Chapter 5

Active Feedback on Atomic Position: Towards a Quantum Servo

The atom-cavity microscope, as described in the preceding two chapters, demonstrates the ability to trap a single atom in an optical cavity and monitor its position in real time. Current experimental efforts in our group and elsewhere (see [87]) focus on the use of this real-time position information to feed back to actively control atomic positions within the cavity mode. Progress in this direction aims to use cavity QED to investigate fundamental questions of quantum measurement and optimal state estimation and control (see, e.g., [88, 89, 90]). The ability to trap an atom with a single photon field is a dramatic illustration of single-quantum physics in optical cavities; however, it is the sensing power and signal-to-noise for real-time observation of the system that makes the optical cavity unique. This chapter discusses feedback strategies and simulations of their performance for experimentally relevant parameters.

The cavity enhancement of sensing was illustrated in dramatic fashion by the two-dimensional trajectory reconstructions of the previous chapter. However, the initial goal of our feedback algorithms will be to control $\rho(t)$; for this purpose it is sensible to ignore $\theta(t)$ for the time being and apply all available signal-to-noise to the task of estimating $\rho(t)$ and $\dot{\rho}(t)$ in real time. The goal of such a program is then to use this information to drive $\rho(t)$ to a constant value, or in other words to circularize an orbit

in the (ρ, θ) plane while not necessarily driving it to the cavity axis ($\rho = 0$). The latter task, which requires an explicit method of breaking cylindrical symmetry for position sensing and for the effective potential, can be considered as a later extension.

5.1 The Atom and Cavity as a Control System: Basic Feedback Strategy

As a guide in the identification of plausible feedback strategies and their limitations, it is useful to restate the problem somewhat in the language of control systems. To this end, we begin by setting aside the issue of axial motion and treating the atom as a particle in a cylindrically symmetric, approximately Gaussian two-dimensional potential whose depth is controlled by the input light intensity:

$$U \approx -U_0 e^{-\rho^2/w^2}. \tag{5.1}$$

Note that the potential waist w is not simply equated with the previously introduced cavity field waist w_0 or with the mode intensity waist $w_0/\sqrt{2}$, but rather is set by the self-consistent interaction of atom and light field in the strong coupling regime. Where the cavity mode profile is exactly Gaussian, U is only approximately Gaussian and has an exact form that is nonanalytical as determined by steady-state solutions to the master equation for an atom at each value of ρ. The potential depth U_0 depends on the intensity of the optical field used to drive the cavity mode. The potential waist w is in fact a (slowly varying) function of the drive strength as well.

The trapped atom is also subject to dynamical noise in the form of friction and momentum diffusion arising from decays and re-excitations of the system on timescales faster than the motion. In the regime of Ref. [25], the contribution of friction is small compared to the momentum diffusion terms in the equation of motion.

Because the two-dimensional potential is cylindrically symmetric, the atom's angular momentum L is constant, or rather, varies only due to dynamical noise. We

can thus write a one-dimensional effective potential in the ρ dimension,

$$V_{eff} = -U_0 e^{-\rho^2/w^2} + L^2/2m\rho^2 \tag{5.2}$$

and thus an equation of motion (for an atom of mass m)

$$\ddot{\rho} = -\frac{2\rho U_0}{mw^2} e^{-\rho^2/w^2} + \frac{L^2}{m^2\rho^3} \tag{5.3}$$

which we notationally simplify to the form

$$\ddot{\tilde{\rho}} = -\tilde{\rho} E e^{-\tilde{\rho}^2} + \frac{\tilde{L}}{\tilde{\rho}^3} \tag{5.4}$$

where $\tilde{\rho}$ is dimensionless ($\tilde{\rho} = \rho/w$), $E = 2U_0/mw^2$ is the input we control by varying the driving field strength, and $\tilde{L} = L^2/m^2w^4$ is constant except for the influence of dynamical noise.

The measurement of light transmitted through the cavity, $T(t)$, is equivalent to a (noisy) measurement of $\rho(t)$. The noise of this measurement is largely fundamental quantum noise (shot noise) of detection. The mapping between T and ρ, derived again from steady-state solutions of the master equation for the coupled atom-cavity system, is not linear and furthermore depends on the value of the driving strength E. The initial objective is to circularize the two-dimensional orbit – in other words, to make ρ constant or to hold $\dot{\rho} = 0$ by varying the control input E.

The simplified system can be described by a block diagram as shown in Figure 5.1. The system exhibits myriad nonlinearities; for example, $(T \to \rho)$ is nonlinear and depends on E, the dynamical noise depends on E, and the equation of motion for ρ is itself nonlinear. Nonetheless, while this statement of the problem does not suggest provably optimal feedback strategies, it does motivate some conceptually simple algorithms based on switching between discrete values of driving strength E. Switching strategies of this sort are often invoked for the sake of robustness, a major consideration in this scenario; robustness to dynamical and measurement noise is certainly important, but perhaps even more relevant is robustness to small uncertainties in sys-

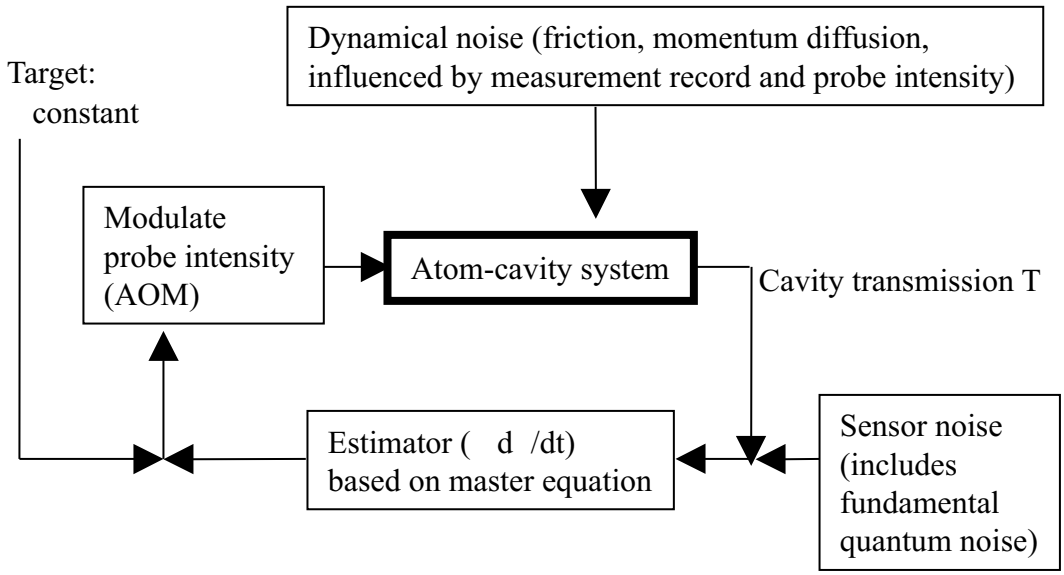

Figure 5.1: Block diagram for the atomic position feedback loop, illustrating sources of noise and system nonlinearities.

tem parameters (e.g., driving strengths or detunings). Switching or "bang-bang" type algorithms have the additional virtue of admitting easy implementation in simulations and in experimental design.

5.1.1 General Feedback Strategy

The feedback algorithms we investigate in this chapter all share the same basic strategy of switching the driving field intensity between two discrete levels. This corresponds to switching between two potential depths (and, incidentally, two different sets of friction and momentum diffusion coefficients as well). The simple objective is to time this switching relative to the atomic motion so that an atom sees a steep potential when climbing out of the trap ($\dot{\rho} > 0$) and a shallow potential when falling back towards the trap center ($\dot{\rho} < 0$), as illustrated in Figure 5.2. The feedback algorithm, then, is based on switching the potential at turning points of ρ, i.e., each time $\dot{\rho}$ crosses zero. Implemented effectively, this approach promises significant dynamical cooling of the radial motion ($\rho \rightarrow constant$ or $\dot{\rho} \rightarrow 0$) in just a few oscillations.

The initial detection and trapping of an atom are accomplished as in Chapter 4 and [25, 26]; a weak probe at driving level *exlo* is used to detect the atom's arrival in

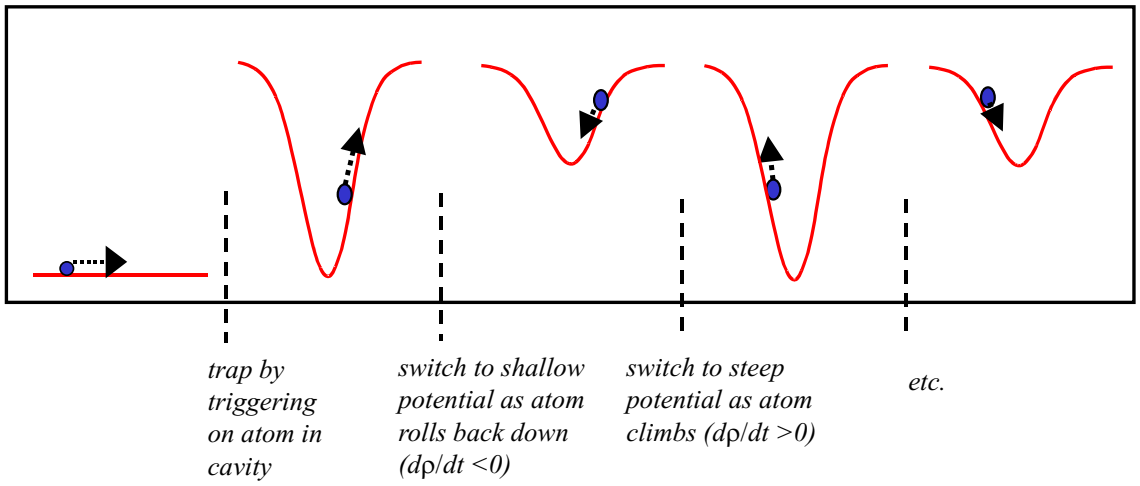

Figure 5.2: General feedback strategy for atomic radial coordinate

the cavity with minimal effect on the motion, and an increase in transmission of this beam triggers a switch to driving level *hi* to populate strong-field-seeking states and trap the atom. Feedback is then implemented by switching the trapping potential between the *hi* level and an intermediate *lo* setting, with switching times based on real-time information about the motion of the single atom.

The simplest algorithm would be to switch back and forth between *hi* and *lo* potentials at the turning points of $\rho(t)$, which are the zero-crossings of $\dot{\rho}(t)$. That is, trap initially in *hi*, switch to *lo* when $\dot{\rho}$ crosses zero from above (i.e., when ρ begins to decrease), switch back to *hi* when $\dot{\rho}$ crosses zero from below (ρ reaching its minimum and increasing), and so on until the atom escapes. However, this strategy calls for a theoretically infinite sequence of switching events, while it is desirable to instead achieve a steady state in some large-time limit. The presence of dynamical noise implies that the exact steady state of $\dot{\rho} \to 0$ is in any case unreachable, so we replace it with a goal of confining $\dot{\rho}$ to some range $[-lim, +lim]$. Thus the feedback strategy is modified to include slight hysteresis: $hi \to lo$ when $\dot{\rho} \to -lim$ from above, $lo \to hi$ when $\dot{\rho} \to +lim$ from below. With this modification, switching stops once $\dot{\rho}$ is confined within the acceptable range. Furthermore, we prefer a steady state with *hi* potential because of the deeper confinement. To bias the system towards this final state, we use asymmetric hysteresis limits: $lo \to hi$ when $\dot{\rho} \to +lim$ from below, $hi \to lo$

when $\dot{\rho} \to -(lim + \delta)$ from above.

5.2 Simulations of Feedback Algorithms in Operation

For our simulations, we have chosen cavity parameters as well as driving strengths and detunings based on Ref. [25], with $\Delta_{ca} = -47$ MHz, $\Delta_{pa} = -125$ MHz, $n_{exlo} = 0.05$ photons in the empty cavity, $n_{hi} = 0.3$ photons in the empty cavity, and $n_{lo} = 0.15$ empty-cavity photons. These driving strengths are chosen to be high enough so that an atom of typical kinetic energy can be trapped by both *lo* and *hi* drives, yet low enough to ensure the increase in momentum diffusion between *lo* and *hi* does not outstrip the increase in potential depth.

Our simulations of the atom-cavity dynamics are based on the treatment outlined in Section 4.4.1 and described in detail in Ref. [26, 83]; the treatment is fully quantized for the atomic internal state and the cavity light field, but considers the atomic center-of-mass motion quasi-classically. This approximation is suitable for the current experimental situation, with more manifestly quantized motion to be accessed by better cooling and/or detection of the atom's axial motion. The dynamical noise of the system, in the form of friction and momentum diffusion, is included in the simulation; the resulting "heterodyne transmission" trace is a perfect record of $|\langle a \rangle|^2$, on which measurement bandwidths and shot noise must be imposed separately. Shot noise is modelled as Gaussian white noise with an amplitude that depends on the size of the (noiseless) transmission signal.

In the presence of sensor noise, we require estimators for the quantities ρ and $\dot{\rho}$. Because one parameter (L) in the equation of motion is unknown and in fact slowly varying, we have chosen not to implement estimators based on a Kalman-filter approach [91]. More sophisticated treatments include Kalman-type approaches to simultaneously estimate ρ, $\dot{\rho}$, and L [92], but these have not been explored in full detail. Possible approaches include various methods such as the extended Kalman filter

or unscented Kalman filter; for an introduction and overview, see [93]. Meanwhile, we choose to estimate $\rho(t)$ and $\dot{\rho}(t)$ directly from the measurement record, with no explicit reference to the equation of motion for the system.

The noisy transmission signal T is sampled at 1 MHz as mentioned above. Digitization of the signal at finite resolution can introduce noise, so the resolution of this step should be kept small relative to pre-existing noise amplitudes. To estimate ρ, we first perform an RC low-pass filter on T at 100 kHz. This step is an infinite impulse response (IIR) filter which consequently introduces only a small delay in the estimator. This filtered transmission signal is then put through a lookup table with linear interpolation to obtain $\rho_{est}(t)$.

The resulting $\rho_{est}(t)$ tracks the actual ρ closely but still with significant noise. Obtaining a time derivative without excessive noise thus requires some care. A variety of methods are mentioned in Ref. [94], in which the authors are concerned with estimating the sign of a time derivative in order to feed back to a system – essentially the same problem we encounter. We employ a simple finite impulse response (FIR) filter that takes the slope of a linear least squares fit to $\rho_{est}(t)$ over a window of fixed size. A detailed implementation of this filter is found in Ref. [95]. The resulting $\dot{\rho}_{est}(t)$ is a good estimator for $\dot{\rho}$ at the middle of the window, so the delay induced is approximately half the window size. We find that a window size of 30 to 40 μs gives a signal $\dot{\rho}_{est}(t)$ which is quiet enough for use in our control. Thus reliable estimation in the presence of noise introduces a delay of approximately 15 to 20 μs in the feedback loop. This time delay can be compared to a typical atomic orbital period of $\tau_r \sim 100$ μs, corresponding to a period of ~ 50 μs for ρ. Feeding back effectively in the presence of such large delays requires a certain amount of adjustment to the naive cooling algorithm of Section 5.1.1, as discussed in detail below.

5.2.1 Actual Dynamics But No Measurement Noise

Before treating the case of actual experimental noise, we explore the performance of our feedback strategy in simulations with noiseless measurement and thus per-

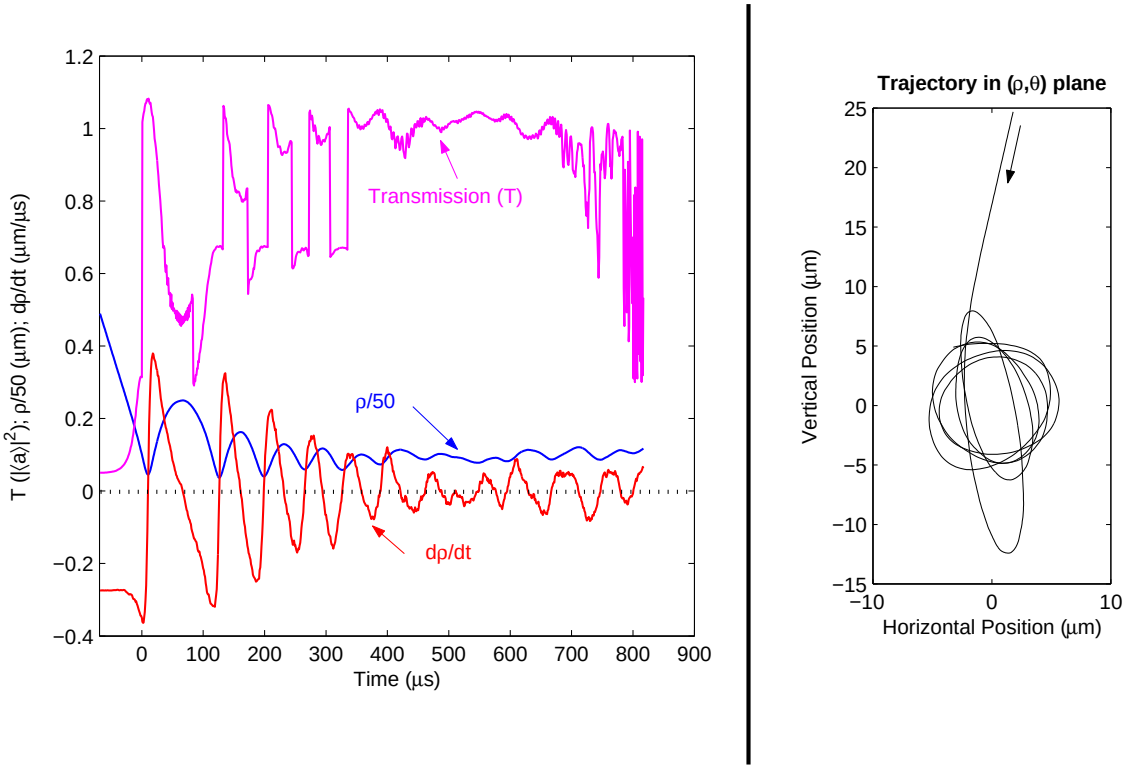

Figure 5.3: Simulated feedback example with perfect sensing of $\rho(t)$.

fect, zero-delay sensing of $(\rho, \dot{\rho})$. Figure 5.3 shows an example trajectory using this asymmetric-hysteresis switching strategy. The values of cavity transmission T and atomic position variables are sampled every 1 μs, but the dynamical timestep is 3,000 times finer than this "information" timestep. Note that axial motion (the x direction our strategy neglects) is included in the simulation, and when its amplitude is large it gives rise to the very fast variations seen in $T(t)$. However, since the period of x motion is similar to the information timestep used, note that these signals are undersampled in the record.

A 10 μs box filter is applied to $\dot{\rho}$ in order to remove some oscillations caused by x motion and also partially to anticipate some effects of noisy detection and delay. In setting the conditions for potential-switching, we employ the asymmetric hysteresis described above, so the potential depth is switched at $\dot{\rho} = +0.05, -0.08$ μm/μs. Switching events, since they correspond to turning the light level up and down, can be seen as sharp edges in the transmitted light T. As the example illustrates, the control

strategy successfully circularizes atomic trajectories within a few orbital cycles. This can be seen from $\rho(t)$ as well as from the trajectory shown in the (ρ, θ) plane. The hysteresis limits are chosen so that variations in ρ due only to dynamical noise tend not to trigger any switching of the drive. This is illustrated by the continued high control level throughout the time $t = 450 - 600\ \mu$s, while ρ is wandering diffusively rather than oscillating with regularity.

The overall trap lifetime is dominated, as in this example, by heating in the x direction. (In the example shown, note the fast, large-amplitude variation in transmission just before the atom escapes; this is a signature of rapid axial heating.) Thus our feedback strategy has little impact (at the level of 10%) on average trapping lifetimes. Circularizing the orbit helps decrease axial heating since the potential depth no longer wanders as ρ varies; however, the feedback is accomplished by sharp switching events which occur at arbitrary times relative to the oscillations in the x direction. The overall impact on lifetimes is therefore small in the simulations we have performed. Since the feedback algorithm is aimed at reducing motion in the ρ direction, its success is best measured by its performance at that task specifically. Lifetime effects can become apparent only if the axial motion is suppressed by some other means; that case is treated briefly in the final set of simulations presented in this chapter.

5.2.2 Adding Measurement Noise Adds Delays

The addition of measurement noise and consequent estimation of $\dot{\rho}$ introduce significant loop delays, as described above. Since $\dot{\rho}_{est}$ can be almost half a cycle behind the actual $\dot{\rho}$, we expect naive switching to be well out of phase with the atomic motion and thus relatively ineffective as a cooling mechanism. Figure 5.4 shows an example in which the feedback strategy is identical to that of Figure 5.3, but applied to $\dot{\rho}_{est}$ rather than to $\dot{\rho}$ itself. The resulting time delay seriously compromises performance, as shown. In the figure, note that the switching events, recognizable as sharp edges in transmission, do not line up with turning points of ρ. As a result, $\dot{\rho}$ is not damped

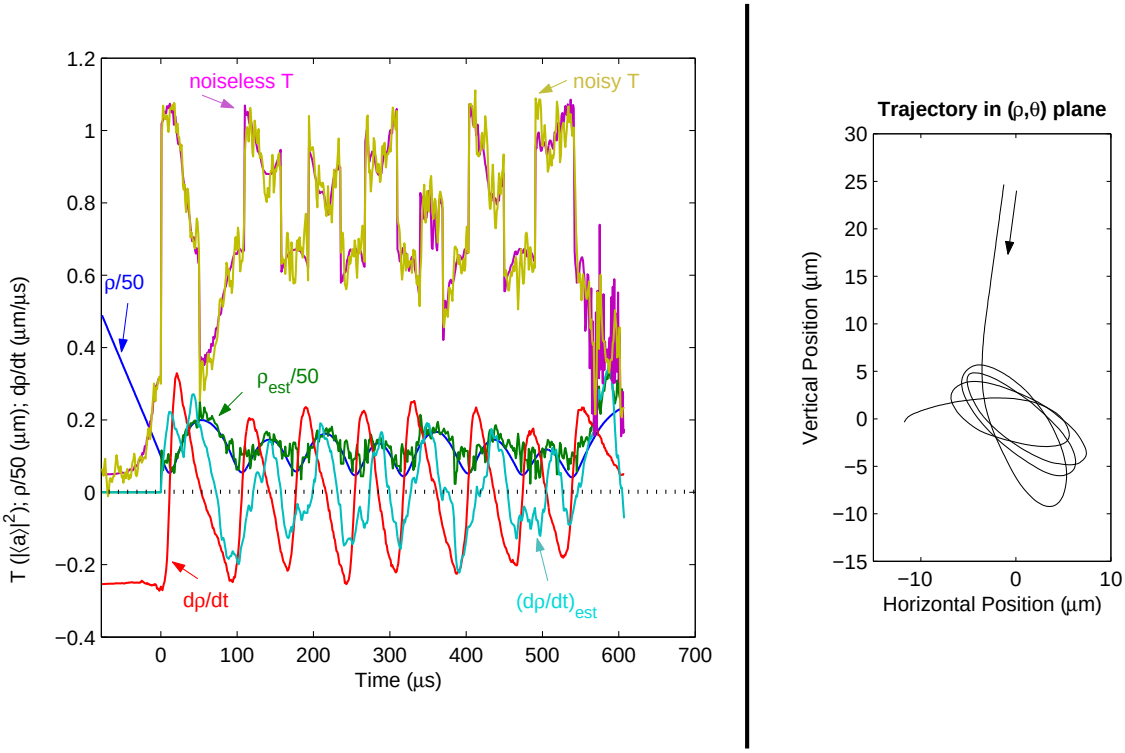

Figure 5.4: Measurement noise leads to filtering and unacceptable loop delay.

and the trajectory remains elliptical.

5.2.3 Account for Delays by Waiting a Cycle

Since it seems clear we cannot simply close the loop with the delays necessitated by measurement noise, we choose to address the problem by adding even more delay – that is, by detecting a switching condition ($\dot{\rho}_{est}$ crossing a hysteresis limit) and then waiting to switch the potential at a time which should catch the *next* oscillation of ρ. A first attempt in this direction would be to assume a fixed period for oscillations of ρ. In this case the additional wait before switching is given by this fixed period minus the estimator delay for $\dot{\rho}_{est}$. Since each switching time is now set by the detected signal from the previous cycle, the first switching opportunity (first minimum of T and maximum of ρ) will be missed in this strategy. Rather than miss this cooling cycle, we impose a single switching event a fixed time after the initial trap turn-on. Thus the potential switches *exlo*→*hi* on the initial trigger, *hi*→*lo* a fixed time later,

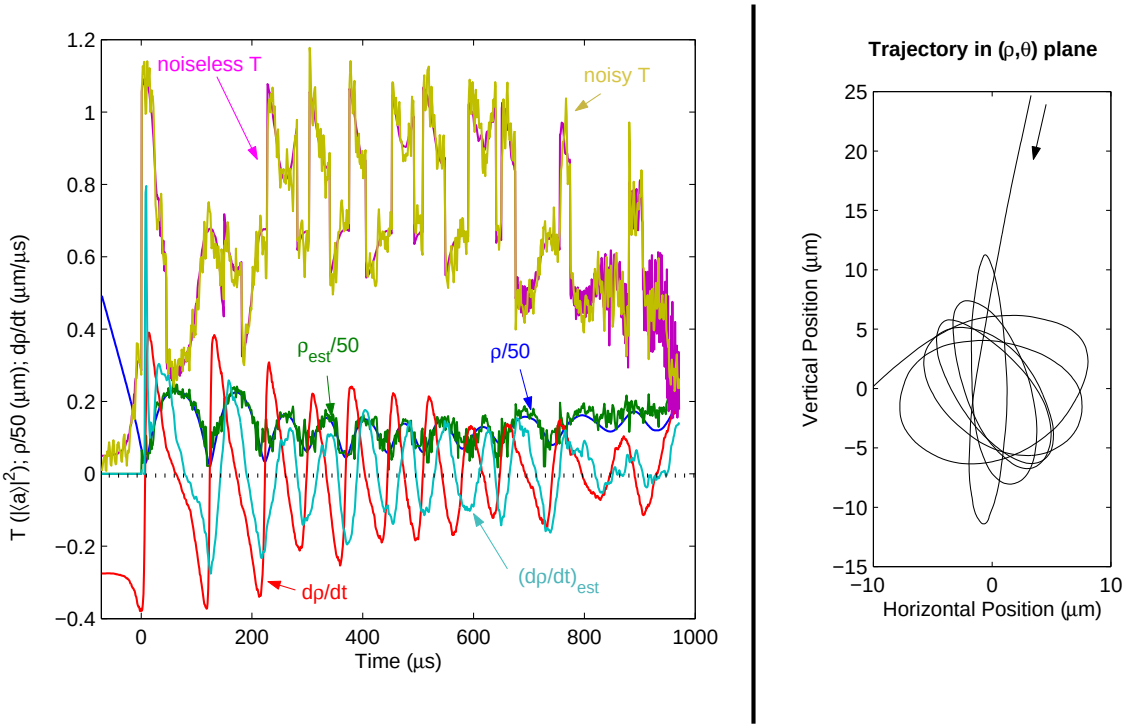

Figure 5.5: Delay is dealt with by tracking ρ turning points and using this information from each cycle to switch the potential at the predicted *next* turning point.

and $lo \leftrightarrow hi$ thereafter based on the last zero-crossing time of $\dot{\rho}_{est}$.

However, the actual dynamical period varies by easily a factor of two over the course of an atom's trapping lifetime due to changing amplitude of oscillation in the anharmonic potential. Thus the fixed-period assumption is a poor one. A better strategy is to record the length of each period in $\dot{\rho}_{est}$ and assume each cycle will be the same length as the previous recorded one. Thus the "waiting time" estimate will adjust itself as the dynamical period changes, though it will in general be one cycle behind. This strategy is employed for the trajectory shown in Figure 5.5. The initial switch occurs 45 μs after trap turn-on, the least-squares window is 40 μs, and the "wait time" between $\dot{\rho}_{est}$ limit-crossings and the resultant potential switches is given by the previous period minus 20 μs. This switching strategy, with deliberate delay based on an active measurement of the ρ oscillation time, appears to be a viable means of performing control in the presence of sensor noise and its associated loop delay.

5.2.4 Comparisons with Open Loop Strategies

To evaluate the effects of feedback more quantitatively, we introduce a figure of merit for the damping of radial oscillations in an atomic trajectory. Since the goal of the control strategy is to confine $\dot{\rho}$ near zero, the performance can be measured by comparing the variance of $\dot{\rho}$ over intervals of equal duration near the beginning of the trajectory and after feedback has been operating for some time. We choose a time window of duration 200 μs as long enough to encompass well over one cycle of the radial motion. The comparison is taken between two such windows separated from one another by 200 μs; this delay is selected as long enough for several cycles of feedback action, yet short enough so that the statistic exists for a large fraction of trapped atom events. Thus our figure of merit for feedback performance is given by

$$M = \frac{\sigma^2_{15\,\mu s \to 215\,\mu s}(\dot{\rho}_{est}(t))}{\sigma^2_{415\,\mu s \to 615\,\mu s}(\dot{\rho}_{est}(t))} \tag{5.5}$$

where times are measured from the initial trapping ($exlo \to hi$) switch. Large values of the quantity M correspond to well-damped radial motion, $\rho(t) \to constant$, though orbits may still be circular at any radius $\rho \geq 0$. (Damping in the sense of actual energy removal is discussed further in Section 5.2.5.) Small ($\sim 50 - 100\,\mu s$) changes in delay time or window size have been investigated and do not appreciably change the nature of the results for M.

Figure 5.6 shows histograms of M for several data sets in which different switching protocols, detailed in Table 5.1, have been employed. Each data set is generated by simulating a fixed number of individual atom drops from the known distribution of initial conditions. Only some fraction of trajectories result in a triggering/trapping event, and of these only a fraction of atoms have dwell times long enough to compute a value for M. Thus, for example, set C2 was generated from 5,000 trajectories, yielding 1,335 trigger events and 147 trajectories for which M could be obtained (i.e., with dwell times at least 615 μs).

While Table 5.1 gives the specifics of each data set represented in Figure 5.6, the essential comparison is between closed loop – i.e., active feedback – algorithms and

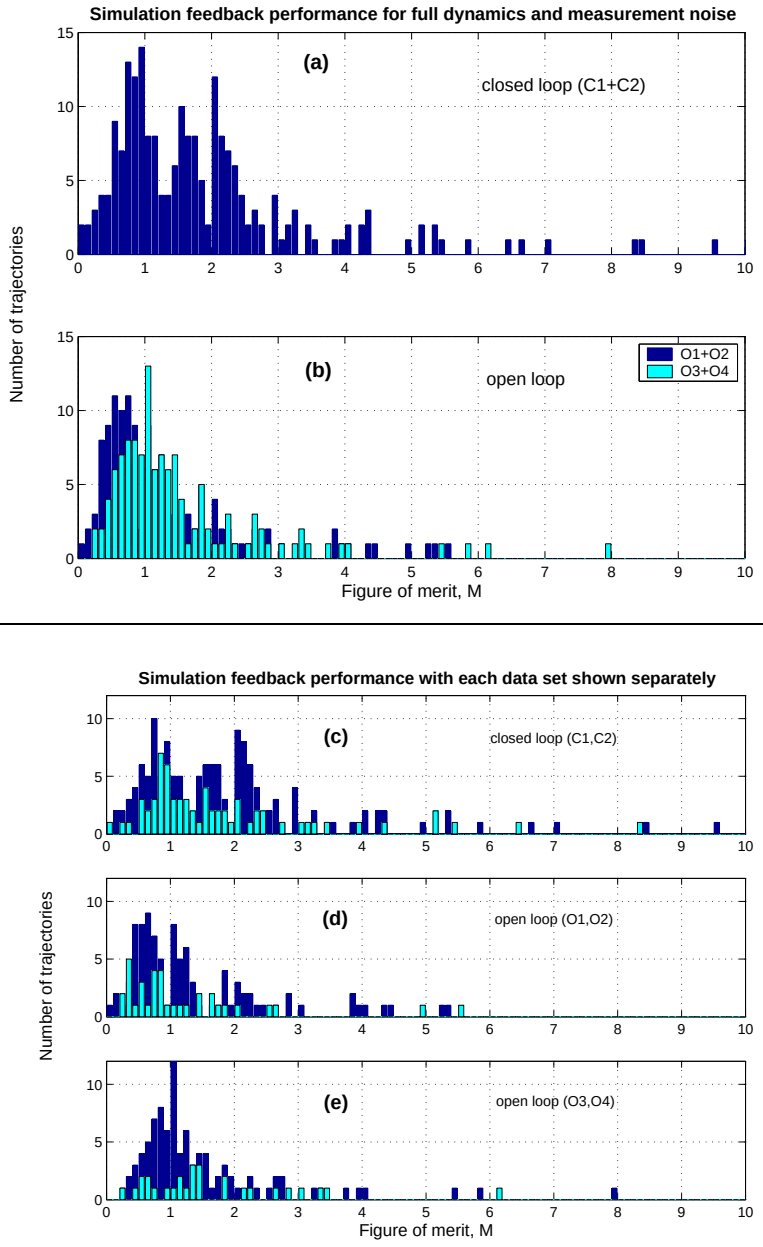

Figure 5.6: Figure of merit for feedback in closed- and open-loop cases (full dynamics). The data of (a,b) are redisplayed in (c,d,e) with each data set shown separately. In (c,d,e), the smaller data sets (C1,O1,O3) are shown with light-colored bars.

Data set	switching protocol	# atom drops	# triggers
C1	Closed loop Switch once 45 μs after initial trigger Thereafter switch (previous cycle length) - 20 μs after $(d\rho/dt)_{est}$ crosses limits	2000	534
C2	Same as above (C1)	5000	1335
O1	Open loop Switch every 45 μs after initial trigger	2000	561
O2	Same as above (O1)	5000	1319
O3	Open loop Switch every 35 μs after initial trigger	2000	552
O4	Same as above (O3)	5000	1322

Table 5.1: Exact conditions used for data of the full-simulation histograms.

open loop counterparts which simply switch potentials in a predetermined sequence independent of real-time position information for the individual atom. The closed loop algorithm is that of Figure 5.5, in which measurement noise and loop delays are dealt with by waiting nearly one cycle to apply the knowledge of motion gained during the previous oscillation. The open loop algorithms, in contrast, simply switch the potential between *hi* and *lo* at fixed intervals following the initial trapping event; the fixed interval is chosen to coincide with a reasonable average value for an atomic oscillation period.

Closed loop, active feedback clearly damps radial oscillations more effectively than its open loop counterpart. The M distributions in Figure 5.6(a,c) have larger mean than those in Figure 5.6(b,d,e), with the closed-loop histograms showing many trajectories pushed out to higher values of M by the active feedback. These results indicate that real-time measurements of $\rho(t)$ can indeed be applied to facilitate cooling of a single atom's motion in that dimension.

Further refinements should improve the performance of the algorithm. For instance, the cycle-length predictor could be changed to allow asymmetries between $\dot{\rho} > 0$ and $\dot{\rho} < 0$ half-cycles. Additionally, the filters themselves could be adjusted or replaced with better estimators which incorporate information about angular momentum L.

5.2.5 Performance with Axial Motion Suppressed

In a final set of simulations, we investigate the performance of our radial cooling algorithm in a setting where the axial atomic motion is independently suppressed. With no (or little) axial motion, axial heating no longer limits trapping times and the effects of radial feedback can be seen more clearly. To achieve this in simulations, we impose an *ad hoc* elimination of the axial dimension; however, this case could be relevant to several future experimental scenarios. For example, trapping and sensing mechanicsms, both currently accomplished with the same probe beam, could be separated to allow a trapping field with a low scattering rate and much-reduced axial diffusion. Alternatively, the separation of axial and radial timescales could be exploited, either to apply axial cooling [96] between cycles of radial feedback or to simply avoid extra axial heating by ramping the potentials up and down at a rate that appears adiabatic to the axial motion while still impulsive in the radial dimension.

Prospects for implementation aside, simulations with no axial motion demonstrate some aspects of the radial feedback protocol that are otherwise less transparent. We explore this regime with a set of simulations that differ in three ways from those presented above. First, the axial dimension is eliminated entirely and the atom is artificially constrained to remain at rest at an antinode of the standing-wave cavity field. Second, since axial heating is no longer an issue, we employ somewhat deeper trapping potentials than in the simulations above. The weak probe level is still $n_{exlo} = 0.05$ photons in the empty cavity, but now we turn on the trap initially at a level $n_{exhi} = 0.6$ photons in the empty cavity, and we feed back by switching between this and the weaker level $n_{hi} = 0.3$ photons. The effective potentials thus generated are $\sim 50\%$ deeper than in the previous simulations using n_{hi} and n_{lo}. Finally, without axial motion we employ a coarser computational timestep of $(1/30)$ μs.

Figure 5.7(a) shows the feedback figure of merit for the cases of closed-loop feedback, constant trapping at n_{hi} or n_{exhi}, and open-loop switching. Taking advantage of longer overall atom dwell times, we can now consider time windows separated by

a greater delay, so the quantity displayed here is

$$M' = \frac{\sigma^2_{15\ \mu s \to 215\ \mu s}(\dot{\rho}_{est}(t))}{\sigma^2_{1015\ \mu s \to 1215\ \mu s}(\dot{\rho}_{est}(t))} \tag{5.6}$$

rather than the original M of Figure 5.6. The result is qualitatively similar to that of the full simulation, with the closed loop strategy performing significantly better than its open loop counterpart. In this case, the mean value of M' for closed-loop feedback is ~ 2.5 times greater than for open-loop switching or for no switching.

The data of Figures 5.6 and 5.7(a) indicate that our feedback algorithm acts to drive $\dot{\rho}$ to zero, i.e., to circularize atomic orbits at a constant value of ρ. Such demonstrations, however, leave open the question of whether this algorithm actually removes total energy from the radial motion. With axial motion eliminated, we can now explore this issue by asking how the feedback algorithm affects trapping lifetimes. Figure 5.7(b) shows atom dwell times for the same three cases of closed-loop feedback, constant drive levels of n_{exhi} or n_{hi}, and open-loop switching. The increase in lifetime for closed-loop feedback is immediately apparent. Indeed, the closed-loop results agree well with an exponential lifetime of 8.9 ms, as contrasted with 2.6 ms for trapping at n_{exhi} alone, 1.9 ms for trapping at n_{hi}, and 1.1 ms for open-loop switching between the drive levels. This effect shows that active feedback as applied to the radial dimension does act to remove radial energy, in addition to simply pinning ρ to a constant value. Note that all these lifetimes are enhanced relative to the full three-dimensional case, in which both experiments and simulations have mean trapping times of only $\sim 400\ \mu$s. Lack of lifetime enhancement from radial feedback in the full simulation is some indication of the very weak mixing between axial and radial motion, so that cooling in one of these dimensions does little to control temperatures in the other.

5.3 Outlook for Experimental Implementation

The feedback simulations discussed in this chapter have been conducted with very close reference to the conditions realized in the experiment of Ref. [25], in particular

(a)

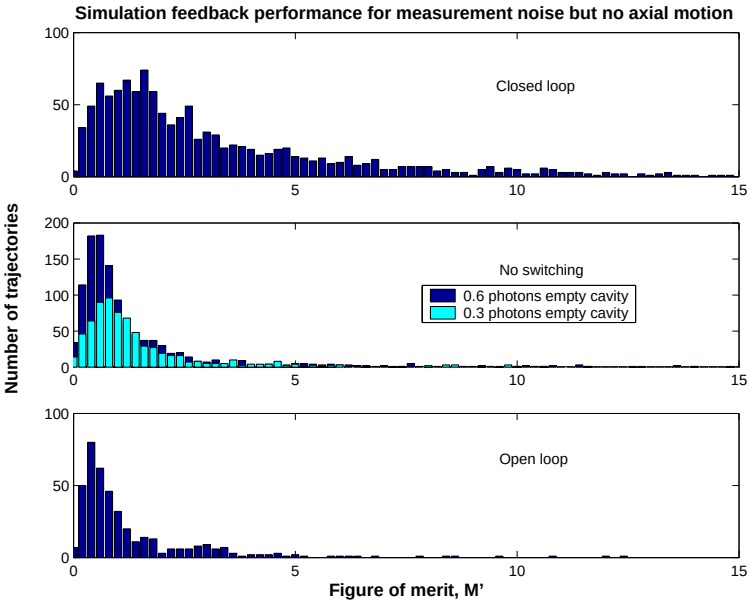

(b)

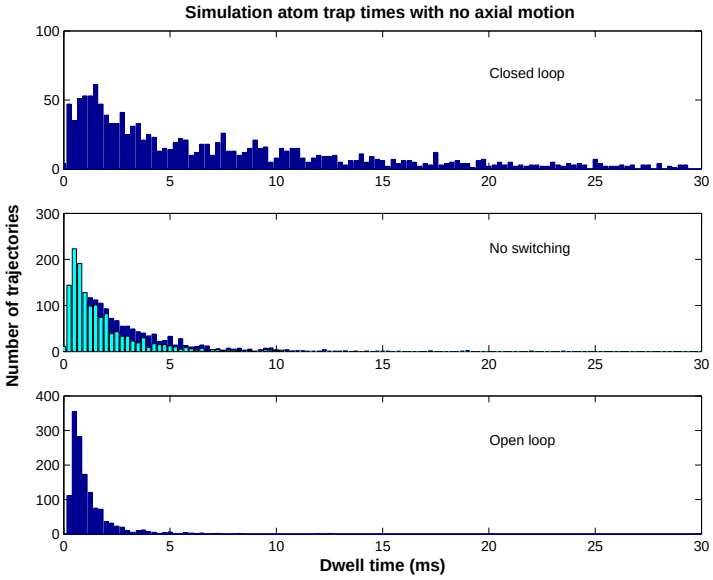

Figure 5.7: Feedback performance and lifetime enhancement in simulations with no axial motion.

for cavity properties, trapping statistics, and signal-to-noise in the balanced heterodyne detection. The current experimental effort, aimed at realizing the feedback strategies described here, employs very similar conditions while incorporating some improvements as in Ref. [51, 52]. Notable changes from Ref. [25] include a slightly shorter cavity, improved vacuum pressure in the cavity region ($3\text{x}10^{-10}$ torr) enabled by a differentially pumped double chamber, and cavity length stabilization via an error signal generated by an independent laser one free spectral range away from the cavity QED light. Implementation of the feedback strategy described above is clearly outside the regime of analog electronics, so an additional modification is the use of digital processing and FPGA technology [97]. With these tools, experimental data similar to the simulation results presented above seem well within reach; the experimental effort is discussed in more detail in the next chapter.

It seems reasonable to ask how much experimental data should be necessary to exhibit a distinction between active feedback and open loop schemes. From the simulations of Figure 5.6(c,d,e) (Table 5.1), we see that with data sets of about 500 trapped atoms the differences between open- and closed-loop schemes already begin to become apparent, and these differences are well demarcated with two or three times that much data. With fairly conservative estimates of one trigger per MOT drop and one MOT drop every 5 seconds, this means significant effects could well be seen with just one to two hours of experimental data at each setting. Much more data collection is experimentally realistic, allowing exploration of a wealth of additional questions.

5.4 Current Limits and Future Directions

The feedback algorithm developed above for the atomic radial position ρ is subject to basic limits arising from dynamical and measurement noise in our system. These limits can be expressed as lower bounds on the (one-dimensional) temperature for ρ. Because the strategy is based on discrete switching, with feedback delayed and timed based on the previous switching-cycle length, the control is always based on informa-

tion gathered over the previous motional cycle. Thus dynamical noise over an atomic motional timescale will set a lower limit on radial temperature T_ρ. Referring to Figures 4.11 and 4.13, we find that momentum diffusion (due to spontaneous emission) gives a typical energy increase per radial oscillation time of $\Delta E_{\tau_r} \sim 0.02(U_0) \sim k_B(50\,\mu K)$. Furthermore, measurement noise places a limit on the detectable amplitude of ρ variations. This amplitude depends strongly on the absolute value of ρ due to the nonlinearity of the $T \rightarrow \rho$ mapping. However, using the measured sensitivity (see Section 4.3) from the atom-cavity microscope, we estimate that over a motional cycle we can resolve ρ oscillations of amplitude $(20\,nm/\sqrt{Hz})\sqrt{1/2\pi\tau_r} \approx 0.77\,\mu m$. On the side of the cavity mode, where the effective potential is steepest, this corresponds to $T_\rho \approx 150\,\mu K$. While this limit corresponds closely to the simulations of the previous section, where axial motion is suppressed, the full simulation never reaches this limit because axial heating cuts off atomic lifetimes too quickly. Thus improvements to address axial heating are of great interest for seeing the full effect even of radial cooling.

Beyond the experimental and algorithmic variations already discussed, a number of broader questions are raised by the use of active feedback to dynamically cool a component of motion for a single atom. One question deals with the ultimate limits of such a cooling mechanism. Within the current experimental setting, limits to radial cooling arise from atomic lifetimes (dominated by axial motion), but are also constrained by the dynamical noise and by the shot noise of detection. Some lifetime and dynamical noise issues could be addressed by separating trapping and sensing, for example by using a far off resonance trap in conjunction with a sub-photon level cavity QED probe [52]. The remaining issues would then center on signal-to-noise for the atomic position measurement, as well as on limits imposed by backaction of the measurement itself as it approaches the standard quantum limit [82, 98].

Extension of active feedback beyond the ρ dimension raises related questions. The question here is one of using various techniques – for example, a symmetry-breaking potential or frequency-domain filtering of the signal – to estimate and control a three-dimensional vector using the time record of a single quantity, the transmitted light

field. Undoubtedly, different driving parameters, detection methods, and data processing will be appropriate depending on the relative importance placed on information about each dimension of the motion. These questions address in various ways some basic issues of optimal state estimation and control for single quantum systems, and this experimental system promises to be a rich one for exploring such issues in further detail.