

# Chapter 6

## Farther Up and Farther In

This chapter describes the current experiment to implement active feedback on the atomic radial position. In contrast to the general experimental summary of Chapter 3, here I treat in detail those aspects of the current experiment which were not present in the atom-cavity microscope setup.

The changes are in four broad categories. First, a more sophisticated vacuum chamber allows the pressure in the cavity region to be in the  $10^{-10}$  torr range even though the MOT loads from a Cesium vapor at several  $10^{-8}$  torr. The motivation for better vacuum in the cavity region is twofold: to allow longer trapping lifetimes for atoms should we reach that stage with the current apparatus, and to help protect the cavity mirrors from the slow death of Cesium coating. Since the chamber now has separate MOT-loading and cavity portions, we use a double MOT as well; the first MOT is loaded from Cesium vapor, cooled, and dropped through the differential pumping hole to be caught in a second MOT just above the cavity.

Second, the generation of probe light and the stabilization scheme for laser frequency and cavity length has been completely overhauled. We abandoned the Ti:Sapph for probe beam generation to avoid dependence on it and its 5-Watt pumping laser in favor of the cheaper, simpler diode laser option. At the same time, the desire for a cavity lock that was quieter and could tolerate longer atom transits led us to introduce an independent “locking” laser whose transmission now provides the error signal for locking the cavity length. As this laser is one free spectral range away from the cavity QED resonant light, we can use substantially more optical power in this

beam for a better error signal and more aggressive locking. The locking laser also has the clear advantage that it (and the cavity lock) can be left on throughout the atom-cavity interaction so that the cavity need not be left to drift for the duration of a transit.

The third change is, of course, what drove the rebuilding of the entire experiment: the absolute need for a new physics cavity. In replacing the ACM cavity we attempted to do better in a couple of ways. We adjusted details of the mirror mounting and gluing process in an attempt to eliminate some birefringence problems that have plagued earlier cavities. We also designed the cavity length exactly in order to leave open the possibility of a FORT at the “magic wavelength” [99, 52]; for a cavity as short as ours, the free spectral range is about 40 nm and so exact planning is required to make the cavity simultaneously resonant at 852 nm and near 935 nm. The new cavity succeeds in these areas, but brings its own unexpected quirks which we have had to deal with in pulling together the experiment.

Finally, getting serious about feedback as in Chapter 5 has led to the replacement of our old data acquisition and signal processing with a new Gage board for data acquisition and field-programmable gate array, or FPGA, for triggering and feedback. The old regime of data acquisition was Labview, notoriously slow and cumbersome to use online. Worse yet, our signal analysis and feedback tools were analog – essentially resistors and capacitors – and were ill-equipped to deal with a fundamentally digital feedback algorithm. In the feedback experiment we are performing right now, analog methods run into the problem of excessive loop delay due to differentiating the noisy signal. However, there is a more fundamental reason for the transition to digital. Referring to Figure 5.1, we recall that when we adjust the driving intensity to implement feedback, we change the whole system in a highly nonlinear way. The effective potential for the atom changes, but so do the dynamical noise for the atomic motion, the measurement noise, and the transfer function from cavity transmission to atomic position. Thus for each value of the control parameter the servo must adjust to reflect the altered transfer functions throughout the system. FPGA’s, with their reconfigurable nature and high speed a/d, d/a, simple signal analysis, and conditional

logic, come to the rescue in the implementation of effective feedback in the system [97].

## 6.1 Vacuum and MOTs

A schematic of the differentially pumped vacuum chamber is shown in Figure 6.1. The main upper chamber is constructed from a Kimball Physics small spherical hex, while the main lower chamber is the custom NorCal vacuum feedthrough retained from the atom-cavity microscope main chamber. They are connected via a long, narrow differential pumping hole to support a sizeable pressure difference between the upper and lower chambers. The lower chamber is pumped with a 55 l/s Varian Starcell ion pump and equipped with a nude Bayard-Alpert ionization gauge to monitor vacuum pressure, while the upper chamber is pumped with a 20 l/s Starcell and has a straightforward B-A gauge. The upper chamber contains the Cs reservoir and is operated as a Cs vapor cell for loading of the upper MOT; the pressure in this part of the chamber is a few times  $10^{-8}$  torr. The lower chamber, on the other hand, is maintained at several  $10^{-10}$  torr and contains the physics cavity itself.

### 6.1.1 Chamber Design and Construction

The chamber itself, surrounded by MOT coils and optics for the MOT and probe, is shown in Figure 6.2. The lower chamber has two large (6") windows for generous optical access along the cavity (and primary MOT) axis; in this direction we couple light into the cavity and also align a retro-reflected MOT beam skimming the top of the cavity substrates. The MOT is also imaged through the large windows. The chamber has eight smaller ports in the transverse plane. Four of these are windows permitting MOT beams in the remaining two dimensions. One goes to the lower ion pump/gauge region, one connects to the upper chamber, and one is used for imaging of the lower MOT. A final window (the bottom) is unused and provides optical access only to the bottom surface of the cavity mount piece.

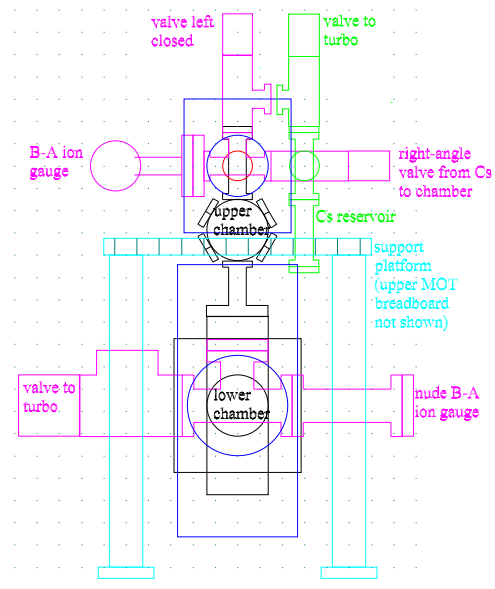
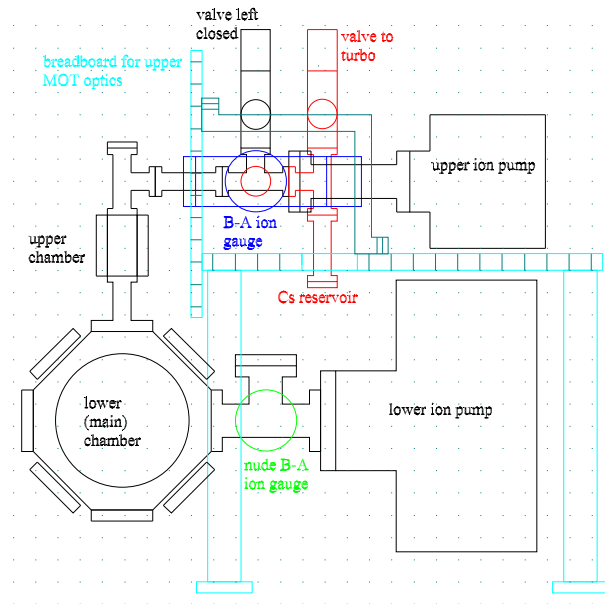


Figure 6.1: The differentially pumped vacuum chamber, seen from along the physics cavity axis (top view) and from the side (bottom view).




Figure 6.2: Actual upper and lower chambers with MOT coils and optics in place.

The upper chamber has two 2-3/4" windows along the upper MOT axis and six 1-1/3" ports evenly spaced around the sides. Four side windows are used for MOT beams, giving us three pairs of MOT beams which are not strictly orthogonal. The effect on MOT performance has been minimal, both here and in the lower chamber where the cavity-axis MOT beam pair are slightly angled as well. The upper chamber's small bottom port connects to the lower chamber, and the top port connects to a window and the upper pump/gauge/reservoir region. Upper MOT imaging is performed through the top window and through the front 2-3/4" window of the hex. All the vacuum windows are antireflection coated from 820-930 nm by Larson.

### 6.1.2 Cesium Reservoir and Vapor Cell Loading

The Cesium reservoir is separated from the upper chamber by an all-metal right-angle valve (from VAT). Cesium is loaded into the (already pumped-down) chamber as illustrated in Figure 6.3. The reservoir region is valved off from the main chamber and opened to air (or nitrogen if Cesium has already been present in it). A complete glass ampule containing Cesium metal and a buffer gas is cleaned and placed inside the bellows, with one end resting against each side of the bellows wall. The reservoir region is connected to a turbo pumping station and pumped down from atmosphere. Bending the bellows (as forcefully as necessary) cracks open the ampule, and the turbo pump load surges as buffer gas is pumped out. After ultrahigh vacuum (UHV) is attained, the reservoir is closed to the turbo pump and opened to the rest of the upper chamber if desired. Cesium is kept in the reservoir by closing the right-angle valve to the chamber and Peltier cooling the bellows end. The vapor cell in the main upper chamber is replenished by opening the right-angle valve and gently heating the reservoir. Initial buildup of the vapor cell took about ten days with the Cesium open to the chamber and kept a few degrees above room temperature. This time is consistent with initial coating of the surfaces between the reservoir and the upper MOT location. Subsequent re-loading of the vapor cell is much faster as surfaces remain mostly coated with Cs.




Figure 6.3: Cesium reservoir and loading procedure

### 6.1.3 Differential Pumping

The pressure difference between the upper and lower chambers is made possible by means of a deliberately low-conductance (long and narrow) connection between them, known as a differential pumping hole. Our differential pumping hole forms a tube of length 122.6 mm and diameter 5.75 mm which connects the upper and lower chambers. The tube begins 1.9 cm below the center of the upper chamber, and does not extend far enough into the lower chamber to cut off any optical access through the large windows there. The dimensions of the hole are constrained by two requirements: support for the desired pressure differential between upper and lower chambers, and clearance for dropping a cloud of cooled Cs atoms through the hole.

To support a ratio of pressures  $P_{upper}/P_{lower} \geq \Delta$  in a system where the lower chamber effective pumping speed is  $S$ , the conductance through the differential pumping hole must satisfy  $C \leq S/\Delta$ . A hole of diameter  $D$  and length  $L$  has a conductance (for Cesium) of  $C = 5.7 \frac{D^3}{L + \frac{4}{3}D}$  where the dimensions are expressed in centimeters and the conductance is measured in liters per second. Thus a minimum pressure ratio  $\Delta$  imposes a relationship between the length and diameter of the hole,  $L \geq 5.7 \frac{\Delta}{S} D^3 - \frac{4}{3}D$ . The lower chamber is pumped by a nominally 55 l/s ion pump, estimated to be close to  $S_{pump} = 25$  l/s in practice for Cesium. The connection to the pump is through an arm of diameter 1-1/2" and total length 7.68", giving a conductance of  $C_{path} = 12.8$  l/s. Thus an overall effective pumping speed for the lower chamber is estimated at  $S = (\frac{1}{S_{pump}} + \frac{1}{C_{path}})^{-1} = 8.5$  l/s. Under these conditions we set a target of  $\Delta \approx 100$ .

In order to drop atoms through the hole with reasonable throughput, we first position the hole as close below the initial cloud position as we reasonably can without cutting off optical access. In reality we have allowed a vertical space of  $s = 1.9$  cm which generously allows for trapping beams and for some error in positioning of the initial MOT. Assuming the pumping hole is aligned perfectly with respect to gravity, a cooled atomic cloud with initial diameter  $d_0$  and transverse velocity spread  $v_s$  clears the hole if  $L \leq [(\frac{D-d_0}{v_s})^2 \frac{a_g}{2} - s]$ . Here  $a_g$  is used to denote the acceleration due to gravity. By assuming an attainable sub-Doppler temperature of



$3 \mu\text{K}$  ( $v_s = \sqrt{2}v_{rms} = 1.8 \text{ cm/s}$ ) based on a similar system quoted in [100], and an initial diameter of 1.5mm, we obtain bounds on  $L$  and  $D$  from this requirement.

Considering these two effects of pumping hole dimension, we plot in Figure 6.4 the constraints on  $L$  and  $D$  imposed by our target requirements. By these arguments our pumping hole should support a pressure ratio of  $P_{upper}/P_{lower} = 102$  and barely allow a cloud of initial diameter 3 mm and temperature  $6 \mu\text{K}$  to pass as described above. A cloud of temperature  $3 \mu\text{K}$  passes even allowing a 1% deviation of the pumping hole axis from the direction of gravity. Clearly alignment with gravity is an important step in the chamber assembly! In practice we maintain a pressure ratio of about 50 as measured with ion gauges in the upper and lower chambers. The transfer efficiency from the upper MOT to the lower chamber has not been calibrated but is certainly sufficient for our single-atoms experiment.

The differential pumping hole is a custom piece mounted to fit into the upper and lower chambers and the standard adapting nipple which connects them. In mounting the piece, care was taken to ensure a good vacuum seal everywhere except through the actual tube. At the same time, a design goal was to promote overall good vacuum pressure by avoiding “virtual leaks” arising from large surface areas close to one another, leading to low-conductance regions which are hard to pump out. Furthermore, any excess surface areas (in close contact or not) should be put on the worse-vacuum side of the design if possible. Finally, based on the bitter experience of Jun Ye and Dave Vernooy in their differentially pumped chamber, care must be taken to ensure that the pumping hole will stay in place during transient pressure differences that arise during initial pumpdown from atmosphere.

The mounting of the differential pumping hole is shown in Figure 6.5. The pumping hole is essentially a hollow tube of 304 stainless steel, with an inner diameter of 5.75 mm over its entire 122.6-mm length. Most of this length is taken up by the top portion, of outer diameter 10 mm, which begins in the upper chamber and extends through the adapting nipple while staying well away from its walls. At the very bottom of the nipple’s length, the differential pumping tube steps up to an outer diameter of 17.1 mm to snugly fit into the end of the nipple. The larger outer diameter persists




Figure 6.4: Pressure-ratio and MOT-dropping constraints on differential pumping hole length  $L$  and diameter  $D$ . For a given value of  $D$ , pressure requirements give lower bounds on  $L$ ; these constraint curves are shown with arrows pointing up from them. Transmission of an atom cloud gives an upper bound on  $L$ , so these constraint curves are denoted by downward arrows.




Figure 6.5: Design and mounting of the differential pumping hole. The adapting nipple attaches to the lower chamber via a 2-3/4" conflat flange, and to the upper chamber via a 1-1/3" conflat flange.

through the portion of the pumping hole length that extends into the lower chamber. Finally, at the junction with the bottom of the adapting nipple, three vanes extend radially from the pumping hole with through-holes for 6-32 screws. Corresponding threaded holes are drilled into the flat surface of the adapting nipple at this end, allowing firm attachment of the pumping hole piece to the chamber with 3/8"-long screws. It is important to plate these screws with gold before assembling the components; otherwise the UHV-clean screws and stainless-steel threaded holes will almost inevitably seize up and get stuck during assembly.

### 6.1.4 Getting Atoms to the Cavity: Upper and Lower MOT's

Atoms are delivered to the cavity mode volume by a cycle of

- loading the upper MOT in the center of the upper chamber, with upper anti-Helmholtz coils on and trap light on (3 seconds).
- cooling the upper cloud, with quadrupole coils off and trap light reduced and detuned (35 milliseconds).
- dropping the upper cloud, with coils and trap light off (100 milliseconds).
- catching the cloud in a lower MOT just above the cavity substrates, with lower antiHelmholtz coils on and trap light on (200 milliseconds).
- cooling the lower cloud, with quadrupole coils off and trap light reduced and detuned (35 milliseconds).
- dropping the lower cloud, with coils and trap light off, observing transits through the cavity, performing triggered trapping and feedback (25 millisecond average drop time to cavity mode, acquire data during 10 or 20 ms window around this).

A diagram of timing for this cycle is shown in Figure 6.6.

The trap light for both MOT's is generated from a single diode laser pumping an SDL tapered amplifier and single-passed through an AOM for intensity-switching capability. The final trapping beam is split to separate polarization-preserving fibers for the upper and lower MOT's. Thus the light in upper and lower chambers is switched as a single unit rather than individually; since the timing of upper and lower MOT's does not overlap anyway, this is an acceptable strategy. A double-pass AOM before the SDL amplifier controls the detuning of the trap light. The trap beam is tuned 11.8 MHz to the red of ( $F = 4$ )  $\rightarrow$  ( $F' = 5$ ) during MOT loading and 23.8 MHz to the red during sub-Doppler cooling. Repump light on ( $F = 3$ )  $\rightarrow$  ( $F' = 4$ ) is likewise generated from a single diode laser, split to pass through the same two fibers as the trapping light as well as through a third fiber which delivers some of it to the cavity input path itself. The repump light is left on continuously.




Figure 6.6: Timing diagram for experiment with double MOT.

The magnetic field gradients for the MOT's are provided by upper and lower sets of anti-Helmholtz coils (main coils), with additional Helmholtz coils in all three dimensions to provide zeroing of the magnetic field. The upper main coils are 14 AWG single-build insulated magnet wire wound onto aluminum forms and mounted without direct contact with the chamber. The lower main coils are of similar construction but wound with 10 AWG heavy-build insulated magnet wire. While anti-Helmholtz magnetic field coils may ideally have their dimensions (coil radius and separation) designed to realize as closely as possible a perfect linear gradient field at their center [101], in practice our coil design is driven instead by crowding and requirements of geometry close to the chamber. The properties of upper and lower main coils are given in Table 6.1. The table gives the average radius ( $R$ ) of each coil, the average separation ( $2A$ ) between the coils, the number of turns in each coil ( $N$ ), and the resulting field strength at the center of each coil for an applied current. The operating current and the resulting field gradient at the MOT location are also shown. A bias coil below the chamber is also used to push the lower MOT field zero up about 3 mm from its natural location, essentially at the physics cavity itself.

|                               | upper MOT            | lower MOT          |
|-------------------------------|----------------------|--------------------|
| trap power (total)            | 43 mW                | 65 mW              |
| loading time                  | 3 s                  | 200 ms             |
| sub-Doppler power             | 2.4 mW               | 4.3 mW             |
| cooling time                  | 35 ms                | 35 ms              |
| repump power (total)          | 4.3 mW               | 7.2 mW             |
| coil radius ( $R$ )           | 4.0 cm               | 9 cm               |
| coil separation ( $2A$ )      | 6.0 cm               | 18 cm              |
| turns per coil ( $N$ )        | 18                   | 118                |
| field at center of each coil  | (2.85 G/A, 2.83 G/A) | (9.0 G/A, 8.4 G/A) |
| operating current             | 8 A                  | 15 A               |
| field gradient on-axis at MOT | 8.4 G/cm             | 16 G/cm            |

Table 6.1: Upper and lower MOT parameters

Optical powers and magnetic field properties for the upper and lower MOT's are given in Table 6.1. The number of atoms in each MOT has not been calibrated, and rough magnetic field zeroing through time-of-flight imaging has so far resulted in hot

sub-Doppler temperatures ( $\sim 60 \mu K$ ). However, as these conditions have produced the desired result of one to several atom transits per MOT-dropping cycle, further optimization has not yet been pursued.

## 6.2 Probe Beam Generation and Locking Scheme

This experiment has seen considerable modification in probe beam generation and in locking of both the probe beam frequency and the cavity length itself. The new lab schematic borrows the concept of the cavity and probe lock used in [51, 52, 100], where the error signal for locking the physics cavity length is generated by transmission of a separate “locking laser” on a different longitudinal mode from the cavity QED light. The frequencies of this laser and the probe laser are both locked to modes of a stable reference cavity introduced for this purpose. Finally, Cesium FM saturated absorption spectroscopy of the probe beam provides an error signal for maintaining the reference cavity at a fixed frequency and tying the entire system to the atomic transition. The probe beam itself is no longer generated from a Ti:Sapph, but comes from a grating-stabilized diode laser at 852nm (total output power 40 mW).

Figure 6.7 shows a schematic of the frequency shifting and locking setup for the experiment. A portion of the probe laser power, as a single-frequency beam with no sidebands, is sent through a fiber to the probe and local oscillator paths. The probe beam is double-passed through AOM’s before being sent to the cavity. The probe is thus shifted down by 2(110 MHz) and up again by 2(95 MHz) to provide the 30 MHz LO-probe frequency offset used for heterodyne detection. The use of two separate AOM’s avoids the presence of strong RF signals at the heterodyne frequency itself; if a single AOM at 15 or 30 MHz is used, electrical pickup of the driving signal can cause excess noise on the heterodyne.

The probe light *not* used for the probe and LO paths is double-passed through an acousto-optic modulator (AOM), sent through a electro-optic modulator (EOM) where it acquires 5.1 MHz FM sidebands for locking purposes, and finally sent to the reference cavity and to Cesium. The probe beam detuning from atomic resonance

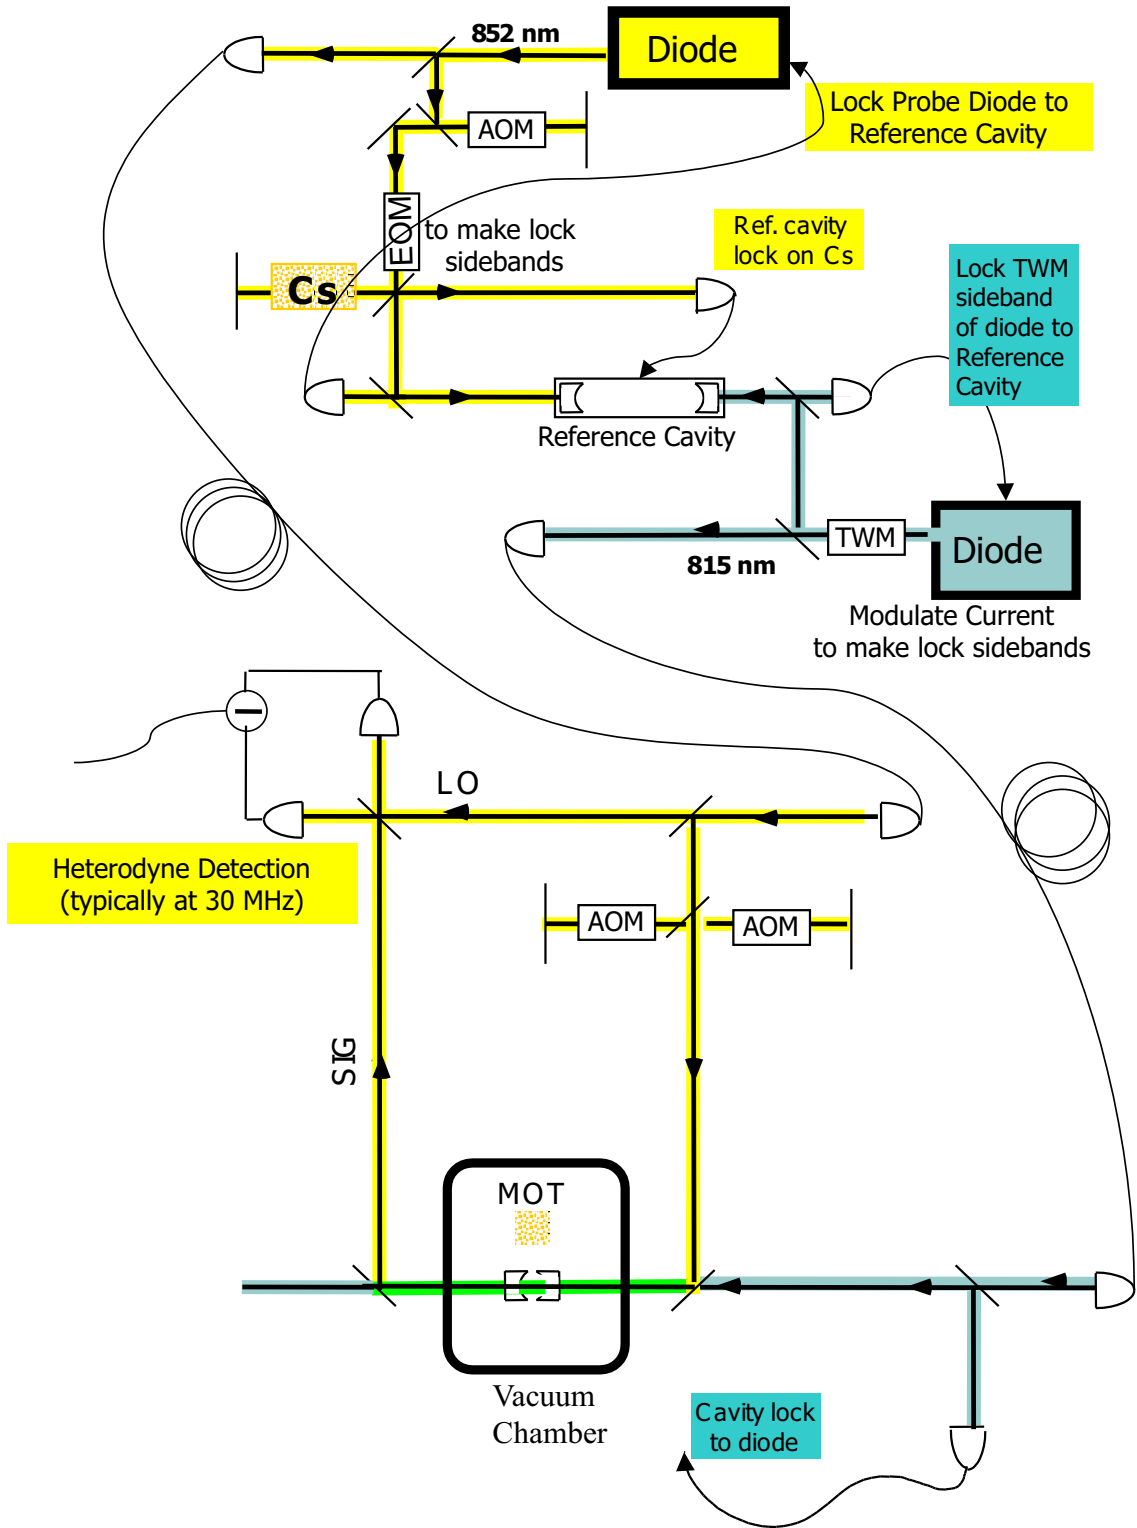


Figure 6.7: Schematic of laser and cavity stabilization



can be set by tuning the double-pass AOM frequency and locking to different Cesium saturated absorption features.

The locking laser operates at a wavelength of 814.8 nm, which is one free spectral range away from 852 nm on the physics cavity. This laser was originally an 825 nm unit pulled to 814.8 nm by cooling to about 12°C. For better linewidth and stability, we have recently replaced it with an 808 nm diode from JDS Uniphase which is grating-pulled to 814.8 nm and operated at 24°C to increase stability at this wavelength. The locking diode is current modulated to produce FM sidebands at 6.2 MHz on the entire lock beam; these sidebands are used for locking of the locking laser to the reference cavity and also of the physics cavity to this laser. The lock laser is also passed through a traveling-wave modulator (TWM) which generates sidebands as far out as 1 GHz; the TWM-generated sideband is locked to the reference cavity but the physics cavity is locked to the carrier. Tuning of the TWM frequency ensures consistency in the resonance conditions of both lasers (probe and lock) with both cavities (reference and physics).

Both lasers are locked to the reference cavity in reflection using the Pound-Drever-Hall technique [102, 103]. The locks employ the lab-standard two-branch locking circuit, worked on by many group members in sequence, including notably Joe Buck and Hanns-Christoph Naegerl, with the most recent version coming from J. M. Geremia. The demodulated error signal drives the diode current via a fast branch and also drives the grating-mount piezo via a slow branch with an integrator.

The reference cavity has finesse  $\sim 10,000$  and length 30 cm for a linewidth of about 50 kHz. The mirrors are mounted on either end of a quartz tube, with one mirror mounted on a PZT for length adjustment. The cavity is solidly held in a series of vibration-isolating materials, but is not kept in vacuum or actively temperature stabilized. The probe laser Cs spectroscopy setup provides an error signal used to lock the reference cavity to an atomic transition. This lock uses the slow branch of the same two-branch locking circuit to feed back to the reference cavity piezo; the important feature in this lock is simply a good integrator for slow drifts.

Finally, the physics cavity itself is locked to the locking laser carrier, as described

above, again in reflection and via the Pound-Drever-Hall technique. The locking circuit used to feed back to the physics cavity closes loop at  $\approx 5$  kHz; this error signal is then added to the voltage-divided output of a 300 V battery and sent to one of the shearing PZT's on the cavity mount. The second shear PZT is grounded to the cavity mount. A battery is used to generate the large DC offset voltage in order to avoid the introduction of electrical noise, particularly near the  $\approx 50$  kHz mechanical resonance of the cavity mount system. Noise near this frequency is unfortunately endemic in the lab, due to power supply switching of, among other things, HeNe power supplies and television monitors. Using a simple chemical-cell battery in place of a more conventional high-voltage power supply helps to reduce the electrical noise driving the cavity PZT. In voltage-dividing the battery output, however, it is important to consider the effect on the transfer function of the servo system. This point is critical because the shear mode PZT itself has a quite large (4 nF) capacitance. Thus the equivalent resistance introduced through the voltage divider creates a low-pass filter that tends to reduce the servo bandwidth. We use a 100 k $\Omega$  resistance in the voltage divider, resulting in a  $\sim 1$  kHz low pass for the final voltage onto the shear mode PZT (400 Hz to 1.6 kHz depending on the setting of the voltage divider).

The physics cavity finesse at 814.8 nm is about 240,000, a factor of two below the finesse at the center of the mirror coating curve (near 852 nm). The circulating lock beam power in the physics cavity must be small enough to cause negligible AC Stark shifts on the relevant cavity QED transitions. Since the lock beam is 16 THz from resonance, we may still use considerably more optical power for locking here than we were able to use with the cavity QED probe beam itself. In particular, a mode-matched power of 0.2  $\mu$ W of lock beam light on the physics cavity will shift the cavity QED transition by  $\lesssim 1$  MHz, at the edge of the frequency resolution we possess in the experiment. This estimate considers only the  $6S_{1/2} \rightarrow 6P_{3/2}$  transition; actually the Stark shifts will even be somewhat smaller due to contributions from the  $6P_{3/2} \rightarrow 8S_{1/2}$  transition at 794 nm (see Figure 3.2).




Figure 6.8: Cavity mount piece: detailed design

### 6.3 New Physics Cavity: Improvements and Difficulties

We replaced the atom-cavity microscope physics cavity with one which is similar but has several changes in its design and construction. Figure 6.8 shows a detail of the cavity mounting piece.

Before discussing this cavity in detail, I briefly review the procedure for assembling a cavity of this design. All components except the mirrors are first cleaned and baked with a view to their eventual introduction into a UHV environment. Mirrors are cleaned and inspected by hand as close as possible to the time of use. Even so, they are shielded from dust in the intervening time by careful storage, and by construction of the cavity in a clean hood. The shear-mode PZT's are glued to the copper mounting block, allowing one day for this glue to set, and leads are soldered to the PZT's and mount. One v-block is then glued on top of its PZT, and mirrors are glued to both v-blocks with small dabs of epoxy. These adhesives are allowed to set for one day. Finally, the cavity is aligned by placing the mounting block on a translation stage and positioning the free mirror/v-block unit in a holder attached to a mirror mount.

The cavity is aligned and set to the desired length; then the translation stage is wound down, glue is applied, and the translation stage is wound up again to form the final bond between the free mirror-v-block unit and its shear-mode PZT. Small adjustments may be made during the first 30 to 60 minutes after this glue is applied. After this bond has set for one day, the holder is removed and the resulting cavity can be measured and moved into place.

The current physics cavity is constructed from mirrors of the same coating run as the previous cavity, giving it a measured finesse of  $F = 480,000$ . Both mirrors have 10-cm radius of curvature, and the cavity length is now  $l_{eff} = 9.2 \mu\text{m}$ , giving a field waist of  $w_0 = 13.3 \mu\text{m}$  for the  $TEM_{00}$  mode. The new cavity has  $(g_0, \kappa)/2\pi = (130, 17)$  MHz, leading to  $m_0 = 2.0 \cdot 10^{-4}$  and  $N_0 = 5.2 \cdot 10^{-3}$ . This cavity was built, as described below, with particular attention to ensuring low birefringence and to selecting a specific cavity length. Unfortunately it also suffers from an unintended length sensitivity to magnetic fields, as chronicled below as well.

### 6.3.1 Controlling Birefringence

Cavity birefringence is discussed in more detail in Chapter 7, but here I simply remark that it has been an issue in the construction of all very high finesse cavities to date, and can largely be traced to mirror mounting techniques rather than to inherent properties of the coatings themselves. Stress induced by gluing (or clamping) transmits through the substrate to the mirror surface and induces a troublesome birefringent phase shift in the coating.

As mentioned in Chapter 3, the previous prescription for reducing cavity birefringence was to glue down the substrates with as little glue as possible, as far from the mirror surfaces as possible. Furthermore, our experience was that a low-birefringence cavity, when subjected to a heating and cooling cycle (to  $\approx 100^\circ\text{C}$ ), often developed birefringence as much as an order of magnitude larger than the original. When such a cavity was disassembled, the same mirrors could again be used to make a low-birefringence cavity; thus blame still focused on stress induced by the mount, this time

through a process of final glue setting or differential thermal expansion/contraction of mounting components. The end result was a prohibition on heating of physics cavities, meaning a complete inability to bake out a vacuum chamber after insertion of the cavity.

To address this concern we maintained the same basic mount design but Kevin Birnbaum did some research and substituted different materials for two components: the v-blocks and the glue connecting mirror substrates to v-blocks. Previously the v-blocks were aluminum pieces and the mirrors were glued to them with small dabs of Torr-Seal (the commonly used Varian catalog name for a UHV epoxy otherwise available as Dexter Hysol 1-C). According to manufacturer specifications, Torr-Seal sets to 90% hardness in 24 hours at room temperature. Complete cure is specified in one hour at 82°C or three days at room temperature. Our cavity construction method takes place entirely at room temperature. During a heating and cooling cycle, then, two dangers arise. First, the dabs of Torr-Seal may asymptotically approach complete hardness at room temperature but in fact “set” completely for the first time during bakeout, reaching new hardness and stress that persist after cooldown. Second, differential expansion of the mirror substrates, epoxy, and v-blocks may cause the components to pull away from one another and result in residual unrelieved stress after the temperature cycle is complete.

We replaced the aluminum v-blocks with identical pieces constructed from BK7 glass, the same material as the mirror substrates themselves. Kevin then spent considerable effort identifying an epoxy which would: (1) be UHV compatible, (2) set rigidly at room temperature, and (3) match thermal expansion properties as nearly as possible with BK7. A guide in these investigations was information from the telecommunications industry, where fiber and optical components must be attached in an industrial process that ensures robustness to *in situ* temperature variations over periods of years. In the end we employed Master Bond EP30LTEND for gluing of mirror substrates to v-blocks. The epoxy is still applied in two small dabs (one on either side) near the back end of the mirror (as far as possible from the mirror surface). This modified mount design caused some improvement, not in initial cavity

birefringence, but in thermal cycling properties, such that a completed physics cavity *may* now be baked gently. The baking tolerance of the overall mount is now set by length considerations and the PZT's (see Section 6.3.3) rather than by birefringence. Reducing the birefringence below the current value ( $\lesssim 1\%$  conversion of off-axis input linear polarization to circular) would involve a drastically different mounting scheme or the use of less glue. However, gluing rigidity is also a primary requirement for cavity length stability, so that further reductions in that area are not feasible for our experiments.

### 6.3.2 Specifically Chosen Cavity Length

For the first time we built a cavity with a fixed target length in mind, i.e., with the desire to have it a certain number of half-wavelengths and not one more or one fewer. Operationally this process involved a certain amount of gluing a cavity and adjusting the length based on how we thought the glue-setting process would pull it after it was too late to actively adjust any more. In other words, it was a procedure with some trial and error involved in hitting the correct length, even with the  $> 1$  FSR tuning range provided by our PZT's as discussed in Section 6.3.3. The motivation for a specific target length was to have the option of eventually including a FORT at (or near) the “magic wavelength” as discussed in [99, 19] and currently pursued experimentally in [52]. At this wavelength the AC Stark shift associated with a FORT actually shifts both the  $6S_{1/2}$  and  $6P_{3/2}$  states *down*, by approximately the same amount, so that roughly speaking a FORT is transparent to cavity QED probing under these conditions. The magic wavelength depends on FORT polarization, ranging from 923 nm for  $\sigma_-$  light to 952 nm for  $\sigma_+$ . A linear FORT – desirable because of its greater stability to polarization errors – has a magic wavelength of 935 to 936 nm, depending on which  $(F', m'_{F'})$  level we consider for the  $6P_{3/2}$  excited state. For a cavity only 10  $\mu\text{m}$  long, 935 nm is only about two FSR's away from 852; if one wants a cavity QED field and a magic-wavelength FORT to be simultaneously resonant in the cavity, cavity length *must* be set to an exact number of half-wavelengths.

Table 6.2 gives a chart of cavity lengths and corresponding resonant wavelengths for the mirrors used in our experiment. The choice of mirrors is important because mirror coatings have dispersion that causes the FSR to deviate from  $c/2l$ . Instead, successive resonant wavelengths are calculated from the detailed coating model presented in [18, 19].

|   |          |            |          |          |
|---|----------|------------|----------|----------|
| $l_{eff} = 8.4 \mu m$ ( $l_{phys} = 18(\lambda/2)$ )  | 811.2 nm | 852.359 nm | 897.8 nm | 945.4 nm |
| $l_{eff} = 8.8 \mu m$ ( $l_{phys} = 19(\lambda/2)$ )  | 813.1 nm | 852.359 nm | 895.5 nm | 941.1 nm |
| $l_{eff} = 9.2 \mu m$ ( $l_{phys} = 20(\lambda/2)$ )  | 814.8 nm | 852.359 nm | 893.5 nm | 937.1 nm |
| $l_{eff} = 9.6 \mu m$ ( $l_{phys} = 21(\lambda/2)$ )  | 816.4 nm | 852.359 nm | 891.6 nm | 933.3 nm |
| $l_{eff} = 10.1 \mu m$ ( $l_{phys} = 22(\lambda/2)$ ) | 817.8 nm | 852.359 nm | 889.9 nm | 929.8 nm |

Table 6.2: Resonant wavelengths for short physics cavities (T95 mirrors). Physical cavity length is given in units of  $\lambda/2$  where  $\lambda = 852.359$  nm is the wavelength of the Cesium transition used for cavity QED.

Thus for a linear FORT in a  $\approx 10 \mu m$  cavity, geometry forces us to choose between wavelengths of 933 nm and 937 nm. Though neither wavelength gives us exactly equal AC Stark shifts for ground and excited states, both still allow negative AC Stark shifts for the two states. We see from Figure 3.2 that the  $6P_{3/2}$  state participates in transitions to higher-lying states at 917 nm and at 921 nm, so its AC Stark shift varies more rapidly with wavelength below 935 nm than above it. For this reason we choose to target a FORT at 937 nm, giving us a design cavity length of  $l_{eff} = 9.2 \mu m$  (set off by double bars in the table). The distinction between  $l_{eff}$  and  $l_{phys}$  need not concern us until Chapter 7;  $l_{eff}$  is the length we measure in the lab while aligning and building a cavity, and has been quoted as the length  $l$  for the cavities discussed so far.

Note that, if another nearby longitudinal mode is to be used for a cavity locking beam, 814.8 nm is the clear choice. Only 814.8 nm and 893.5 nm fall well within the mirror coating curve, assuring good cavity finesse and a narrow line for locking. The 893.5 nm mode must be avoided because of the  $6S_{1/2} \rightarrow 6P_{1/2}$  transition in atomic Cesium at 894 nm.

### 6.3.3 The Shear Mode PZT's

The active element in the cavity mount system for this experiment, as for the atom-cavity microscope, is a shear mode piezoelectric transducer (PZT) from Staveley Sensor. Each cavity mirror/v-block unit is glued to its own shear mode PZT, with the two PZT's oppositely directed so positive voltage on each acts to lengthen the cavity. The PZT's are made of a lead zirconate titanate material known as EBL3, and we obtain them in 1" by 1" sheets of thickness 0.020" (20 mil). In the construction of a cavity, these sheets are scored and cleaved to produce the individual pieces we use, which are approximately 1/2" long by 1/4" wide (and still 20 mil thick). Each PZT is bonded to the main mount piece with electrically conductive epoxy, spread over the entire bottom surface of the PZT and clamped down aggressively while drying since a good bond is essential for maximum mechanical stability in the system. A v-block is glued to the top surface of the PZT with Torr-seal. The PZT's are driven by applying voltages to the top surface of one or both PZT's, while they share a common ground through the cavity mounting piece.

The shear mode PZT's were a vast improvement over the previous design (cylindrical piezo within the mount) in terms of mechanical stability; the mounting system with a large cylindrical PZT exhibited a forest of mechanical resonances at only a few kHz, where we now obtain a lowest mechanical resonance of 50-60 kHz with the current design. However, there are numerous characteristics which might well be improved upon in a future design. Toward that end, I take some space here to present some advantages and disadvantages of the present design and of the piezos in particular. The chief overall design considerations are vacuum performance (including bakeability), length tuning range, and stability; as can easily be imagined, the latter two considerations are often at war with one another.

From a vacuum pressure point of view, the weak link in the mounting chain ought to be the electrically conductive epoxy, which is only rated by the manufacturer for use at  $10^{-3}$  torr. However, we have had no noticeable difficulty with this component over years of use both in this design and in the old one. Either the manufacturer's



rating is extremely conservative, or we simply benefit from the fact that we use the epoxy only in small quantities and in a setting where an extremely small surface area of it is exposed.

A bigger limitation is the relatively low baking temperature set by the piezo material itself. EBL3 is a Staveley material very similar in properties to some other common materials, e.g., PZT-5H, EC-76, and Navy Type VI. One less-than-ideal characteristic is its unusually low Curie temperature, 190°C. Because of this very low Curie temperature and our observations of stability problems after baking of mounts, we now bake these elements at no more than a few degrees above 100°C; even this may be a bit aggressive for this material. Most other EBL types have much higher Curie temperatures ( $\approx 300^\circ\text{C}$  is common), and should perhaps be considered in future cavity mounts for this reason.

The main attractive feature of shear mode PZT's, and EBL3 in particular, is the large cavity length tuning range that they provide as compared with other flat PZT's. Given that the flat, glued-down piezo paradigm is helpful for mechanical stability, the challenge is then to assure final length adjustability over at least one free spectral range so that the finished cavity can always be tuned to the Cesium D2 resonance. This corresponds to a length range of at least 426 nm to be obtained from some reasonable voltage applied to one or both piezos. For the shearing mode of displacement we are using, the directly relevant property is the "d-constant"  $d_{15}$ , which describes shear displacement between top and bottom surfaces as a function of the voltage between them. EBL3 is quoted as having  $d_{15} = 730 \cdot 10^{-12}$  m/V, which is the same order of magnitude as related materials but at the upper end of the range. Since the recommended maximum DC applied voltage for the material is 10 V/mil or 200 V, it seems marginal that even the combined effect of two piezos should give us our desired tuning range.

In practice, however, we see displacements well above this specified value; in fact, in the current physics cavity we tune over one free spectral range of 426 nm with a voltage of zero to about 250 V, suggesting  $d_{15} = 1,700 \cdot 10^{-12}$  m/V or more than twice the nominal value. Some of this tuning range may be related to hysteretic effects we

observe in the piezo, and some may in fact be due to even our gentle baking procedure. While the generous range is in many ways a pleasant surprise, the discrepancy most likely indicates some unintended alteration of the material properties which may have other, less benign consequences. One possible choice in future designs might be to resign ourselves to high voltages and obtain the requisite tuning range with a different material of lower  $d_{15}$  but larger allowed voltage.

A final consideration is the frequency of the lowest mechanical resonance in the system, and the way in which this all-important quantity is tied to the choice of PZT. Technical noise on cavity transmission signals is substantially due to electrically excited mechanical resonance of the cavity mounting structure. The piezo itself has a native resonant frequency  $\nu_0$  set by its thickness and by the material’s “transverse frequency constant,” which for EBL3 is 1105 Hz·m. This implies  $\nu_0 = 2.2$  MHz for the elements we use. For a piezo (spring) of mass  $m$  mechanically loaded with a mass  $M_L \gg m$ , the overall resonant frequency becomes  $\nu = \nu_0 \sqrt{\frac{m/3}{m/3 + M_L}}$ . Thus a loading mass of  $M_L = 1600(m/3)$  would bring the 2.2 MHz resonance down to the 50 kHz frequency we observe. In fact, however, the mirrors and v-blocks constitute a load only a few times more massive than the piezos themselves, so the predicted resonance is still well above what we observe. Thus we see that other spring constants are in action and it remains crucial to push the resonance as high and small as possible by rigidly attaching the mount elements and minimizing the height of the mirrors above the PZT’s themselves to fight tuning-fork effects in the mount. The observed importance of rigid gluing in ensuring a high resonance frequency, as mentioned earlier, confirms that the native resonance of the loaded PZT is not yet the biggest issue.

### 6.3.4 Length Variation with Magnetic Field

The current physics cavity was installed in the vacuum chamber, the chamber was closed and pumped down, MOT coils and optics were built up around the chamber, and Cesium was introduced. The first time MOT’s and cavity lock were operated simultaneously, we encountered a completely unexpected (and, needless to say, un-

welcome) variation of the cavity length with applied magnetic field. This effect had not been seen in previous cavities. The lack of a corresponding voltage jump on the PZT's eliminated the hypothesis of PZT shearing in response to the field. Instead the effect suggested the presence of a magnetic material in the cavity mount or somewhere nearby where the mechanical stress could be transmitted to the cavity.

After considerable hunting for the source of the effect, I picked up a spare cavity mounting block, applied an external magnetic field, and measured with a Gaussmeter an unmistakable induced field from the mounting piece itself. The explanation became simple. Hoping to ensure good vacuum pressure in the cavity region, and concerned about oxidation of the copper mount over the several days between cleaning of the piece and insertion of the finished cavity in the vacuum chamber, we had taken the protective measure of sending the copper mount pieces to be *gold-plated* before cavity construction. Unfortunately we were completely ignorant of the standard gold-plating process for copper, which involves application of a thin coating of *nickel* between the copper substrate and the gold top layer [104]. Thus our cavity mount piece itself is magnetic; applied fields induce magnetization in the nickel layer, force on the mount, and a jump in length of the cavity itself.

With this fact established, future mount construction can be altered in one tiny way; gold-plating can be specified without the nickel-plating step in the process. The chief purpose of the nickel layer is to prevent gradual diffusion of the gold into the copper substrate, an effect which is completely unimportant in our application. However, with our cavity already assembled, in UHV, and with Cesium and working MOT's around it, we chose to proceed with it in our experiment and compensate for the length jump rather than starting over with a new cycle of building up cavity, chamber, and optics.

The measured effect consists of a cavity length jump of  $4 \text{ \AA}$  when the lower MOT main anti-Helmholtz coils are turned on. When the coils are switched off, the cavity returns to its original length with remarkable accuracy, at the level of a linewidth. Nevertheless, this length jump is as rapid as the field switching and is huge compared to a linewidth ( $\approx 10 \text{ fm}$ ); it requires an applied voltage of 250 mV to the PZT to

compensate for the length change. The cavity lock, with gain designed to deal with small-amplitude vibrations and integration to deal with slow drifts, cannot possibly follow this jump during lower MOT loading or recover the lock quickly after the MOT is turned off. Thus the cavity locking circuit was modified to include a magnetic field compensator which directly samples the current in the main lower MOT coils and feeds that signal forward to send a corresponding additional voltage to the PZT. With this modification, the cavity length lock is still lost transiently at the switching of the MOT coils, but recovers fully in less than 5 ms after switching. Since atoms do not arrive in the cavity region until much later (after sub-Doppler cooling and a 3-mm drop), the cavity is amply prepared for atom-trapping by the time transits occur.

A second concern, however, is the lingering presence of the induced magnetic field of the mount itself and how it might alter the atom-cavity interaction by shifting the atomic levels. I have put some bounds on this effect by using the spare cavity mount to measure the size of the induced field at the cavity position. In the experiment, the top face of the mount sits at the center of the antiHelmholtz coils (in three dimensions) and extends  $\pm 1.6$  cm along the main coil axis. Given a field gradient of 16 G/cm along the main axis, we can use an applied field of 25 G to approximate the situation. Using a Gaussmeter and a permanent magnet, I applied this field to the cavity mount and subsequently measured an induced field of 0.1 Gauss (vertical) at the position of the cavity mode. This induced field would cause precession between  $m_F$  levels as otherwise defined by the circularly polarized cavity probe.

For magnetic fields small enough to be perturbative relative to hyperfine splittings, we are in the regime of the anomalous Zeeman effect, and levels are shifted by  $\Delta E_{|F m_F\rangle} = \mu_B g_F m_F B$ . Here  $\mu_B = h \cdot 1.40$  MHz/G is the Bohr magneton, and  $g_F$  is the hyperfine Landé  $g$ -factor. For  $6S_{1/2}$  and  $6P_{3/2}$  we have  $g_F = \frac{1}{4}$  and  $\frac{2}{5}$ , respectively [74]. Thus we estimate Zeeman shifts of  $\approx m_F \cdot 0.035$  MHz for the ground state and  $\approx m'_F \cdot 0.056$  MHz for the excited state. These Zeeman shifts between  $m_F$  levels as defined by the vertical magnetic field correspond to precession between  $m_F$  levels in the orthogonal basis defined by the circularly polarized cavity QED probe. The rate

of precession can be estimated as  $\approx 0.05 \text{ MHz} \cdot m_F$ , small compared with the scattering rate between field and probe with one photon in the cavity, but not vanishingly small and certainly disturbing from the point of view of future FORT experiments in this cavity.

On the other hand, the measurement described above is certainly an overestimate of the mount-induced field at the cavity position. In the actual experiment the mount is not magnetized by a constant field, but by a gradient field with its zero at the mount center. Thus two mitigating factors should make the induced field significantly *less* than 0.1 Gauss at the cavity position. First, the two halves of the mount are oppositely magnetized (at least along the primary axis) by the gradient field. Second, the portions of the mount closest to the cavity are also those which see the smallest applied field. An improved diagnostic for the induced field would be to apply the actual quadrupole field to the test mount, turn off external fields, and measure the induced field at the cavity position in this setup.

## 6.4 Into the Digital Era

While the overall timing of the experiment is still controlled from Labview as in the atom-cavity microscope, data acquisition and triggering/feedback have been changed. The spectrum analyzer video output signal is sent to a Gage (CompuScope 1450) board for digitization and storage of the transmission data. The Gage board can acquire up to 50 megasamples per second, with 14-bit resolution. We digitize at only 1 MHz and acquire data during a 20 ms interval for each MOT drop. Triggering and feedback are now accomplished through the digital processing capabilities of an FPGA board. The motivations for this change lie in the implementation of active feedback algorithms, due both to near-term technical issues and to more basic properties of the servo loop, as discussed in the introduction to this chapter. We use a VirtexE DSP hardware accelerator board (GVA-290a), which incorporates four dedicated A/D and D/A channels as well as three Xilinx FPGA's (two fully configurable). The A/D and D/A conversion has a depth of 12 bits and a maximum rate of 100 megasamples per

second; the FPGA clock rate is 100 MHz.

The programming of a very similar device for use in a quantum optics experiment is discussed in detail in [97]. FPGA's are an attractive choice for inherently digital processing, and additionally excel at rapid digitization, lookup tables, and simple but fast on-chip logic producing output pulses conditioned on the input data stream. The incorporation of the FPGA in our experiment has been largely undertaken by Kevin Birnbaum. The function of the board in our experiment is to (a) digitize the signal from the spectrum analyzer video out, (b) convert it to units of intracavity  $|\langle a \rangle|^2$ , and (c) implement the triggering and feedback scheme of Chapter 5, whose chief computational elements are a lookup table converting  $|\langle a \rangle|^2$  to  $\rho$  and an FIR filter for estimating turning points of  $\rho$  (zeros of  $\dot{\rho}$ ).