

## Chapter 3

# Experimental Tools and Technical Challenges

My involvement in cavity QED research as a student has encompassed three experimental setups for cold-atom cavity QED with very short cavities. All three stages share a basic experimental scheme and many technical tools and challenges. In this chapter my intent is to present a general summary of the experiments. The details of the atom-cavity microscope experiment are presented thoroughly in Christina Hood's thesis [19], while an even more in-depth discussion of several experimental issues can be found in Quentin Turchette's thesis [27]. Chapter 6 presents in detail the new or altered aspects of the active feedback experiment. Chapter 7 returns to some technical issues, particularly those related to the all-important high-finesse cavities, and expands on some aspects which may be useful for future efforts though not directly essential for understanding current work.

### 3.1 Principal Components of the Experiment

A schematic depicting the heart of the cold-atom cavity QED experiment is shown in Figure 3.1. Cesium atoms are trapped in a magneto-optical trap (MOT) [64, 65], cooled to a few microkelvin using standard sub-Doppler cooling techniques [66, 67, 68], and released to fall through the cavity mode volume. The cavity mode is continuously driven by a probe laser coupled into the cavity at one end; the light transmitted out the other end is detected with an optical balanced heterodyne setup, providing real-




Figure 3.1: Core of experimental setup: cavity, probe, detection, and MOT

time information on the atom-cavity interaction. The detected transmission signal is acquired and digitized for storage, and the signal is also used to condition changes in the probe field for trapping or feedback. Throughout the experiment, stabilization of the lasers and in particular of the physics cavity itself are crucial to the generation of meaningful data in the lab.

## 3.2 Magneto-Optical Trap and Sub-Doppler Polarization Gradient Cooling

The magneto-optical trap [64, 65, 69, 70] has become a standard workhorse of modern optical physics with neutral atoms. Its success hinges on the elegant combination of spatial trapping with a simultaneous cooling mechanism. Counterpropagating light beams of opposite circular polarization are tuned slightly below resonance with an atomic ground to excited state transition. This red-detuning technique by itself provides Doppler cooling [71, 72], so named because the Doppler effect causes an atom to absorb light and thus momentum preferentially from the beam it is moving towards, while spontaneous emission occurs in a random spatial direction. The result is an average light force, proportional to velocity and oppositely directed, which cools the

sample. In three dimensions, energy is removed as moving atoms absorb photons from the red-detuned beams ( $\omega < \omega_a$ ) and re-emit with a frequency spread of  $\gamma_\perp$  centered around  $\omega_a$ . The limiting temperature of this cooling mechanism is the Doppler temperature  $k_B T_D = \hbar\gamma_\perp$  which for Cesium gives  $T_D = 120 \mu K$ .

The spatial restoring force relies on the introduction of a quadrupole magnetic field, produced by an anti-Helmholtz coil configuration in combination with Helmholtz bias coils to zero the magnetic field. The magnetic field along an axis is  $B_z \propto z$ , producing a Zeeman shift in the atomic transition  $m_F \rightarrow m'_F = m_F \pm 1$  which is linear in atomic distance from the trap center. When the frequency of a given transition is decreased, it becomes closer to resonant with the trapping light, and the atom absorbs preferentially from one of the circularly polarized beams. If the circular polarizations and magnetic field gradient are chosen correctly, a displaced atom absorbs more strongly from the beam which will push it back to the center, providing the restoring force for the MOT.

The relevant levels for Cesium are shown in Figure 3.2. Physics takes place on the  $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F = 5$  transition of atomic Cesium, with a closed or ‘‘cycling’’ transition on  $m_F = 4 \rightarrow m'_F = 5$ . The excited state lifetime is 30 ns, corresponding to decay rate  $\gamma_\parallel = 2\pi(5.2 \text{ MHz})$  or  $\gamma_\perp = 2\pi(2.6 \text{ MHz})$  [73, 74].

A final ingredient is the addition of a repumping laser on the  $F = 3 \rightarrow F' = 4$  transition to recycle population which otherwise slowly accumulates in the non-trapping  $F = 3$  ground state due to off-resonant transitions. The repumping laser remains on at a constant strength throughout the MOT loading and the sub-Doppler cooling described below.

After a MOT is loaded (a process of some hundreds of milliseconds to a few seconds), the magnetic field gradient is turned off and the light intensity is turned down to accomplish a short pulse of sub-Doppler cooling on the atoms [66]. Standard sub-Doppler cooling via polarization gradient effects brings Cesium fairly easily to final temperatures of a few to a few tens of microKelvin, depending largely on the degree of heroics associated with magnetic field zeroing [67, 68].

Diagnostics on the final temperature are performed via time-of-flight imaging of




Figure 3.2: Cesium level structure for cooling, trapping, and cavity QED. The  $6S_{1/2}, F=4 \rightarrow 6P_{3/2}, F'=5$  transition is at 852.359 nm.

the expanding cloud after trapping and sub-Doppler cooling, using a brief pulse of the trapping light beams. Magnetic field trim coils are adjusted to ensure that the cloud falls straight down and with as little expansion as possible. A slight difficulty arises because our MOT is constructed as close as possible to the cavity, giving us little time to observe expansion before the cloud hits the mirror substrates. In practice we perform final field zeroing by raising the MOT position a few millimeters with a vertical bias coil; we then have about 30 ms of accessible fall time before the atoms leave the beam profile and hit the cavity substrates. Cooling is optimized with the trap raised in this way and we assume fields will be nearly zeroed for the lower trap position in the actual experiment.

### 3.3 Probe Beam and Local Oscillator Generation and Stabilization

The cavity QED probe beam in these experiments has a very modest requirement for optical power. Since we typically work at intracavity photon numbers of one or

smaller, the mode-matched optical power delivered to the cavity is on the order of  $(\hbar\omega)\kappa \sim 10$  pW or less. This means in practice that the probe beam is maintained at some sensible working strength on the order of a milliWatt and then severely attenuated with a stack of neutral-density filters just before being sent to the cavity. The probe laser also serves as the source for a strong local oscillator (LO) beam used in the heterodyne detection of the cavity-transmitted light, as described later in this chapter. We typically use an LO strength of 5 mW.

The probe beam in the atom-cavity microscope and previous experiments was generated by a Ti:Sapph laser pumped with 5 Watts of green light (532 nm) from a Spectra Physics Millennia V solid-state diode-pumped unit. With no water cooling of the Ti:Sapph crystal, single-mode output power at 852 nm was typically 300-350 mW. This light was split to: (1) a Cs FM saturated absorption spectroscopy [75] path to generate an error signal for locking the laser to Cesium, (2) a local oscillator path intensity stabilized using a ThorLabs liquid-crystal-based intensity stabilizer, and (3) the probe beam path including frequency shifting, intensity stabilization, rapidly switchable intensity adjustment, and mode-matching. The final probe beam in this setup consisted of a carrier  $\sim 800$  MHz blue-detuned from  $(\omega_a, \omega_c)$  which received sidebands at  $\sim \pm 800$  MHz through a traveling-wave modulator so that the lower sideband acted as the near-resonant probe. The carrier and upper sideband were ignored as too far detuned to interact significantly with the atoms. In practice, care was required to ensure negligible effect from the carrier, since it was far-detuned but typically a factor of a hundred stronger than the probe itself. While a given probe strength could be achieved by many combinations of overall optical power and RF modulation index, the optical power had to be set low enough and RF power high enough to keep the carrier from noticeably altering the atom-cavity system.

In the active-feedback experiment, the probe beam originates from a grating-stabilized diode laser at 852 nm. The shifting, switching, and stabilization of this beam are described in detail in Chapter 6.

### 3.4 The Physics Cavity

The heart of the cold-atoms cavity QED experiment is the physics cavity itself. The atom-cavity microscope and active-feedback experiments use cavities of length  $l \approx 10 \mu\text{m}$  and finesse  $F \approx 480,000$ . The mirrors have radius of curvature  $R = 10 \text{ cm}$ , leading to cavity field waists  $w_0 \approx 14 \mu\text{m}$  for the  $TEM_{00}$  mode. Mirror properties and cavity mounting are described in detail in [19] and in Chapters 6 and 7; here I outline the basic construction and some important technical details related to length stability and birefringence.

The cavity is mounted on a solid OFHC copper block, which sits on a vibration-isolation stack within the vacuum chamber. Each (cylindrical) mirror substrate sits in a v-block to which it is attached by glue; the v-blocks in turn are glued atop shear mode PZTs which are glued to the copper mount piece. Cavity length is tuned and stabilized via voltage applied to the shearing PZTs. The strategy of these experiments is to use the cavity's transmission of a fixed-frequency probe beam to infer a measurement of atom-cavity interactions; thus stabilization of the cavity length (resonance frequency) is essential. In fact, cavity length excursions are the principal source of technical noise in the measurements. In particular, mechanical resonances of the mounting stack must if possible be engineered out of the frequency range where meaningful atomic motional signals are expected. This issue, which I addressed via overall mount/PZT selection and rigidity of attachment, is dealt with more thoroughly in [19] and in Chapter 6.

A second nagging concern in cavity design is birefringence, treated more fully in Chapters 6 and 7. Experience suggests that glue applied to the mirror substrates causes stress-induced birefringence in the mirror coatings, so that the cavity alters all input polarizations save the two linear cavity eigenpolarizations. This is inconvenient for experiments in which we wish to work with circularly polarized light to maintain the cycling transition  $m_F = 4 \rightarrow m'_F = 5$ . In practice we attempt to minimize cavity birefringence by applying the minimum amount of glue necessary, at the farthest possible distance from the mirror surfaces. We then use circularly polarized probe

light, but a small fraction of the opposite circular polarization in fact builds up in the cavity. This situation allows a slow drain of population from the interacting ( $F = 4$ ) ground state, motivating the use of a repumping beam on  $F = 3 \rightarrow F' = 4$ . This beam comes from the same laser as the MOT repumping light, but is overlapped with the cavity QED probe for injection into the cavity.

Finally, this discussion of cavity construction and geometry is a good place to include a note on delivery of single atoms to the cavity mode. Given a MOT with  $10^4$  or more atoms cooled and dropped directly onto the cavity, how do we obtain a situation where atoms traverse the cavity mode one at a time? Figure 3.3 provides a sense of how this occurs. The figure shows a to-scale rendering of the cavity mirror surfaces and  $TEM_{00}$  mode profile in the center of the mirrors. The cavity geometry itself clearly cuts off a large fraction of the atom flux to the interaction region, by imposing a very strict limit on atomic velocity along the cavity axis. Indeed, for an atom to enter the gap at the top of the mirror substrates and reach the cavity mode without hitting a mirror, the axial velocity must be less than about  $4 \text{ mm/s}$  corresponding to an axial temperature of  $\sim 0.1 \mu\text{K}$ . Thus single atoms are easy to come by in the experiment.

### 3.5 Vacuum Chamber

The experiments are conducted under UHV conditions in a chamber where steady-state pressure is maintained by ion pumping with Varian Starcell pumps. Pressure within the chamber is measured using Bayard-Alpert ionization gauges. The atoms are supplied from a Cesium reservoir initially loaded with a one-gram ampule of atomic Cesium. The entire chamber (in the atom-cavity microscope) or just the upper chamber (in the active-feedback experiment) is operated as a Cesium vapor cell replenished periodically by opening a valve to the reservoir and/or heating it gently. The pressure in the chamber under these conditions is a few times  $10^{-8}$  torr. Pressure in the actual cavity region is difficult to accurately measure, but it could plausibly be significantly higher than the overall chamber pressure due to (1)

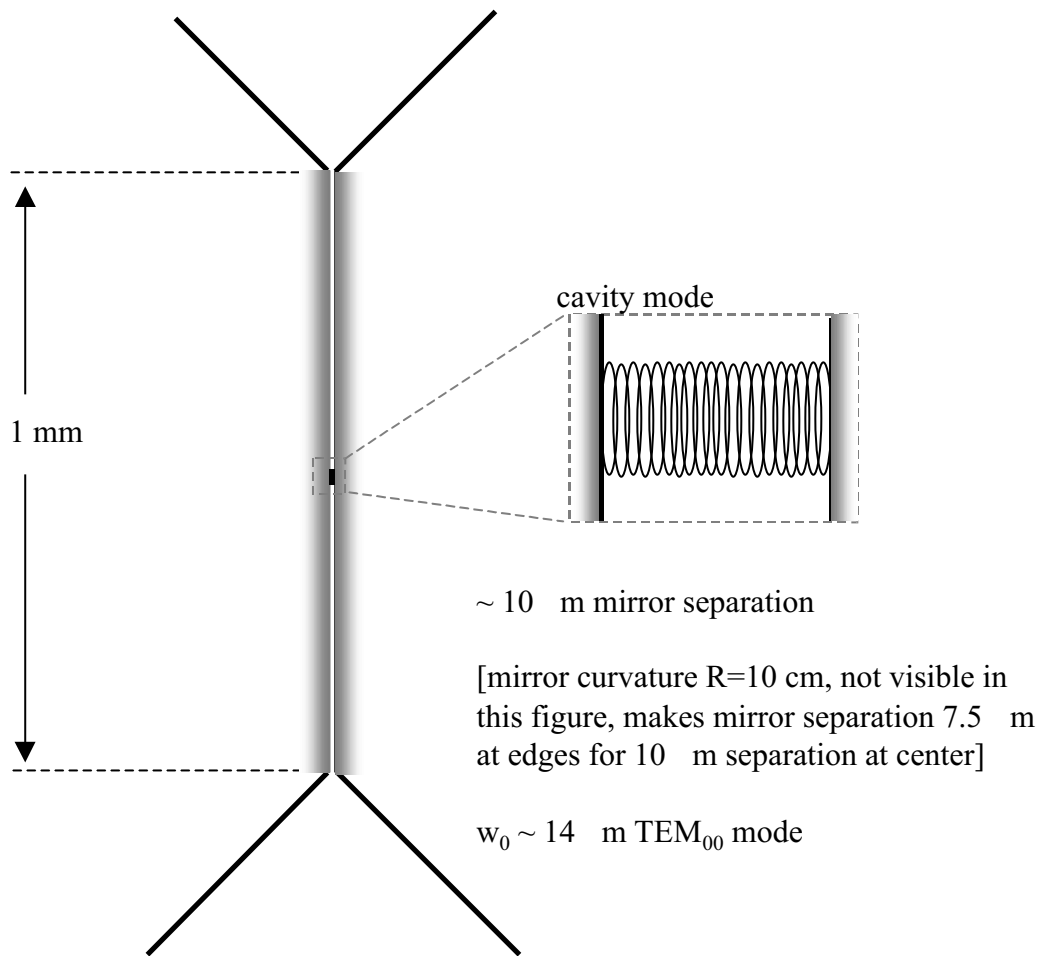


Figure 3.3: Cavity mirrors, spacing, and field mode depicted to scale.



low conductance given the geometry shown in Figure 3.3 and (2) concentration of marginal UHV materials close to the cavity because they are involved in the cavity mount itself. The most realistic prospect for an intracavity pressure measurement is trapped atom lifetime itself. This prospect is not sensible for our single-photon trapping experiments, since lifetimes are  $\lesssim 1$  ms and even a background pressure of  $10^{-8}$  torr allows a collisional lifetime of about 0.5 s. However, experiments using a far-off-resonant trap in a similar setting now reach lifetimes of  $\gtrsim 2$  s and demand more careful consideration of background limits [76]. Better understanding of other loss mechanisms in these experiments could yield some useful limits on the intracavity vacuum pressure in the near future.

As mentioned, background collisions do not limit trapped-atom lifetimes in our current single-photon trapping experiments. From this point of view, better vacuum only becomes an issue as we consider the future implementation of feedback, cooling, or auxiliary trapping schemes that will drastically enhance atom-cavity dwell times. However, one other advantage also accrues from an improved vacuum at the cavity position. This is the possible increase in cavity longevity by slow-down of the process whereby Cesium vapor atoms coat the mirror surfaces. Previous experiments have seen degradation of cavity finesse after approximately 1.5 to 2 years in an experiment; while the coating process is not known to be responsible, two pieces of circumstantial evidence suggest it is worth guarding against. First, after vacuum was broken on these past experiments and the mirrors were examined, “stars” of oxidized Cesium were prominent on the mirror surfaces (Figure 3.4). Second, the cavity of [51, 52], which is sheltered from direct contact with the Cs vapor cell region in that apparatus, has already seen an unusually long period of nearly four years of experimental usefulness. The combination of all these considerations motivated the transition to a differentially pumped chamber for the active-feedback experiment, as discussed in Chapter 6.




Figure 3.4: Oxidized Cesium on the surface of a mirror removed from a well-used physics cavity. A clean mirror surface is shown for comparison. Each mirror surface has a diameter of 1 mm.

### 3.6 Locking of Laser Frequency and Cavity Length

The probe beam frequency  $\omega_p$  and cavity resonance  $\omega_c$  must be stabilized to Cesium for fixed detuning from  $\omega_a$ , but need not be locked to much better than 1 MHz given the physics cavity linewidths of more than 10 MHz; this large cavity linewidth, despite the high finesse, is a result of very short cavity length. The probe frequency is referenced to Cesium via FM saturated absorption spectroscopy [77, 78, 75].

To obtain a rough length-stability requirement for the cavity, we note that while the full atom-cavity coupling  $g_0$  shifts the system resonance by  $\sim 10\kappa$ , we want to detect the effects of small fractional changes in  $g(\vec{r})$  which shift the system resonance by small fractions of  $\kappa$ . If for a cavity of width  $\kappa/2\pi \sim 14$  MHz we require  $\omega_c/2\pi$  to be stabilized to about 1 MHz, this translates to a length stability on the order of  $10 \text{ fm} = 10^{-14} \text{ m}$ . While this stability requirement on the order of atomic nuclear dimensions may initially seem nonsensical, this is of course the position averaged over the spatial cavity mode profile and over the fastest dynamical and measurement timescales of the system.

The realization of laser and cavity locks in the atom-cavity microscope experiment is treated in [19]; the Ti:Sapph frequency is locked to Cesium and the physics cavity

length is locked to the probe transmission itself, with lock blanking during an atomic transit. The rather different methods of the active-feedback experiment are presented in Chapter 6.

### 3.7 Heterodyne Detection and Calibration

Experiments with intracavity photon numbers of roughly 0.05 to 1 imply a cavity transmission signal on the order of 1 to 10 piconWatts at 852 nm. Signal powers in this range suggest the use either of photon counting or of optical homodyne/heterodyne methods for detection with reasonable signal-to-noise. While photon counting provides a measurement of  $\langle a^\dagger a \rangle$  for the output field, the balanced heterodyne detection we employ produces a photocurrent proportional to  $\langle a \rangle$ . It can be used to measure both amplitude and phase of the field [47], but in our experiments we measure only  $|\langle a \rangle|^2$  as discussed in the next section. For a coherent field  $\langle a^\dagger a \rangle = |\langle a \rangle|^2$ , but for our experimental parameters there is a small ( $\approx 20\%$ ) difference between these quantities when an atom is strongly coupled to the cavity mode [19]. This distinction is much more pronounced when detecting either  $|\langle a \rangle|^2$  or  $\langle a^\dagger a \rangle$  with a probe close to the cavity resonance frequency, rather than near the lower vacuum Rabi sideband as in our triggered-trapping experiments.

Figure 3.5 illustrates the principle of balanced heterodyne detection for a weak signal beam. The most basic idea is to encode the signal field as an RF beatnote on a much stronger local oscillator field. The signal and LO are phase-coherent, coming from the same laser, but are deliberately shifted away from one another by an RF frequency shift, in our case in the range  $\approx 10 - 100$  MHz. The signal field  $E_s$  and LO field  $E_{LO}$  are spatially overlapped and mixed on a 50/50 beamsplitter, yielding fields  $E_{LO} + E_s$  and  $E_{LO} - E_s$  at the two beamsplitter output ports. These fields are incident on balanced photodetectors whose photocurrents are then subtracted, yielding a difference photocurrent proportional to the product of the two fields:

$$i_- = i_1 - i_2 \propto (E_{LO} + E_s)^2 - (E_{LO} - E_s)^2 \propto E_{LO}E_s. \quad (3.1)$$

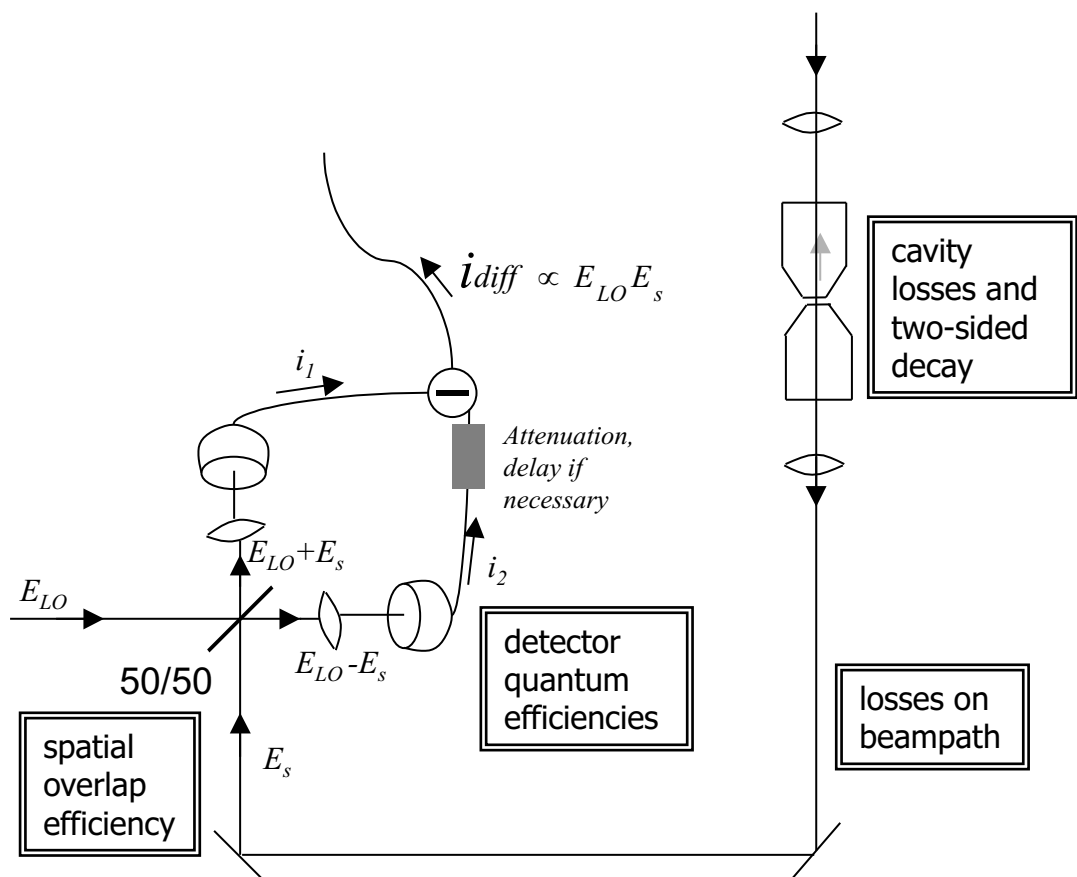


Figure 3.5: Heterodyne operation and factors contributing to detection efficiency.

If we further break down the LO and signal beams into their amplitude and noise,  $E_{LO} = A_{LO} + N_{LO}$  and  $E_s = A_s + N_s$ , we have

$$i_- \propto A_{LO}A_s + A_{LO}N_s + N_{LO}A_s + N_{LO}N_s. \quad (3.2)$$

The terms above are written in decreasing order of importance in the desired configuration of a balanced heterodyne detector; the relegation of  $N_{LO}$  to the obscurity of the last two terms allows us to write  $i_- \propto A_{LO}(A_s + N_s) \approx (Constant)E_s$ . This is the vaunted LO-noise immunity of the balanced heterodyne detection method.

This noise immunity operates on (at least) two distinct levels. First, the balanced configuration and subtraction of  $i_1$  and  $i_2$  eliminates photocurrent terms containing only the large LO field, so that the difference photocurrent of Equation 3.1 is strictly proportional to the signal field  $E_s$ . As long as  $A_{LO}/N_{LO} \gg 1$ , we can write  $i_- \sim (Constant)E_s$ . Second, however, it is possible to ask for more. We may hope for a situation in which the noise characteristics of the photocurrent truly reflect, to leading order, only the signal noise  $N_s$  and not the LO noise  $N_{LO}$ . In this case, referring to Equation 3.2, we require more specifically that  $N_{LO}/A_{LO} \ll N_s/A_s$ . The inequality will certainly hold for shot-noise dominated weak signal and strong LO but may be more questionable if technical noise dominates. In practice the separation in size between the four terms of Equation 3.2 is not necessarily huge and the noise immunity must be carefully evaluated in operation.

In an experimental system one contends in general with imperfect balancing, both of optical powers (beamsplitter not exactly 50/50) and of detector efficiencies and gains. In optimizing the overall performance we proceed by (1) balancing the LO and signal powers on both detectors as nearly as possible and (2) attenuating the signal from one detector if necessary to achieve balancing of the resulting photocurrents at the subtractor. ‘‘Unnecessary’’ attenuation certainly carries the possibility of contributing excess electronic noise if that is a primary concern. Another alternative would be to deliberately unbalance the optical powers to match the detector imbalance, i.e., adjust the beamsplitter transmission directly to cancel the  $E_{LO}^2$  term in

the difference photocurrent. This would result in a nonzero  $E_s^2$  term, but insofar as  $E_s \ll E_{LO}$  this term is unimportant anyway. The one-step overall balancing method is actually a bit more difficult to carry out with certainty that one has balanced the right thing. In terms of noise immunity, though, the relative merits of step-by-step vs. overall balancing must be weighed given the details of each case. A treatment of diverse concerns in non-ideal homodyne detection can be found in [79].

Our heterodyne detection and its calibration are presented in detail in [19], and have not changed significantly from the atom-cavity microscope to the active-feedback setup. However, I revisit here the issue of overall efficiency for detection of an intracavity photon. In that experiment we measured a heterodyne efficiency (due to detector quantum efficiency and spatial mode overlap) of 48%. The cavity mirrors each had transmission of 4.5 ppm and absorption/scatter losses of 2 ppm, while we measured just the field transmitted through one mirror. Finally, propagation losses from cavity output to heterodyne were estimated at 25%; this rather large figure comes mostly from a beamsplitter which sent 20% of the light to a photomultiplier tube for locking the cavity while optimizing the heterodyne alignment. Taking into account all these contributions, we find an overall efficiency of  $\eta = (0.48) \cdot \frac{4.5}{2(4.5+2)} \cdot (0.75) = 0.125$ , or 12.5%. This is the efficiency for detecting the *total* cavity decay, which is a signal of size  $(2\kappa)|\langle a \rangle|^2$ .

However, if we neglect the two-sided nature of the cavity and quote the total signal as  $(\kappa)|\langle a \rangle|^2$ , the corresponding efficiency will be  $\eta' = 2\eta = 0.25$  or 25%. This is the origin of the 25% efficiency quoted in [25, 19] and used in simulations of the experiment. Since the detection efficiency is relevant in determining signal-to-noise for sensing of intracavity dynamics, either set of definitions can be chosen in a given calculation. However, the mutual consistency of the definitions must be carefully checked in each situation.

I conclude the whirlwind tour of heterodyne detection with two related points. First, the strength of the local oscillator beam is determined by several considerations. Clearly the LO should be very much stronger than the signal, but there is little danger of violating this criterion in our experiment. LO power must also be high enough to

ensure that overall signal-to-noise is limited by signal shot noise ( $\propto A_{LO}N_s$ ) rather than electronic noise unrelated to the light. An important upper limit in practice is the saturation of the photodetectors, not only in terms of DC power but in the (more relevant) RF band of the actual heterodyne beatnote. The LO-signal beatnote must not saturate the detectors either for the actual experimental signal (10 pW) *or* for the much larger (10 nW) signals needed to calibrate the heterodyne-to-photon-number relationship. Another cautionary note about large local oscillator power is that a tiny fraction of the LO leaking back along the signal path to the physics cavity can wreak havoc on the atom-cavity system itself. To deal with this issue for the LO power we use, the focusing lenses and detectors in both arms of the heterodyne are tilted to avoid back-reflection of the LO.

### 3.8 Data Acquisition

The heterodyne difference photocurrent is amplified and fed into a commercial Hewlett-Packard (Agilent) spectrum analyzer for detection of the component at the LO/signal RF beat frequency [80]. An initial detection bandwidth is set via the resolution and video bandwidths of the spectrum analyzer, both typically set to 100 kHz. The video output of the spectrum analyzer then measures the difference photocurrent *power* at the heterodyne beat frequency, making it proportional in the end to  $|\langle a \rangle|^2$  since the photocurrent amplitude is proportional to  $|\langle a \rangle|$ . This signal constitutes the experimental data which is digitized, processed, saved, and used for triggering/feedback via any of several different schemes. The data acquisition and triggering protocols for the atom-cavity microscope are presented in [19], while those for the active-feedback experiment are given in Chapter 6.