magnetic field is applied along the cavity axis, so that only $\Delta m = 0$ Raman transitions are permitted. It is important to stress, however, that Raman-based optical pumping also works for an arbitrary applied field, in which $\Delta m = \pm 1$ transitions are also possible. The point is that given an applied field of any direction and magnitude, it will have an associated Raman spectrum which reflects transitions between different pairs of

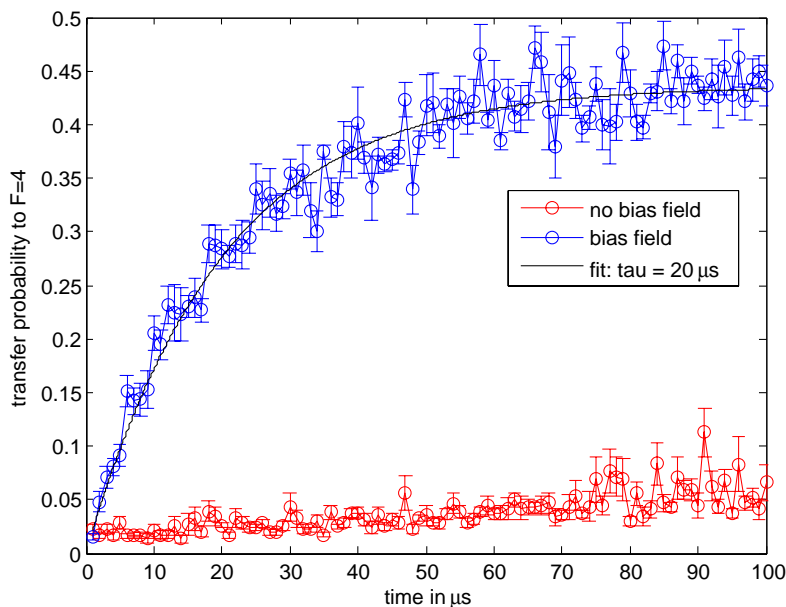


Figure 4.7: Rate of population transfer for incoherent Raman transitions: atoms are prepared in $F = 3$ and then driven with the notched incoherent Raman spectrum of Figure 4.5(a), in the presence (blue) or absence (red) of a magnetic bias field. The probability to transfer an atom to $F = 4$ is shown as a function of pumping time.

Zeeman levels. Based on this Raman spectrum, one can tailor a noise spectrum such as 4.5(a) which will address transitions between all but one Zeeman level.

In implementing our optical pumping scheme, we took several steps to optimize the relevant parameters. First, we characterized the preparation of atoms in $F = 3$ via lattice light and linearly polarized side light; we found that by using short (300 ns) interleaved pulses and resonant light, we could transfer 99.9% of the population to the $F = 3$ manifold in 10 μs .⁵ We next measured the rate at which the incoherent Raman spectrum transferred population from $F = 3$ to $F = 4$. This is shown in Figure 4.7, where atoms are prepared in the $F = 3$ manifold and then subjected to incoherent Raman pulses of varying duration, either in the presence or absence of a

⁵Switching between two fields with different polarizations is necessary to avoid creating a permanent dark state in the $F = 4$ manifold. An alternate but slightly less efficient scheme is to drive both lasers continuously but with a small relative detuning; the detuning then sets a precession rate for the dark state. In the case of an applied magnetic field, only one $F = 4 \rightarrow F' = 4$ laser is necessary, since now the atom's dark state changes as it travels through the field.

bias field. With the magnetic fields nulled, we expect that the the notched Raman spectrum will be unable to drive any resonant transitions between Zeeman states, and the slow leakage we observe into $F = 4$ is probably due to off-resonant driving from the edges of the notch. In the presence of a bias field, the Raman pulses should be able to address atoms in every Zeeman level but $|3, 0\rangle$ and to transfer population until an equilibrium is reached. We see that this population transfer takes place with time constant $\tau = 20 \mu\text{s}$ and that it asymptotes to a state in which approximately 45% of the population is in $F = 4$; presumably, another 45% resides in $|3, m_F \neq 0\rangle$ and the final 10% remains in $|3, 0\rangle$. The measured value of τ is in reasonable agreement with the calculated value $\Gamma = 0.084 \mu\text{s}^{-1}$ in Section 4.1.7, which should be weighted by the initial distribution of atoms among Zeeman manifolds as in equation (4.20).

The time constants for these two processes — repumping to $F = 3$ and incoherent Raman transfer to $|4, m_F \neq 0\rangle$ — allowed us to design a series of pulses in which population is shuttled back and forth between the two hyperfine manifolds while gradually accumulating in $|3, 0\rangle$. In order to optimize both the pulse durations and the number of iterations, we measured the final population transferred to $|3, 0\rangle$ as a function of these parameters; we settled on 10 μs of incoherent Raman light followed by 4.2 μs of $F = 4 \rightarrow F' = 4$ repump fields (7 cycles of paired 300 ns pulses), repeated 40 times. Thus, the time required for optical pumping of each atom is roughly 600 μs .

We also investigated the fraction of population that we could prepare in $|3, 0\rangle$ as a function of the applied axial magnetic field. Recall that our incoherent Raman spectrum has a low-frequency cutoff at $(2\pi)(500 \text{ kHz})$ and a high-frequency cutoff at $(2\pi)(5 \text{ MHz})$. In the limit of very small axial fields, all of the $\Delta m = 0$ transitions are contained within the low-frequency notch; we expect that the Raman spectrum will be unable to transfer population to $F = 4$, and optical pumping will be ineffective. However, also for small applied fields, the $|3, 0\rangle \rightarrow |4, 0\rangle$ Raman pulse that we use to measure population should in fact address all Zeeman states. Thus, as we see in Figure 4.8, the measured population transfer is very high at the smallest field values. As the $\Delta m = 0$ transitions with $m \neq 0$ are pushed out of resonance with the

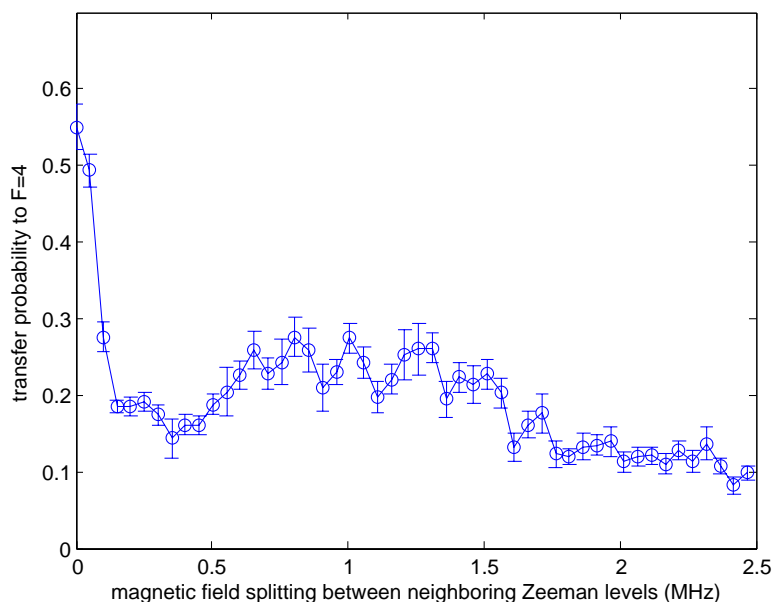


Figure 4.8: We measure the success of our Raman optical pumping protocol as a function of axial magnetic bias field. After optical pumping, we apply a Raman pulse to transfer population from $|3, 0\rangle$ to $|4, 0\rangle$, followed by state detection. For small fields, the pulse addresses all Zeeman levels, while at larger values, the number of Zeeman transitions inside the Raman envelope is field-dependent.

$|3, 0\rangle \rightarrow |4, 0\rangle$ detection pulse, this value declines steeply. It then rises again for field splittings greater than $(2\pi)(0.5 \text{ MHz})$, as one by one, Raman transition frequencies are pushed outside the notch and can now be driven by the incoherent Raman fields. Finally, at roughly $(2\pi)(\frac{1}{3} \times 5 \text{ MHz})$, the $|3, \pm 3\rangle \rightarrow |4, \pm 3\rangle$ transitions are pushed past the high-frequency cutoff and again out of range of the incoherent Raman; here the transfer probability begins again to decline. Based on this measurement, we set our axial field to generate a splitting of $(2\pi)(910 \text{ kHz})$ between neighboring Zeeman levels. However, it is clear that optical pumping should be effective over $\sim 1 \text{ MHz}$ of field values, and we could of course use different filters in order to access a wider range.

Despite our efforts at optimization, we are able to transfer at best 60% of atoms into a target Zeeman state. We remain unsure as to what limits our pumping effi-

ciency, although we have narrowed the list of possible suspects. We have ruled out scattered light from stray fields, as we have seen that atoms pumped into $F = 3$ only leak into $F = 4$ very slowly (Figure 4.7). The small peak at 85 MHz in the power spectrum shown in Figure 4.5(a), due to carrier frequency leakage through our RF mixer, is also not a cause for concern; we have confirmed that the optical pumping efficiency does not change when we shift the center frequency of the spectrum out of resonance with the $|3, 0\rangle \rightarrow |4, 0\rangle$ transition. We also have replaced the 500 kHz low-frequency notch with one at 100 kHz with no observed change in pumping efficiency, suggesting that off-resonant excitation of $|3, 0\rangle \rightarrow |4, 0\rangle$ due to the finite notch width is not the problem. Remaining candidates include fluctuating magnetic fields and the finite suppression (~ 20 dB) of Raman power which the notch provides.

A final unresolved question concerns the timing of our optical pumping scheme. We would expect that atoms would be pumped to the dark state $|3, 0\rangle$ by applying the incoherent Raman pair and the $F = 4 \rightarrow F' = 4$ repumping light simultaneously. However, in practice we have found that this is much less successful than the pulsed scheme and that any overlap between the Raman and repump pulses degrades the scheme's effectiveness.

4.3 Trapping and detecting multiple atoms

4.3.1 Cavity loading

Through mid-2007, we relied on optical lattice Sisyphus cooling to load atoms into the FORT after their release above the cavity [13]. The details of this method are as follows: 30 ms after the end of polarization-gradient cooling in the lower MOT, when we expect the atoms to be falling through the cavity mode, we turn on $F = 3 \rightarrow F = 3'$ and $F = 4 \rightarrow F = 4'$ lattice beams for 5 ms, where both beams are blue-detuned by 10 MHz. The best observed loading probabilities for ≥ 1 atoms are $\sim 60\%$ with this technique, in which case (assuming Poissonian loading statistics) $\sim 30\%$ of all loading attempts would consist of single atoms. We hypothesize that one of the two

lattice beams provides Sisyphus cooling [67], while the other acts as a repump (and thus in principle does not need to be in the standing-wave lattice configuration). In experiments where the cavity is tuned to the $F = 4 \rightarrow F = 4'$ transition, we have found that we cannot load efficiently using $F = 4 \rightarrow F = 4'$ lattice light; we believe that this is due to resonant scattering into the cavity mode. Instead, in this case we employ the $F = 3 \rightarrow F = 3'$ lattice beam in conjunction with a resonant probe beam along the cavity axis, which presumably repumps the atoms with less scatter-induced heating.

We have recently found that by instead using the FORT-Raman pair during the 5 ms loading window in conjunction with the $F = 4 \rightarrow F = 4'$ lattice beam, over 99% of all MOT drops result in cavity loading. (Note that as the intracavity FORT is always present, the beams we turn on for loading are the Raman beam and the $F = 4 \rightarrow F = 4'$ lattice.) In contrast with the previous method, which required a dense, bright lower MOT, here we observe excellent loading even when the lower MOT quality is visibly poor. Furthermore, the measurements described below confirm that we are loading multiple atoms almost every time.

We hypothesize that by replacing the $F = 3 \rightarrow F = 3'$ beam with Raman transitions, we are still able to recycle the atom to $F = 4$ while avoiding the spontaneous emission inherent to a $F = 3 \rightarrow F = 3'$ repump; the $F = 4 \rightarrow F = 4'$ lattice beam then provides the necessary Sisyphus cooling. Repumping via Raman transitions thus results in more efficient atom cooling, because spontaneous emission was previously heating the atom throughout the cooling process. Loading efficiency is found to increase as a function of Raman beam intensity, presumably because it allows faster recycling, and we currently use a Rabi frequency of 2.2 MHz for atom loading. We have explored loading efficiency of the Raman scheme as a function of various parameters, as discussed in Section 4.3.4.

Despite the significant improvement in loading efficiency, we do not observe direct loading from the upper MOT (with the timing sequence adjusted appropriately, so that the cooling beams address the atoms while they are within the cavity). In the process of setting up the lab 11 experiment, Joe Buck and Jason McKeever observed

transits from atoms dropped from both the upper and lower MOTs, with narrower transits corresponding to the faster velocities of the upper MOT atoms [102], but we no longer see upper MOT transits. Presumably this effect is highly sensitive to the position of the upper MOT with respect to the gap between the cavity mirrors, as well as to the forces which the falling atoms experience as the magnetic fields are ramped off.

4.3.2 Raman-based determination of atom number

The “Raman repump” technique outlined above can be used not only to load atoms into the cavity but also to determine the number of atoms present in real time. In order to explain this measurement, we first consider the “by the numbers” effect originally demonstrated in Ref. [18]. In that experiment, the cavity was tuned to the $F = 4 \rightarrow F = 4'$ transition and probed continuously with a resonant on-axis beam in tandem with a $F = 3 \rightarrow F = 3'$ lattice beam. While the vacuum-Rabi splitting tells us that the presence of a single two-level atom in our strongly coupled cavity will completely suppress the transmission of a resonant probe beam, the two-level model is insufficient for a non-cycling transition such as $F = 4 \rightarrow F = 4'$. An atom may be in the $F = 4$ manifold, in which case it participates in the cavity-QED interaction, or it may be in $F = 3$, in which case the probe laser sees an empty cavity. With the introduction of a lattice beam as a repump, the atom moves back and forth rapidly between these two possibilities, and over timescales $\gtrsim 1$ ms, the cavity transmission represents a time-average of high ($F = 3$) and low ($F = 4$) values. When $N > 1$ atoms are present, there is a greater probability that at least one atom will be in $F = 4$ and that transmission will be suppressed. Thus, as atoms leave the cavity one by one, we observe a climbing “stair-step” transmission with distinct plateaus corresponding to atom number.

In addition to the initial $F = 4 \rightarrow F = 4'$ measurements [18], we have subsequently observed this “by the numbers” effect with the cavity and probe tuned to the $F = 4 \rightarrow F = 3'$ transition (though with worse signal-to-noise and shorter atom lifetimes

in the cavity). Jason suggested in his thesis [39] that this technique could be used to prepare one atom deterministically in the cavity, either by actively discarding cases where $N > 1$ or by heating out the additional atoms. Unfortunately, most of our experiments to date have taken place with the cavity tuned to the $F = 4 \rightarrow F' = 5$ resonance in order to exploit the physics of the cycling transition, whereas “by the numbers” relies on a non-cycling transition for atom detection.

Raman transitions between the cesium ground states allow us to have our cake and eat it too, or in this case, to implement atom number determination on a cycling transition. The key concept here is that while the trapped atoms would ordinarily remain in the $F = 4$ manifold, we can use the FORT-Raman combination to cycle each atom between $F = 4$ and $F = 3$ at the Rabi frequency of the Raman pair. As in the original experiment, the probability to have at least one atom in $F = 4$ increases with atom number. The result is the same “telegraph signal” of high (all atoms in $F = 3$) and low (one or more atoms in $F = 4$) values, which we observe as an averaged stair-step pattern. One might imagine that the atom-number plateaus would be noisy due to this averaging; in Ref. [18], their smoothness is due to the fact that the timescale of decay to $F = 3$ and recycling to $F = 4$ ($\sim \gamma^{-1}$, ten of nanoseconds), is much shorter than the lifetime of atoms in the cavity (hundreds of milliseconds), and that the data are filtered by the heterodyne detection bandwidth of 1 kHz. For Raman-based atom number determination, cycling between atomic levels also needs to be fast with respect to the number determination time window. In practice, as for atom loading, we operate with a Rabi frequency of 2.2 MHz. As we have replaced heterodyne detection with photon counting, we bin the photon counts on a time frame that is short compared with atom lifetimes in the cavity.

Figure 4.9 shows typical traces from a data set of 1,200 MOT drops in which the FORT-Raman pair and a blue-detuned $F = 4 \rightarrow F' = 4'$ lattice were used to load atoms into a cavity tuned to the $F = 4 \rightarrow F' = 5$ transition, followed by 280 ms of the Raman pair in conjunction with a probe resonant with the cavity. Photon counts are binned in 1 ms intervals for the data presented. Here we observe the expected stair-step pattern and well-defined plateaus for single atoms, suggesting that $t_{detect} = 1$ ms




Figure 4.9: Two sample traces from a data set of 1,200 MOT drops: the FORT-Raman pair and $F = 4 \rightarrow F = 4'$ lattice are used to load multiple atoms into a cavity on the $F = 4 \rightarrow F = 5'$ transition, followed by 280 ms of Raman and a resonant probe beam on the cavity axis. Photon counts are binned in 1 ms intervals.

should be sufficient in order to determine $N \leq 1$.

Figure 4.10 shows a histogram of counts per 1 ms bin for the entire data set. Because the histogram is dominated by multi-atom (highly suppressed) events, the second plot shows a closer view of the y-axis near the origin in order to resolve zero-, one-, and two-atom peaks clearly. We observe that the one-atom peak is easy to distinguish from the zero-atom and two-atom cases, and that it is also plausible that we could separate out two-atom events. More quantitatively, we can fit the data to a sum of Gaussians corresponding to photon count distributions for each atom number, as shown in Figure 4.11; the center of each Gaussian is then the mean number of counts for a given atom number plateau. (We initially attempted to fit the data as a sum of Poissonians, but it was clear that the distribution widths were greater than \sqrt{N} fluctuations alone could explain. Other contributions to the spread may include noise on the cavity lock and variation in the probe intensity over time.)

Note that the mechanism which determines the level spacing between N -atom plateaus is different in the lattice (non-cycling) and Raman (cycling) cases. For the lattice, the intensity of an N -atom plateau is given by $\frac{I_0}{\sum_{k=0}^N (k!)^2 y^k}$, where the sum is from $k = 0$ to N and I_0 is the intensity of the zero-atom plateau. Here y is the ratio $\frac{\gamma_{3 \rightarrow 4}}{\gamma_{4 \rightarrow 3}}$,




Figure 4.10: Histogram of photon counts per 1 ms time bin for the data set of Figure 4.9, with 280 ms of continuous Raman and probe beams. The data are shown twice, with a smaller range in the second plot.

where $\gamma_{n \rightarrow m}$ is the rate at which one atom is transferred from level n to level m . In Ref. [18], these two rates could in principle be controlled independently by adjusting the intensities of the lattice and probe beams. In the current experiment, however, the probability that we will measure an atom driven by Raman transitions to be in $F = 3$ is simply $\frac{1}{2}$, and the probability to find N atoms in $F = 3$ is $(\frac{1}{2})^N$. Thus we expect $I_1 = \frac{I_0}{2}$, $I_2 = \frac{I_0}{4}$ and in general,

$$I_N = \frac{I_0}{2^N}. \quad (4.22)$$

The fits in Figure 4.11 are roughly consistent with this expectation, although we have observed day-to-day variation of $\sim 20\%$ in the peak ratios. For the purposes of atom discrimination, we are only concerned with setting a strict enough lower bound so that the case $N \geq 2$ is always excluded.

In Figure 4.12, the histogram data of Figure 4.10 are separated out by time bin. This plot displays the same basic features as Figure 3 of Ref. [18], namely (a) clustering of data in well-defined atom-number plateaus, and (b) time-dependent evolution of the plateaus from $N \geq 3$ to $N = 0$ as atoms leave the trap. In Ref. [18], we observe a clear departure of atoms from one plateau to the next, and by the end of the probing interval, the trap population resides entirely in the $N = 0$ plateau.




Figure 4.11: Fit of the data in Figure 4.10 to a sum of four Gaussians. The data have been truncated for small count rates in order to consider atom number $N \leq 3$.

Here, in contrast, while some traces show the expected progression through $N = 1$ to $N = 0$, there are many cases for which $N \geq 3$ even at the end of the probing interval. The distribution of counts around these high-atom-number events narrows as a function of time, which might suggest that poorly coupled atoms leave the trap quickly while well-coupled ones remain. Furthermore, the data at late time intervals ($t > 150$ ms) indicate a direct progression from the multi-atom case to $N = 0$; that is, atoms may sometimes leave in rapid succession due to a “catastrophic event” rather than independently.

We can further explore the time dependence of atom populations by assigning approximate atom-number boundaries to the histogram data. We have seen that the $N = 0$ plateau (Figure 4.11) is centered at 335 counts per ms bin and that we expect this plateau value to scale as 2^{-N} . We define four regions as follows:




Figure 4.12: Photon count rate histograms (y axis) as a function of time bin (x axis) for the data set of Figure 4.9

Atom number	Min. counts per bin	Max. counts per bin
0	$\frac{3}{4} * 335$	
1	$\frac{3}{8} * 335$	$\frac{3}{4} * 335$
2	$\frac{3}{16} * 335$	$\frac{3}{8} * 335$
≥ 3		$\frac{3}{16} * 335$

By summing the number of counts within each region for every time bin, we can then plot atom populations as a function of time, as shown in Figure 4.13. Again, we find a point of comparison in the equivalent plot (Figure 4) of Ref. [18]. For that data set, the $N \geq 3$ population exhibited a steady decay from its initial value at $t = 0$. The $N = 2$ and $N = 1$ populations first peaked, then decayed, with the $N = 2$ peak occurring prior to that of $N = 1$. The $N = 0$ population grew steadily over time, while the other three populations had almost reached zero by the end of the interval. In our current Raman data set, we begin almost every interval with $N \geq 3$ atoms and rarely with only two. After 250 ms, the most likely case is that we have no atoms,




Figure 4.13: Time evolution of the N -atom populations for for the data set of Figure 4.9

but the second most likely case is that we still have $N \geq 3$ atoms. Furthermore, despite the continual departure of atoms from the $N \geq 3$ region, there is no growth in the $N = 1$ and $N = 2$ populations after early times. Only the $N = 0$ population is increasing, and its growth can not be accounted for solely by the departure of atoms from $N = 1$.

4.3.3 Simulations

Russ has written a Matlab simulation in which N atoms (on average, with Poissonian statistics) are loaded into the trap, and at any time after loading, there is a fixed probability for an atom to leave. Transmission is given by equation (4.22), with Poissonian noise added. While the time dependence found in these simulations mirrors the observations of Ref. [18], we are not able to adjust the simulation parameters to account for the long multi-atom dwell times we observe in Figures 4.12 and 4.13. However, we do observe qualitative agreement with a second simulation, in which the likelihood for atoms to leave the cavity is described by a simple “catastrophic

event” model; that is, a model in which the initial departure of one atom increases the probability for all other atoms to leave.

A physical basis for such catastrophic events may be the variation in g over the range of possible FORT wells. Atoms in the best wells couple strongly to the cavity and completely suppress a resonant probe beam, while atoms in poorly coupled wells allow some resonant light from the probe to enter the cavity, and this intracavity field can now heat the atoms out of the trap. We can imagine a scenario in which multiple atoms are trapped in the FORT over a distribution of wells: if only one atom is well-coupled, then it functions as a gatekeeper. Once it leaves, however, probe light enters the cavity, initiating a cascade in which the remaining atoms leave rapidly.

4.3.4 Optimizing atom loading

Atom loading takes place within a 5 ms window during which we apply both lattice light from the side of the cavity and Raman beams along the cavity axis. Loading parameters thus include which lattice beams are used ($F = 3 \rightarrow F' = 3$, $F = 4 \rightarrow F' = 4$, or both), and the frequency detuning and intensity of these beams. Assuming that our goal is to load as many atoms as possible, we would like to have a means to quantify the number of atoms loaded as a function of these parameters. Unfortunately, because by-the-numbers plateaus scale as 2^{-N} , they provide good discrimination for the $N = 1$ case but do not allow us to distinguish easily between, for example, 5 and 6 atoms in the cavity.

We introduce the following measurement to compare loading parameters: after loading multiple atoms, we first apply $F = 4 \rightarrow F' = 4$ lattice light to pump all atoms into the $F = 3$ manifold. We then drive a short Raman pulse which we expect to transfer the population of each atom to $F = 4$ about 10–20% of the time; for example, for the data presented in Figure 4.15, we use a Rabi frequency of 140 kHz corresponding to a π pulse time of 3.5 μ s, and we drive a 700 ns pulse, where $\sin^2(\frac{\pi}{2} * \frac{0.7}{3.5}) \approx 0.1$. Finally, we apply the state detection scheme of Ref. [30] in order to read out with a probe beam whether any atom is in $F = 4$. The idea is that the

probability for one or more atoms to be in $F = 4$ will scale linearly with the number of atoms in the cavity.

Each state detection measurement has a binary result: the atom is either in $F = 3$ (zero) or $F = 4$ (one). We perform a series of 1500 measurements per atom, where each measurement lasts about $110 \mu\text{s}$, including $5 \mu\text{s}$ of pumping to $F = 3$, one 700 ns Raman pulse, and $100 \mu\text{s}$ of state detection. After accumulating data over hundreds of MOT drops, we average the number of 0 and 1 counts in each of the 1500 trials to find the probability for Raman transfer. From the early data (i.e., trial number < 25), we can infer how many atoms were initially loaded. The decay of transfer probability over time provides a gauge of how quickly the atoms leave the trap during the measurement process.

Figure 4.14 shows atom loading data for various combinations of loading beams. The most significant finding is that loading with Raman and $F = 4 \rightarrow F' = 4$ is about a factor of 5 more effective than our previous method of loading with the two lattice beams alone. Loading with Raman and both lattice beams is worse than the loading with the lattice alone; presumably the heating effects of the lattice repump are still present, but now the Raman may slow its efficiency. Loading with Raman and $F = 3 \rightarrow F' = 3$ works about half as well as Raman and $F = 4 \rightarrow F' = 4$, which may be due to the relative power in the two lattice beams. Additionally, the $F = 4 \rightarrow F' = 4$ lattice is not very sensitive to detunings between 10 MHz and 20 MHz, but the loading probability drops significantly when the detuning is reduced to 5 MHz. When the lattice beam is resonant with the cavity, we found that loading almost never occurs, and so we were unable to acquire data at that setting. (In these measurements, both $F = 3 \rightarrow F' = 3$ and $F = 4 \rightarrow F' = 4$ beams are assumed to be 10 MHz blue-detuned unless otherwise specified.) For this data, the single atom transfer probability is about 20%. However, we did not calibrate the data with respect to a single-atom π pulse and so cannot extract information about overall loading efficiencies.

A second set of atom loading data is shown in Figure 4.15. In this case, the yellow curve (no Raman pulse) establishes a background transfer probability of $P_{\text{bgd}} = 0.01$,




Figure 4.14: Probability for transfer of one or more atoms to $F = 4$ following a short ($t < t_\pi$) Raman pulse, as a function of trial number, where each trial lasts about $110 \mu\text{s}$. The measurement is repeated for various loading beam combinations in order to optimize the number of atoms loaded into the FORT.

and the pink curve (conditional loading of only one atom, discussed in Section 4.4) provides a single-atom reference of $P - P_{bgd} = 0.06 - 0.01 = 0.05$. We see that increasing the Raman power from -12 dBm to -2 dBm only improves the loading rate; -2 dBm represents an upper limit on the amount of RF power we can safely use to drive the amplifier for the Raman AOM. Detuning changes to both beams had no appreciable affect. Under the best loading conditions, we find a transfer probability of $P - P_{bgd} = 0.37 - 0.01 = 0.36$, corresponding to an average of $\frac{0.36}{0.05} \approx 7$ atoms loaded per drop. After $165 \mu\text{s}$ of interrogation, about 4 atoms remain.




Figure 4.15: Probability for transfer of one or more atoms to $F = 4$ following a short ($t < t_\pi$) Raman pulse, as a function of trial number, where each trial lasts about $110 \mu\text{s}$. The measurement is repeated for various loading beam intensities and detunings in order to optimize the number of atoms loaded into the FORT. A conditional loading measurement provides a single-atom reference probability $P = 0.05$.

4.3.5 Atom detection with Raman and lattice light

The data in Figure 4.10 suggest that we can use a Raman/probe scheme to identify zero, one, and (less efficiently) two atoms. We have also explored the possibility of introducing $F = 4 \rightarrow F' = 4$ lattice light in conjunction with the Raman and the probe, in order to shift the histogram features to resolve higher atom numbers. To understand this shift, consider a three-level atom as in Figure 4.16 with ground states $|a\rangle$ and $|b\rangle$ and excited state $|e\rangle$, where we can drive transitions between $|a\rangle$ and $|b\rangle$ at rate α and between $|b\rangle$ and $|e\rangle$ at rate β . (In our system, $|a\rangle \rightarrow F = 3$, $|b\rangle \rightarrow F = 4$, $|e\rangle \rightarrow F' = 4$, and α and β are determined by the Rabi frequency of the Raman pair and the intensity of the lattice light, respectively.) We also include decay from the




Figure 4.16: Diagram of three-level atom with transition rates α , β and decay rates γ_a , γ_b . Histogram of photon counts per 1 ms time bin, with 250 ms of continuous Raman, probe, and $F = 4 \rightarrow F' = 4$ lattice beams.

excited state to the two ground states at rates γ_a and γ_b . Then the coupled equations which describe the populations of the three levels are given by

$$\begin{pmatrix} -\alpha & \alpha & \gamma_a \\ \alpha & -\alpha - \beta & \beta + \gamma_b \\ 0 & \beta & -\beta - \gamma_a - \gamma_b \end{pmatrix} \begin{pmatrix} P_a \\ P_b \\ P_e \end{pmatrix} = 0. \quad (4.23)$$

In order to shift the histogram towards higher count rates, we want to minimize the population in $F = 4$, that is, P_b . Solving equation (4.23) for P_b , we find

$$P_b = \frac{1 + \frac{\gamma_a + \gamma_b}{\beta}}{3 + \frac{2(\gamma_a + \gamma_b)}{\beta} + \frac{\gamma_a}{\alpha}}. \quad (4.24)$$

If $\beta > \gamma_a, \gamma_b$ and $\gamma \gg \alpha$, then P_b can become very small. Thus, we should reduce the Rabi frequency of the Raman pair and use a high-intensity lattice beam.

Unfortunately, both of these steps tend to increase the noise of the signal. An intense $F = 4 \rightarrow F' = 4$ beam will heat the atoms, and as we saw in Section 4.3.2, we need a high Rabi frequency to smooth out the atom plateaus. Figure 4.16 shows a histogram from an attempt to use both Raman and lattice beams continuously;

while the $N \geq 1$ peaks have clearly shifted toward higher count rates, the peak widths prevent effective discrimination. After unsuccessful attempts to find optimal intensity settings for both beams, we have concluded that this is probably not a useful technique for atom number determination.

4.4 Conditional feedback to the experiment timing

In the previous sections, we have demonstrated that based on the number of photon counts in a 1 ms window, we can make an accurate determination about the presence of a single atom in the cavity. Here we outline how this information can be used to run the experiment conditioned on real-time atom detection.

The ADwin-Gold system which controls the experiment timing has input and output channels, but until the current application it had only been used to generate output signals (both analog and digital). In the past, we wrote down a sequence of timing commands as a single control program, `timing.in`, using a syntax designed by Dave Boozer. Dave's compiler `hirez.exe` then translated these commands into a text file `foo.txt` written in the ADbasic language. This text was embedded in a larger ADbasic program `pulses.bas` and compiled by the ADwin system. The commands in `timing.in` collected and released the upper and lower MOTs, loaded atoms into the FORT, and then initiated the series of pulses necessary for a particular experiment. Once compiled, a series of commands will loop continuously until a new program is loaded to the ADwin.

To implement conditional loading of the experiment pulses, it was useful to partition the `timing.in` commands into three smaller programs: `timing_load_atom.in`, `timing_one_atom.in`, and `timing_experiment.in`. These programs are then separately compiled to the three text files `foo_load_atom.txt`, `foo_one_atom.txt`, and `foo_experiment.txt` and embedded in `pulses_ten.bas`, which replaces `pulses.bas`.

The command sequence in `timing_load_atom.in` begins with the formation of the upper MOT and ends with loading atoms into the FORT using the Raman and lattice beams. The next step in `pulses_ten.bas` is to initiate a do-until sequence which

calls the `timing_one_atom.in` pulses. These consist simply of 1 ms of continuous Raman and probe fields and a trigger pulse sent to an SR400 Dual-Channel Gated Photon Counter (Stanford Research Systems). After the three pulses, the for-loop then instructs the ADwin to read in the value on the first of its analog channels, which is connected to the analog output from the SR400 and scales linearly with the number of photon counts during the previous ms. Based on histograms such as Figure 4.10, we establish in advance a lower bound of counts corresponding to $N \leq 1$ atom. The timing program repeats the pulse/read combination until the count threshold is exceeded, that is, until there is at most one atom in the cavity. (For redundancy, the do-until sequence is embedded in a for-loop which requires the count rate to exceed the lower threshold over three 1 ms intervals.) Note that the probe and Raman intervals serve two purposes: they measure the atom number, but they also provide a heating mechanism which induces extra atoms to leave the cavity; this usually occurs within about 50 ms.

Once we have eliminated multi-atom events, we introduce the experiment-specific pulse sequence. However, it is possible that all atoms left the trap during the testing interval. For the sake of efficiency, we only want to run the experiment pulse sequence if we think there is an atom present, though no-atom events will also be detected and eliminated in data processing afterwards. We set an upper threshold for counts corresponding to $N \geq 1$ atom; only if the number of photon counts per 1 ms interval exceeds this threshold do we load the final sequence of pulses in `timing_experiment.in`. Otherwise, the experimental cycle is skipped and we begin the process again with `timing_load_atom`.

Figure 4.17 shows an histogram for atom-number determination with data acquired as in Figure 4.10, except that here we have implemented the conditional loading scheme described above. We see that peaks associated with $N \geq 2$ atoms have been eliminated, and only zero- and one-atom peaks remain. There are also a small number of bins with very low count rates, which we have traced to 3 multi-atom events in this data set of 1200 MOT drops. We are unsure as to how these events survived the screening process but are satisfied for the time being with this low contamination




Figure 4.17: Histogram of photon counts per 1 ms time bin, with 280 ms of continuous Raman and probe beams. In contrast to the data presented in Figure 4.10, here we have implemented conditional loading of $N \leq 1$ atom before the experiment begins, with lower and upper count thresholds indicated.

rate.

Having outlined the conditional loading scheme, I now discuss further details of the SR400 photon counter implementation. We continue to send TTL pulses from our two Perkin-Elmer SPCMs to the P7888 photon counting card for computer-based data acquisition. However, we tee off these signals before the card and send them also to the two channels of the SR400. (Before the SR400, each channel passes through a NOT logic gate which serves as a buffer, to prevent potential damage to the P7888 card.) During each 1 ms gate pulse sent from the ADwin, the SR400 sums the counts on its two inputs. At the end of the gate pulse, it sends out an analog signal between 0 and 10 V which is linear with respect to total counts between 0 and 999; that is, 1 count = 10 mV. This signal returns to the ADwin, which provides 16-bit analog-to-digital conversion of inputs between -10 and 10 V, that is, according to the formula

$$\text{value} = 32767.5 \left(1 + \frac{\text{volts}}{10}\right) = 32767.5 \left(1 + \frac{\text{counts}}{1000}\right). \quad (4.25)$$

We have currently set a lower threshold for conditional detection at 160 counts/ms (value = 38010) and an upper threshold at 220 counts/ms (value = 39976).

We have sometimes found it useful to incorporate a second ADwin analog channel input for real-time feedback to the timing. The second channel is needed when we want to generate a sequence of pulse lengths that vary from atom to atom; for example, the data in Figure 4.7 was obtained with Raman pulse lengths between 0.1 and 100 μ s. As discussed in Section 3.5, we generate these pulses by using serial commands to program a DG535 pulse generator, then triggering each pulse sequence with a TTL signal from the ADwin. However, the ADwin timing program also needs to know when the pulses are finished; for pulse lengths which vary from atom to atom, a simple way to do this is to have the DG535 put out a signal for the ADwin at the end of its pulse cycle. After the ADwin sends its initial trigger to the DG535, it then polls its second analog channel input until the value crosses a TTL threshold, at which point it continues with the experiment.

We first put conditional feedback into practice for the Raman optical pumping results presented in Section 4.1 and in [32]. We were pleased to see a number of immediate results: not only did it speed up the process of data acquisition, but also we found that our single-atom Raman scans had reduced background levels, and that the measured Zeeman populations summed to the expected value. (Both of these effects can be understood in the context of multi-atom contamination: when more than one atom is trapped in the cavity, there is a greater probability that an off-resonant Raman pulse will transfer population to $F = 4$, increasing background levels, and we also expect an increased probability to find one atom in any given Zeeman level.) Implementing conditional feedback into our field-nulling protocol (e.g., Figure 2.2) has also improved the contrast of these scans and allowed us to set magnetic field values more reliably. Further applications for this scheme include improved generation of single photons; we would expect a stronger suppression of two-photon events than observed in Ref. [16]. Meanwhile, single photon generation combined with optical pumping might allow us to produce photons with narrower temporal pulse shapes, as the long tail observed in Ref. [16] was attributed to atoms trapped in dark states.