Chapter 6

A new generation of cavities

For several years, the workhorse of cavity QED experiments in our research group has been the optical cavity in lab 11. This cavity was built in 1999 and since then has remained in an unbroken ultra-high vacuum (UHV) environment, an exceptional stroke of good luck and a testament to the remarkable experimental skill of David Vernooy and Jun Ye. More recent cavity-building efforts in our group have had multiple goals: to construct cavities with different characteristics than those of the lab 11 resonator, e.g. shorter or asymmetric; to improve upon the design and assembly procedure for the cavity and vacuum chamber system; and to keep knowledge about the process of cavity building — something of a black art — alive in our group.

In this chapter, I discuss our group's two most recent cavity projects. While the lab 1 feedback project never produced any publishable experimental results, we learned a great deal in our effort to get it running; this information has not been detailed elsewhere and should be of use in the design of new experiments. Meanwhile, the future of the asymmetric cavity project is still an open question. In the second part of the chapter, I summarize progress to date in this most recent endeavor and attempt to highlight outstanding questions and challenges.

6.1 Active feedback to atomic motion

Chapter 6 of Theresa Lynn's thesis provides thorough documentation about the design of the feedback experiment [22], so I will only touch on this briefly before discussing

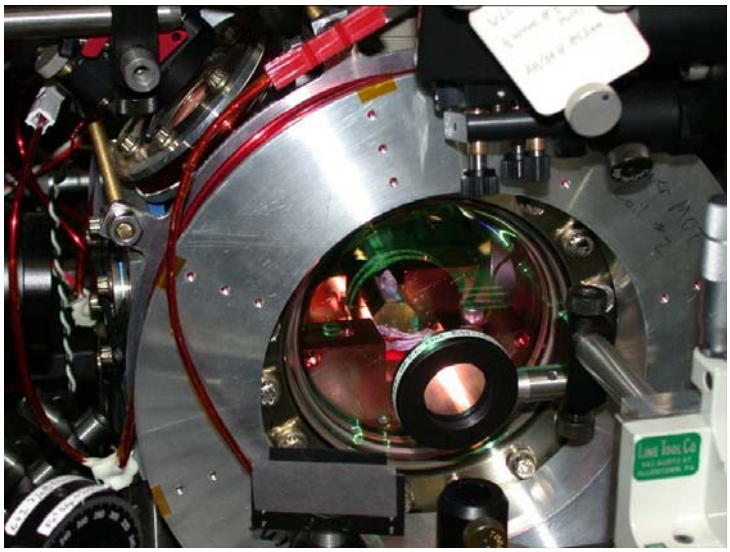

Figure 6.1: The lab 1 cavity inside the lower vacuum chamber

its implementation. In 2000, Theresa and Kevin Birnbaum began to assemble a two-chamber UHV system similar to the one already in place in lab 11. The Ti:sapphire laser previously used both for probing and locking the cavity was replaced with diode lasers at 852 nm (probe) and 815 nm (lock), both stabilized to the cesium $D2$ line via an external reference cavity. The physics cavity inside the new vacuum chamber was the first one ever built to a specifically chosen length (in this case, 9.2 μm), and special care was taken to select the cavity mount materials in order to minimize birefringence. Another important reference for experiment design is Appendix B of Kevin Birnbaum's thesis, a collection of information about vacuum systems compiled during the construction of the new chamber [4]. A separate, unpublished set of notes by Kevin describes the cavity-building procedure used in December 2001. (Yat Shan Au's undergraduate thesis includes the most recent updates to this procedure [105].)

6.1.1 Characterizing the active-feedback cavity

In June 2002, when I joined Kevin, Theresa, and Dominik Schrader in lab 1, the new cavity had already been placed inside the lower vacuum chamber, which had just been pumped down to 10^{-10} torr. We worked through the summer on optimizing the

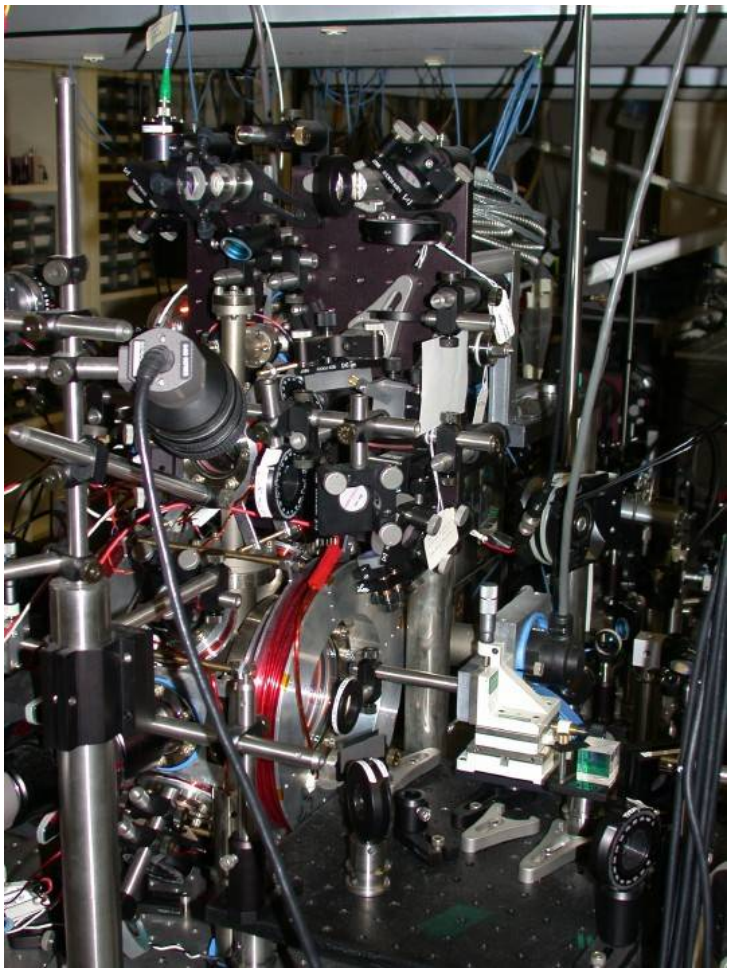

Figure 6.2: The lower and upper vacuum chamber of the lab 1 experiment, surrounded by associated optics

MOTs in the upper and lower chambers and on positioning the lower MOT above the mirrors so that atoms would fall through the cavity. As we began to implement active stabilization of the cavity length, we discovered that the cavity mount had been inadvertently plated with nickel, a ferromagnet, when it had been sent out for gold plating. Each time we switched on or off the magnetic fields used for trapping atoms, the cavity length jumped by roughly 4 Å. In order to compensate for these jumps, we built a feed-forward input to our cavity length servo which anticipated the magnetic field switching. With the new servo in place, locking the cavity in transmission to the probe laser field, we observed our first atom transits in October 2002.

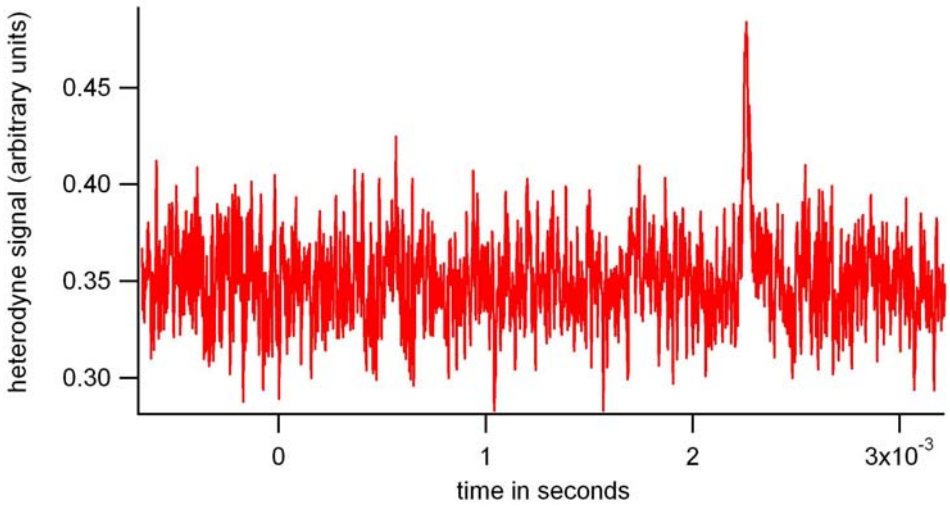

Figure 6.3: A sample "upgoer" atom transit observed in lab 1 via heterodyne detection. Here the probe laser is tuned red of the atom and cavity resonance such that $\omega_p = g = -130$ MHz. The transit is observed about 23 ms after release of the atom cloud and has width ~ 50 μs.

Much of the following year was spent assembling the frequency metrology setup, analogous to the one used in lab 11. This rather complicated stabilization chain required us to lock our probe laser to a reference cavity, then lock the reference cavity to the cesium $D2$ line, lock our locking laser to the reference cavity, and finally lock our physics cavity to a frequency sideband of the locking laser. It was June of 2003 before we could accomplish all of these steps simultaneously and several more months before we could do this reliably.

Once we had observed transits, both "downgoers" and "upgoers," while locking to the 815 nm sideband, our next projects were to improve their frequency and to characterize them. Kevin and I were not particularly successful in this first task, despite our repeated efforts to optimize the size, position, and temperature of the MOT, and even to tilt the optical table. Subsequently, Kevin wrote a simulation to estimate the number of atom transits per MOT drop as a function of cavity angle, included as section 7.9 of Ref. [4]. He found that the angular dependence of transit probability was fairly weak: a tilt of 4° from horizontal was necessary to reduce the number of observed transits by half. (It is interesting to note that in another cavity

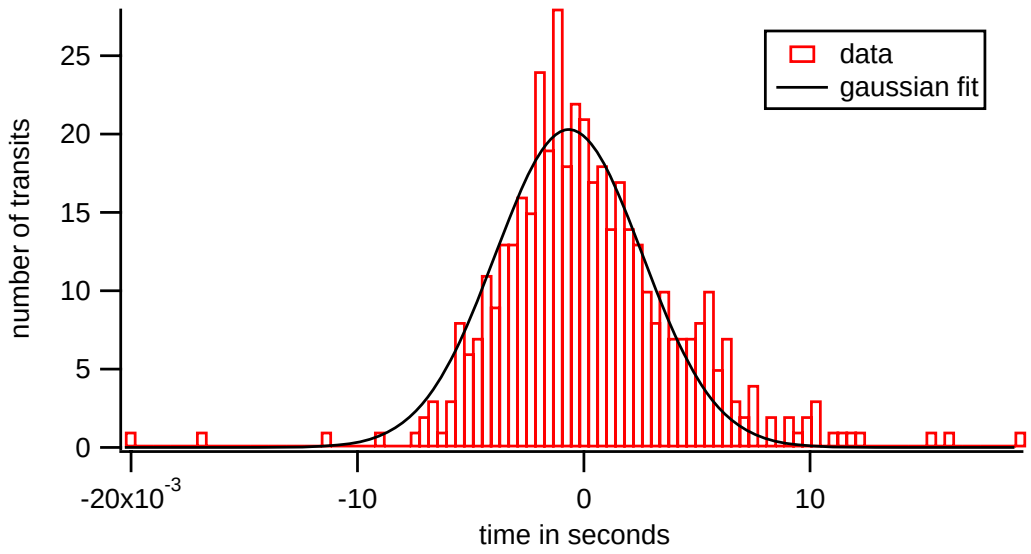

Figure 6.4: Histogram of 442 recorded transit arrival times. The trigger at time $t = 0$ occurs 25 ms after the atoms are released from polarization-gradient cooling.

QED experiment, tilting the optical table has in fact been used to optimize the atom count rate through the cavity [106].)

One step in characterizing the transits was to assemble a histogram of transit arrival times, based on 442 transit files recorded over 7 hours in November 2003. This data is shown in Figure 6.4, where the transit arrival times are referenced to an external trigger from the LabView control program, 25 ms after the end of polarization-gradient cooling. A Gaussian fit to the data gives a mean arrival time of -0.8 ms and a distribution half-width at half-maximum (HWHM) of 4.6 ms. Note, however, that the Gaussian fit is not entirely satisfactory: it overestimates the number of atoms arriving early and underestimates those arriving late.

To investigate the observed distribution — both its asymmetry and its implications for the temperature of the atom cloud — we returned to Kevin's transit simulation program. The program takes input variables which include the standard deviations of initial x and y positions (σ_x, σ_y) and velocities (v_x, v_y) of atoms in the MOT as well as the height of the MOT above the cavity $(\langle y \rangle)$. It then samples atoms from this distribution and records the arrival time of each atom within the cavity. (This is

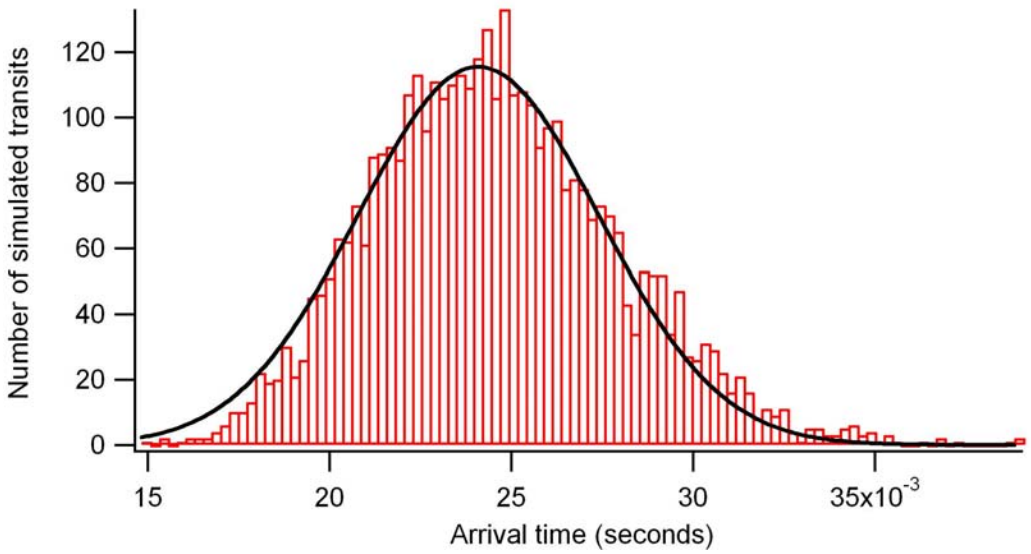

Figure 6.5: Simulated distribution of transit arrival times. Atoms are released at $t = 0$ from a cloud above the cavity. The height and size of the cloud are determined from images of the MOT, while the initial velocities of the atoms are adjusted for good agreement with the data.

simply a calculation of moving particles under gravity, constrained by the mirror geometry, and does not take into account any light forces on the atoms.) We determined σ_x, σ_y, and $\langle y \rangle$ independently from calibrated frame-grabber images of our MOT. We then adjusted input parameters v_x and v_y, assuming them to be equal, in order to obtain good agreement with the measured distribution width. Figure 6.5 shows a simulated histogram with $\sigma_x = \sigma_y = 0.174$ mm, $\langle y \rangle = 2.9$ mm above the center of the cavity, and $v_x = v_y = 3.13$ cm/s, for a distribution HWHM of 4.7 ms. From these velocities, we were able to infer a lower MOT temperature of 16 μK, consistent with effective polarization gradient cooling. Moreover, it is interesting that the simulation reproduces our experimental asymmetry, with more atoms piled up at the slow tail.

We also wanted to understand the observed time for an atom to transit the cavity, which was faster than we had initially expected. For example, in November 2003, we measured transits with full-width at half-maximum (FWHM) between 20 and 36 μs, corresponding to a velocity of 0.6 – 1 m/s across the cavity mode. However, an atom falling from rest 3 mm above the cavity (the measured location of the MOT)

would only have $v = 0.25$ m/s. The answer was that this classical picture neglects the trapping force provided by the atom-cavity coupling g, which in fact accelerates the atoms as they approach the center of the mode. A calculation using the cavity parameters $\omega_{atom} = 0, \omega_{cavity} = 0, \omega_{probe} = -130$ MHz, $g_0 = 130$ MHz, $\kappa = 17$ MHz, and empty cavity photon number $\langle a \rangle = 1$ found that an atom falling 3 mm would cross an antinode of the cavity mode with $t_{FWHM} = 29$ μs. As it turns out, there is only a very weak relationship between the initial height of an atom and its transit time, which is almost entirely determined by g. For example, an atom falling from 2 cm would have a transit time $t_{FWHM} = 22$ μs; for an atom with no initial fall velocity, $t_{FWHM} = 31$ μs. An apparent discrepancy with the 74 μs average transit time reported in Ref. [20] was resolved when Theresa explained that their criterion had been the amount of time that atoms were within one mode waist of the cavity center, rather than the FWHM.

A third transit measurement looked at the difference between peak and empty cavity transmission for "upgoers" as a function of probe intensity (Figure 6.6). The results were roughly consistent with a Matlab calculation using the quantum optics toolbox and thus confirmed that our heterodyne calibration (which we used to scale cavity transmission to photon number) was about right.

Finally, we were concerned about the polarization of the intracavity light in connection with observations that a repump beam from $F = 3$ to $F' = 4$ improved the frequency of transits. (We intended to drive the cavity with circularly polarized light on the closed $F = 4 \rightarrow F' = 5$ transition. The fact that atoms were present in the $F = 3$ manifold seemed to suggest the presence of linearly polarized light, though we later concluded that the repump beam probably just helped us see atoms that were in $F = 3$ as they entered the cavity.) After reworking the cavity input path so that the last optical elements were a polarizing beamsplitter cube (PBS) optimized for 850 nm and a zero-order quarter waveplate at 850 nm, we determined that only 0.003% of the input light had the wrong circular polarization. In December 2003, we measured $\kappa_1 = 14.4 \pm 0.6$ MHz along one birefringent axis of the cavity and $\kappa_2 = 17.0 \pm 0.8$ MHz along the other, where the difference in the center frequency between the two

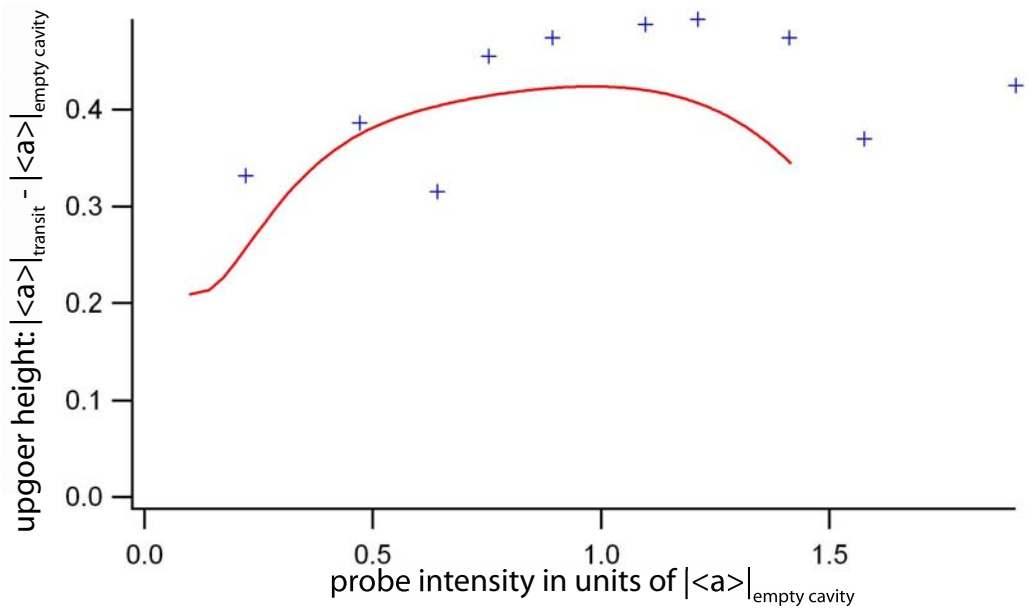

Figure 6.6: Upgoer height as a function of probe intensity, expressed in units of $|\langle a \rangle|$ (proportional to the square root of the heterodyne detection signal) for comparison with a Matlab quantum optics toolbox calculation (red). As the empty cavity value of $|\langle a \rangle|$ also increases with probe intensity, this is subtracted from each data point.

resonances (i.e., the birefringent splitting) was $\Delta\nu = 4$ MHz. Thus $\frac{\Delta\nu}{\kappa} \approx 0.25$, in contrast with $\frac{\Delta\nu}{\kappa} \approx 1$ for the lab 11 cavity [29]. As $\frac{\Delta\nu}{\kappa}$ is proportional to the birefringent phase shift per round trip, we can conclude that the efforts made by Theresa and Kevin in constructing the new cavity had succeeded in reducing birefringence effects by a factor of four.

In the spring of 2004, we assembled a diode laser at 936.8 nm, two cavity free spectral ranges red of the cesium resonance, with the intention of trapping atoms within an intracavity FORT. Because of the narrow gap between the cavity mirrors, we would not be able to focus lattice beams through the cavity from the side in order to cool and load atoms into the trap. However, with an EOM providing frequency sidebands for locking the laser to the physics cavity and an AOM for switching the trap, we planned to trigger the FORT on in the presence of an atom, as had been done in the early days of lab 11 [40]. The FORT linewidth and other cavity parameters are listed in Table 6.1.

lab 1 cavity parameter	value
effective length l	9.2 μm
mode waist w_0	13.3 μm
κ_1, κ_2	$2\pi\times$(14.4 MHz, 17 MHz)
finesse	570,000 (along axis 1)
g_0	$2\pi \times 130$ MHz
$\kappa_{815\text{nm}}$	31.2 MHz
$\kappa_{936\text{nm}}$	4.9 GHz

Table 6.1: Parameters for the lab 1 active-feedback cavity. All parameters are for 852 nm unless otherwise stated.

6.1.2 Difficulties in implementation

We encountered a number of difficulties which slowed our progress on the active-feedback experiment and would have made data acquisition a challenge. Here I highlight what seem in retrospect to have been the most significant of these problems and suggest how they might be avoided in future experiments.

6.1.2.1 A nickel-plated cavity mount

The nickel-plated cavity mount described in Section 6.1.1 has received a great deal of blame. The copper mount was gold-plated through the machine shop in order to prevent oxidation during cavity assembly, after which it could no longer be baked. Unknown to us, nickel is generally used as an underplate in this process to avoid gold diffusion into the copper, and we should have specifically requested its omission. While it was frustrating to learn that magnetic material had found its way into the vacuum chamber, in practice we were able fairly quickly to devise a feed-forward locking circuit that compensated for the cavity stabilization difficulties caused by induced fields. Although the cavity length jumped each time the current in the MOT coils switched, a signal sent to the locking circuit provided a synchronous offset voltage to the cavity piezo.

However, this adaptation made the locking circuit inherently less stable. In addition, the induced magnetization in the nickel would make it challenging to implement controlled transverse and axial bias fields in future experiments. This sort of con-

trol over the magnetic field seen by atoms trapped in the cavity has emerged as an important tool in lab 11.

6.1.2.2 Complexity of the stabilization network

The network of lasers and servos required for the experiment to run was fairly complex. The MOT trapping and repump lasers had to be locked separately to cesium lines. The cavity probe laser and 815 nm locking laser were each locked to a 30 cm, 50 kHz reference cavity, and the reference cavity was then locked to cesium. A traveling wave modulator was used to put frequency sidebands (tunable up to 1 GHz) on the locking laser, and the physics cavity length was then stabilized by locking to one of these sidebands. Later, the 936 nm FORT laser was also locked to the physics cavity. All of these locks used the Pound-Drever-Hall technique, and all five lasers were home-built external-cavity diode lasers.

This complexity was to some extent inevitable, but it meant that we spent a great deal of time just trying to keep everything locked simultaneously, and the instability of the cavity lock due to the feed-forward circuit described above only made things worse. We learned to avoid unwanted amplitude modulation (which causes DC offset drift in an error signal) by using EOMs to put frequency sidebands on the laser light instead of modulating the laser current; in addition, it was helpful to put a PBS cube before each EOM for a clean input polarization. At first, we locked the physics cavity in reflection but then switched to locking in transmission, which allowed us to avoid noise on the reflected signal.

In this situation, we could have learned from lab 11, where Andreea Boca implemented a robust injection-locked laser setup for the MOT trapping laser [29]. By replacing the grating-stabilized $F = 4 \rightarrow F = 5$ laser with a slave laser seeded by injection light from the cavity probe, we could have eliminated one cesium lock from our system. We could also then have bypassed our SDL tapered amplifier, which came with its own alignment headaches, since an injection-locked laser can provide a few hundred mW of power. (Using a grating to form an external cavity reduces the laser output power by roughly half, since the rest of the power is fed back into the

diode.) More recently in lab 11, we replaced the cesium lock for the $F = 3$ repump laser with a phase lock to the probe laser, as described in Chapter 3. This was done for the purposes of having phase-stable $F = 3$ and $F = 4$ lasers for the coherent state transfer experiment; however, we have found as a side benefit that the lock has been much more robust.

6.1.2.3 A dearth of transits

This final problem is perhaps the most fundamental: we simply never saw very many atoms in our cavity. During successful data runs, it was more likely to see no transits than one transit in a four-second experimental cycle, and it was uncommon to see more than one transit per cycle. When lasers drifted out of alignment, or the MOTs were not well positioned, it was difficult to see any transits at all.

Perhaps this should have been more cause for concern beforehand, given the short length of the cavity: Theresa notes in her thesis that on average 40 transits per cycle were observed in lab 11, and that the flux of atoms through the lab 1 cavity was expected to be 60 times smaller than in lab 11. This had not been an issue in the previous atom-cavity microscope experiment, also based upon a $\sim$ 10 μm cavity, because that experiment had used only one MOT, formed directly from cesium background vapor above the cavity. The cost of switching to a double-MOT design is a reduced atom number in the final MOT, by roughly a factor of 100.

The current focus in our group has been on cavities with enough space to admit lattice beams from the side as in lab 11, that is, at least 30 μm in length. However, if there is future interest in a very short cavity with large g, I would recommend developing a new strategy for obtaining a large lower MOT. For example, one might incorporate a "pusher beam" into the upper MOT in order to create an "atom faucet" [107] that would load the lower MOT with a high-flux, low-divergence atomic beam.

6.1.3 Technical improvements

There are also some positive lessons to be drawn from the lab 1 experiment. The design of the two-chamber vacuum system offers much more physical and optical access than is available in lab 11. The technique of coupling both the trapping and repump beams for the MOT into polarization-maintaining optical fiber, then bringing them out of fiber in a compact setup close to the chamber, proved very successful. I would like to emphasize the usefulness of having a systematic procedure in place for MOT alignment, that is, the ability to center all three transmitted MOT beams by eye on their output windows and the use of irises for retro-reflection alignment.

Two optimization procedures for the MOT are worth mentioning. The first concerns the positioning of both the upper and lower MOTs, in the first case above the differential pumping tube, and in the second case above the cavity. The main MOT coils provide a gradient trapping field for the MOT, while three pairs of bias coils ensure that the magnetic field is nulled at the MOT location. If this condition is not met, then when the main coils are turned off for polarization-gradient cooling, the atoms will be pulled away from the center of the trapping beams by the residual fields and will not fall correctly into the second chamber or through the cavity mode. In order to optimize the currents through the bias coils, we used a Labview control program which, after forming a MOT, gradually stepped down the current through the main coils over the course of a few seconds. This allowed us enough time to observe the trajectory of the atom cloud on a camera as the MOT dissipated. We could then adjust the bias coil currents until the cloud expansion appeared uniform and symmetric.

The second procedure measured the temperature of the cesium cloud after cooling. In January 2004, as part of an effort to improve the frequency of transits, we decided to quantify the polarization-gradient cooling, which consisted of 35 ms of illumination from the trapping beams, at a lower intensity and increased detuning from their MOT settings and with the MOT coils off. We used the following method: after waiting a variable time delay t after the fields were turned off, we applied a detuned pulse of light

a few ms in length in order to illuminate the atom cloud. For each time t, we took a pair of pictures of the illuminated cloud with our camera frame-grabber. One picture of each pair was taken with the MOT coils always turned off (i.e., in the absence of atoms) and then subtracted from the other in order to remove background light from the chamber. After obtaining image pairs at a series of delay times, we analyzed the data in IgorPro: following the image subtraction, we took a line profile through the center of the remaining cloud, which we fit to a Gaussian. Plotting the Gaussian widths as a function of t, we fit the data to determine both the initial radius of the cloud and the expansion time constant τ, from which we could extract a temperature. In practice, once we had improved the cooling by adjusting the intensity and detuning settings of the beams, it was no longer possible to measure a temperature accurately because the cloud fell under gravity out of the range of the camera before it had much chance to expand. In this case, we could make further improvements to the upper MOT settings by instead observing the number of atoms collected in the resulting lower MOT.

One interesting consequence of our difficulties in observing transits was that Kevin began to wonder if there was some underlying physics at work that we hadn't considered. He realized that in the case of such a small cavity, the atom-cavity coupling of $2\pi \times 130$ MHz was on the order of the cesium hyperfine excited state splittings. This brought into question the assumption that we could model our system as a two-level atom interacting with a single cavity mode (the Jaynes-Cummings model [3]); even though our circularly polarized probe beam addressed a two-level cycling transition between $F = 4$ and $F = 5$, the cavity could now couple to other hyperfine levels. Working with Scott Parkins, who was visiting from the University of Auckland, Kevin calculated the transmission as a function of probe and cavity detunings for a cavity coupled to multiple excited state levels, as well as to the entire $D2$ transition [28, 5]. While our optical cavity was just on the border where these couplings become relevant, it may be possible to observe some of the phenomena predicted by his calculations in the microtoroid experiments currently underway in our group.

6.2 An asymmetric cavity system

In the fall of 2004, we abandoned work on the active feedback experiment and began instead to construct a new cavity. This next-generation cavity would address several shortcomings of the current lab 11 system. While we hoped that it could someday run in parallel with the lab 11 experiment, the system also had a practical function as a backup cavity for lab 11, which at that point had been running continuously for six years (now nine).

In particular, the new cavity would be "single-sided," that is, built with one "open" mirror more transmissive than the other ("closed") one, so that light escaping from the cavity would exit primarily through this open mirror. At a minimum, this design would allow us to collect roughly twice the intracavity photons as in lab 11, where photons exit from both cavity mirrors but are only collected at one port. We could also hope to realize protocols for quantum computation, such as a quantum nondemolition (QND) measurement of photon number [44]. (This protocol relies on a single-sided cavity because it requires a photon either to reflect off the cavity — acquiring a phase shift in the process — or to enter the cavity and then exit through the same mirror it entered.) The most ambitious vision was that with a second functioning optical cavity in our group, we might entangle the states of two atoms trapped simultaneously in both cavities [78].

Previous cavity-building efforts had attempted to minimize birefringent splitting in the cavity modes by avoiding any nonuniform stress to the mirrors. Techniques for achieving this included minimizing the use of epoxy which bonds the mirrors to their v-blocks, keeping the epoxy as far as possible from the mirror surfaces, selecting an epoxy with optimal thermal expansion properties, machining the v-blocks out of BK7 in order to match the thermal expansion coefficient of the mirror substrates, and not baking the vacuum chamber once the cavity was inside [21, 22]. As we have seen, in the active-feedback cavity, the result was a reduced but still measurable birefringence. We hoped in this case to explore stressing the cavity mirrors on purpose in order to induce a birefringent splitting. Ideally, we could then tune this stress, perhaps with a

piezoelectric transducer, either to (a) null it entirely, or (b) make it large enough (i.e., $\Delta\nu \gg g$) so that our intracavity atom would only interact with one cavity mode.

Finally, we were not sure what was limiting the recently demonstrated 2–3 s trap lifetimes in the lab 11 cavity, but a prime candidate was background collisions [102]. In the new vacuum chamber, we would aim for a background pressure of 10^{-12} torr, two orders of magnitude lower than in previous experiments, in the hope that this would permit longer intracavity storage times. However, this new target pressure suggested that we would have to bake the vacuum chamber more aggressively than in the past. It would be necessary to reconsider how well the chamber components could withstand baking.

Kevin had meanwhile joined Andreea, David Boozer, and Russ Miller on the vacuum Rabi splitting experiment in lab 11. I took on the project of collaborating with mirror manufacturers to develop low-loss mirrors for our cavity, the subject of Chapter 7. Meanwhile, I worked with undergraduates Yat Shan Au and Travis Bannerman on vacuum chamber and cavity redesign; the project later grew to include Toby Burrows, Andrey Rodionov, and Dal Wilson.

6.2.1 Cavity redesign

At the heart of our cavity-QED experiments is a high-finesse Fabry-Perot resonator consisting of two mirrors fabricated on BK7 substrates. As described in more detail in Refs. [21, 22], in the most recent generation of cavities, the mirrors are attached with epoxy to BK7 v-blocks which have been glued on top of shear-mode piezoelectric transducers. Another layer of electrically conductive epoxy bonds the piezos to the cavity mount, a copper block.

This cavity design presented two problems in light of our plan to bake the vacuum chamber: the low Curie temperature (190 °C) of the piezo material, EBL3, and the temperature ranges of Torr-Seal and Master Bond EP30LTE-ND, the epoxies used for bonding the v-blocks to the piezos and the mirrors to the v-blocks. After extensive research, Yat was able to suggest replacement materials. Pz23 and Pz27 from Ferrop-

erm Piezoceramics (Kvistgaard, Denmark) have relatively high Curie temperatures of 350 °C and a recommended maximum working temperature of 250–300 °C. For a UHV epoxy, Caburn MDC (West Sussex, U.K.) sells a product called H27DUHV which it claims is vacuum-compatible to 10^{-11} torr and can be baked to 270 °C. Dal returned to the question of epoxy last year in the process of building several test cavities, where in addition to UHV compatibility and the coefficient of thermal expansion, he considered the importance of a good consistency and reasonable curing time for cavity-building. He has settled on Epotek H21D as a cheaper alternative to the Caburn product; in addition, its volatile organic outgassing specifications are well-documented. (After a cavity baking experiment in which the finesse dropped by a factor of two, we are now concerned about outgassing not only because of its effects on background pressure but also because it may deposit material on the mirror faces.) It is interesting to note that he has also successfully used thermo-compression to achieve a glue-free bond between two flat surfaces, for example, by clamping together a silver-evaporated piezo and a copper cavity mount for an hour at 150 °C.

Would the new high-temperature piezos allow us to tune the cavity length over a full free spectral range (426 nm of travel)? This is a difficult question to answer, as the shear mode piezos in both the previous lab 1 cavity and the current lab 11 cavity exhibit two to three times as much travel as one would expect [40, 22]. Here the relevant quantity is the shear displacement d_{15} of one piezo surface relative to the other, which is 335 pm/V for Pz23 and 500 pm/V for Pz27. In comparison, the nominal value of d_{15} for EBL3 is 730 pm/V, but measurements in lab 1 indicate $d_{15} = 1.7$ nm/V; one free spectral range corresponds to an applied potential of about 250 V. Initial interferometer measurements in the lab were consistent with specified d_{15} values; a test cavity which Dal later built with Pz27 in October 2006 was found to travel 425 nm over 400 V, or $d_{15} \approx 1$ nm/V. It remains to be seen whether the Ferroperm products will display a similar behavior in a cavity under vacuum. As for the response of the piezos to temperature stress, Yat tested Pz23 after baking for 48 hours at 198 °C and observed no change in maximum travel. Now that we have an oven capable of higher temperatures, it would be interesting to revisit this question.

Given the uncertainty about the range of the piezos, for several months we discussed a more significant redesign of the cavity mount. The impressive passive drift of < 10 fm/s reported by Michael Chapman's group at Georgia Tech [47, 108] prompted us to consider a single expansion-mode piezo design, which they believed was the source of their stability [109]. Here as well we were confronted by the short travel range of the piezo for our small cavity as well as questions about stress induced by gluing the mirrors directly to the piezo. Yat then designed a mount in which an expansion-mode piezo drove a spring system formed by electrical discharge machining of an aluminum piece; the cavity mirrors were glued directly to the metal on either side of the spring. The design had several appealing features, including a large range of travel, mechanical stability because the mirrors shared a common mount, and the possibility of avoiding glue altogether. However, in initial tests, piezo expansion was found to tilt the mirror alignment along with changing the length of the cavity, probably due to nonuniform spring compression [105]. Finally, we have considered the option of giving up long-range piezo tuning altogether and instead relying on slow temperature tuning, either with a laser or a heating element inside the vacuum chamber. But as we have not yet demonstrated a satisfactory replacement, we have decided for the time being to rely on the previous cavity mount design, with the simpler substitution of high-temperature-compatible components.

6.2.2 Controlling birefringence

The challenge of controlling the cavity birefringence is of course connected to cavity redesign, since any mechanism to actively induce stress in the mirrors would have to be incorporated into the cavity mount, along with the usual considerations of UHV compatibility and optical access for the MOT and cavity beams. Toby spent the summer of 2005 engaged in two projects: using finite element analysis to model how stress on the mirror substrates leads to strain and a change in the mirror's index of refraction, and trying to induce and quantify birefringence with test cavities in the lab using jigs that I had designed. Unfortunately, due to the short length of his stay,

neither project was conclusive.

Yat and Dal then took on this project in the following year and spent some time characterizing the birefringence of test cavities. By first securing a mirror with a wire in order to pre-load it, then applying a voltage to a piezo wedged between the mirror and the wire, they successfully observed an induced birefringence. More sophisticated mounts for inducing stress were then employed, one in which a mirror was pre-loaded with set screws against two piezos which applied pressure from orthogonal directions, and a second (a possible candidate for use in UHV) in which the usual cavity mount design was adapted by placing an expansion-mode piezo inside one of the v-blocks. In all cases, the observed change in the birefringent splitting between the two orthogonal cavity modes was on the order of $\frac{\Delta\nu}{FHWM} = \frac{\Delta\nu}{2\kappa} = 1$. This is in rough agreement with what we would expect from a finite element analysis model that Yat developed [105].

While these results are encouraging, at this time we still have no reliable means of controlling the birefringent splitting of our cavities. Moreover, the magnitude of the induced birefringence suggests that while we might (with a clever design) be able to eliminate the intrinsic birefringence of a cavity, we would have to increase the stress on the mirror coatings by at least two orders of magnitude in order to observe splittings large compared to g. One interesting prospect is the use of cylindrical piezoelectric transducers which could be custom-ordered with the same diameter as the mirror substrates. Under an applied voltage, the piezos would then squeeze the substrates, and a slit in the piezos could perhaps break the cylindrical symmetry, providing a preferred birefringent axis.

6.2.3 Cavity length

The lab 11 cavity, 42.2 μm long, supports 99 half-wavelengths at $\lambda_{Cs} = 852$ nm and 90 half-wavelengths at $\lambda_{FORT} = 936$ nm. A cavity which is resonant at both wavelengths must therefore have a length which is an integer multiple of $\frac{11\lambda_{Cs}}{2} = \frac{10\lambda_{FORT}}{2}$. In order to maximize g, we would like to build the smallest such cavity that will still allow us to focus lattice beams through the cavity side gap, as in lab 11.

Imagine a Gaussian beam focused through a gap which is tens of microns wide and 1 mm long. Let the x-axis lie along the path of the beam with its origin at the center of the gap. If z is the beam waist at the focus (at $x = 0$), then the beam waist as a function of x, z, and the wavelength λ is given by

$$w(x, z, \lambda) = z\sqrt{1 + (\frac{x\lambda}{\pi z^2})^2}. \tag{6.1}$$

We want to minimize w at $x_{gap} = 0.5$ mm, the edge of the gap. We thus solve

$$\frac{dw(x_{gap}, z, \lambda_{Cs})}{dz} = 0 \tag{6.2}$$

to find $z_{min} = 11.6$ μm. The minimum gap between the mirrors is then given by

$$2w(x_{gap}, z_{min}, \lambda_{Cs}) = 32.9\mu\text{m}. \tag{6.3}$$

The smallest co-resonant cavity that can accommodate this gap supports 88 half-wavelengths at λ_{Cs}, 80 half-wavelengths at λ_{FORT}, and is 37.5 μm long. This is then the target length for our cavity.

The preceding discussion has not addressed the problem of diffraction. Siegman notes that an aperture criterion of $4.6w$ is necessary to reduce spatial intensity variation to 1% [9]. In practice, we need to compromise between the demands of strong coupling and the desire for a uniform intensity along the cavity axis. What we have seen in lab 11 is that a 42.2 μm gap allows us to perform effective lattice-based cooling, but that diffraction may be at the root of our difficulties in using side beams for optical pumping.

6.2.4 Rethinking passive stabilization

In past cavity systems, passive vibration isolation from the chamber and optical table has been provided by a series of heavy copper pieces separated by rubber bumpers of Viton and RTV. [21] We had two concerns regarding this design. First, it seemed

likely that outgassing from the Viton and RTV would limit the background pressure we could attain in our new system. Second, Viton can only be baked to 200 °C, and we hoped to bake out the system at higher temperatures.

On the recommendation of Ching-Tzu Chen, a graduate student in the Yeh group, we used an accelerometer to analyze the sensitivity of the cavity system. By measuring acceleration as a function of driving frequency at the location of the cavity mount, we established a baseline for the current system which we could then use to evaluate the efficacy of new designs. An alternate method to analyze the passive stability of the cavity is to lock it to a laser and then monitor the error signal from the lock; we also explored this option, using both the 852 nm probe and 936 nm FORT lasers.

The simplest redesign consisted of replacing both the Viton and RTV with Kalrez, a material which Travis discovered in his research. Viton and Kalrez are both manufactured by Dupont Dow; Viton is a fluoroelastomer, while Kalrez is a perfluoroelastomer, about ten times more expensive, has superior outgassing properties, and can be baked to 316 °C. (RTV stands for "room temperature vulcanization" and may refer to a number of silicone products.) In addition, the bumpers cut from Kalrez O-rings were loaded radially rather than axially, a technique known as streamline loading which produces a more uniform compression and improved damping. The measured sensitivity of the new system was found to be as good as the old one, with some improvements at low frequencies [110].

A more ambitious redesign came out of our discussions with LIGO staff scientist Riccardo DeSalvo. He suggested that we suspend our cavity mount from the walls of the vacuum chamber by employing the mini Geometric Anti-Spring (mini-GAS) system currently in use at the TAMA gravitational wave antenna in Japan [111]. In collaboration with Travis and Riccardo, Italian engineers who had produced the TAMA suspension designed a four-spring system made of beryllium copper that was compatible with our vacuum chamber. Its all-metal construction ensured that it would not be the limiting factor in determining bake-out temperature, and the anti-spring design avoids the long-term creep that ordinary spring systems experience. Although the design would provide effective passive suppression of high-frequency

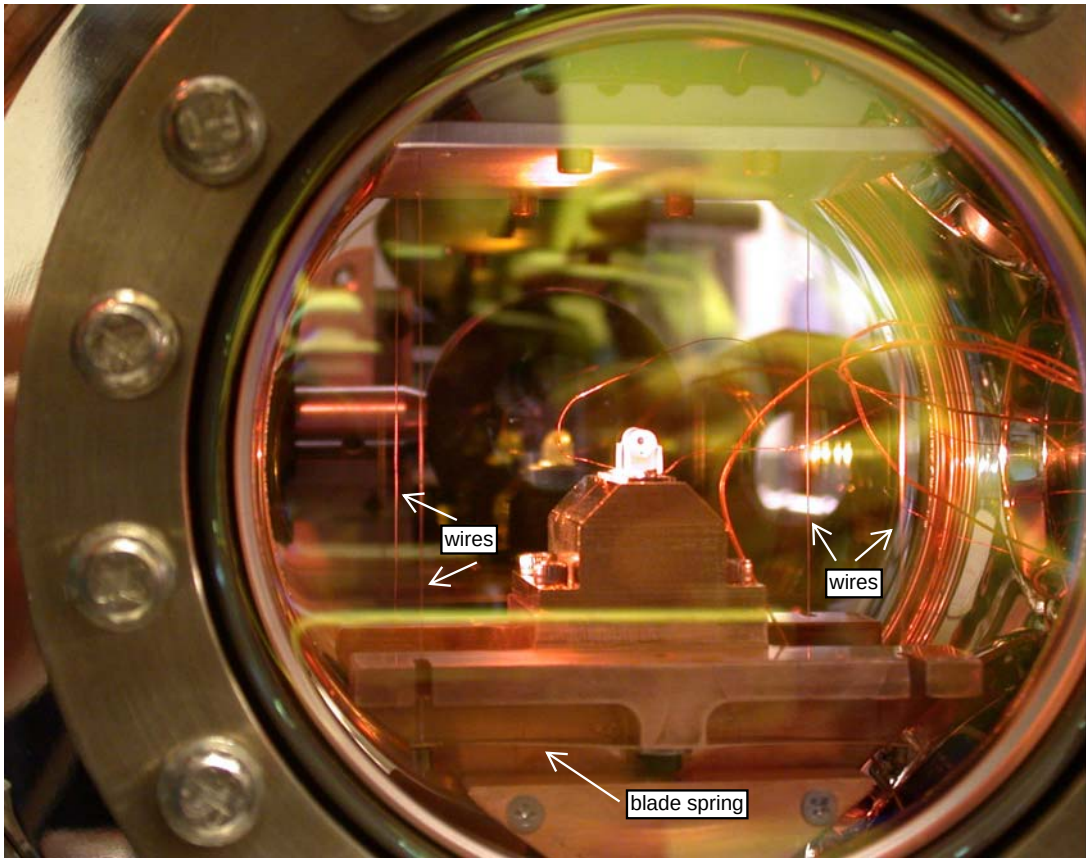

Figure 6.7: The mini-GAS suspension mount supports a cavity inside the lab 1 test vacuum chamber in June 2006. Note the four thin BeCu wires which are secured to blade springs beneath the mount.

noise, it would have a $\sim$ 2 Hz pendulum mode and a $\sim$ 8 Hz leaf spring resonance which would both require an active servo for compensation; we hoped to achieve this initially by using magnetic fields from the MOT coils to induce eddy-current damping in the mount.[1] The system arrived from Italy shortly before Travis's graduation. It was assembled within a test chamber (Figure 6.7) and and characterized with both the accelerometer (Figure 6.8) and cavity lock techniques, which confirmed that the low frequency resonances were substantial; an attempt to observe damping from the MOT coils was unsuccessful. One possibility which might circumvent this difficulty

[1]The slow damping time of pendulum modes and daily creep are two reasons why the Chapman group, for example, has moved away from their earlier cavity suspension, despite its demonstrated excellent passive stability [109].

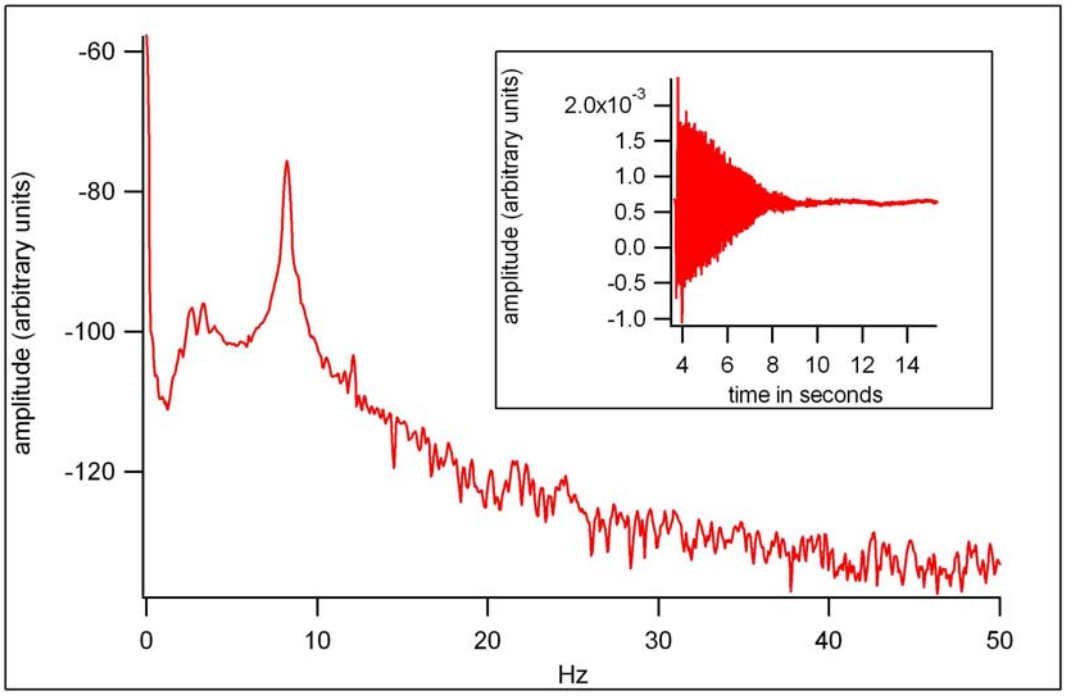

Figure 6.8: Accelerometer measurements of the mini-GAS suspension. The suspension mount is displaced and its ringdown as a function of time is recorded (inset). A Fourier transform of this data reveals expected resonances at around 2 Hz and 8 Hz, corresponding to the pendulum mode and the leaf spring resonance. Vibrations at higher frequencies are quickly damped by the spring system.

would be to use the suspension to support future microtoroid experiments. As light is coupled in and out of the microtoroids via optical fiber rather than in free space, the motion of the cavity with respect to the vacuum chamber is no longer a cause for concern, and perhaps low-frequency active damping could be avoided entirely.

6.2.5 Vacuum chamber

We decided to purchase new vacuum chamber components rather than attempt to reuse the old cavity chamber. This allowed us not to worry about damaged knife edges from old components that might compromise vacuum, and it provided us with a backup chamber in which we could test vibration isolation.

I attempted to stay as close to the previous chamber design as possible, for simplicity and ease in reusing input and output optics. The lower chamber had been

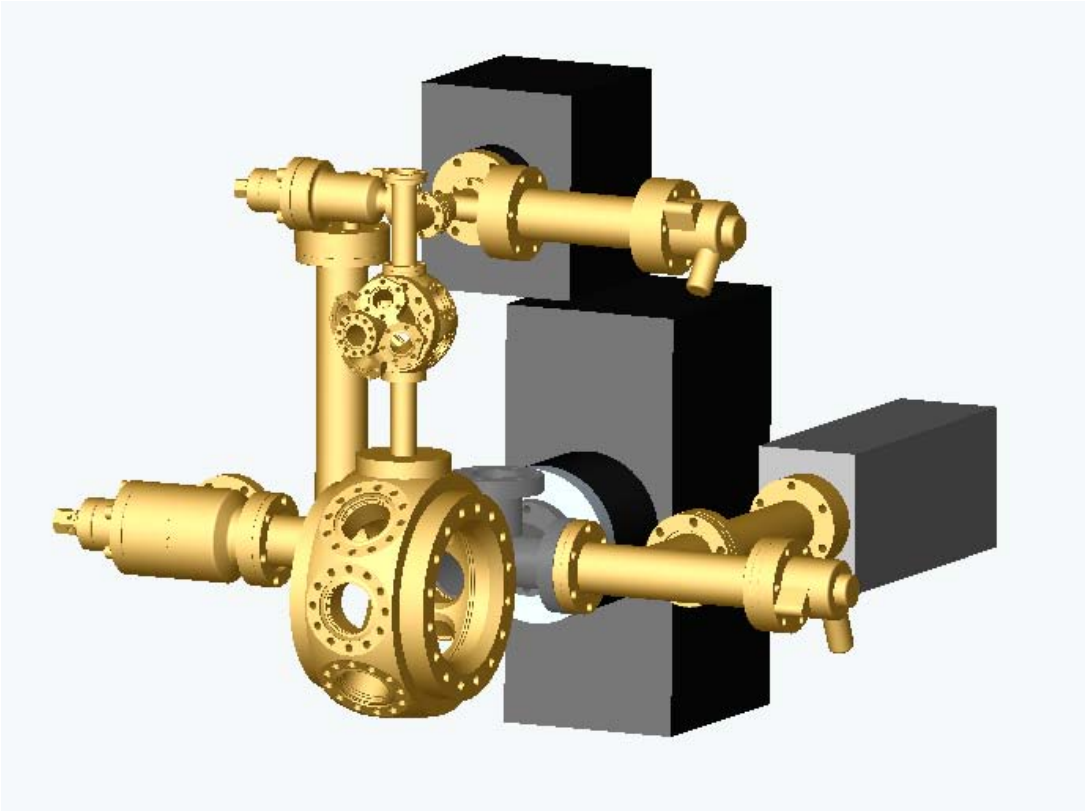

Figure 6.9: Vacuum chamber design for a new asymmetric cavity. Much of the earlier lab 1 design remains intact, with an upper and lower chamber connected by a differential pumping tube and brought to UHV with two Varian ion pumps. New components include a residual gas analyzer (RGA), a titanium sublimation pump, and a multiplexer port on the upper chamber in order to incorporate getters.

custom-built by NorCal and was replaced with a Kimball Physics 6" spherical octagon of similar dimensions. The upper chamber was again a 2.75" Kimball Physics spherical hexagon. The 55 l/s and 20 l/s Varian VacIon Plus Starcell ion pumps that had been used for the lower and upper chambers were also replicated. The new differential pumping tube was identical to the older one, as it was a backup version that had been machined at the same time. Both chambers now had nude Bayert-Alpert ionization gauges from Varian, model UHV-24 for the upper chamber and UHV-24p (lower pressure limits 2×10^{-11} torr and 5×10^{-12} torr, respectively) for the lower chamber. All-metal gate valves from VAT could be used to pump down the upper and lower chambers simultaneously.

From his experience building vacuum chambers at Michigan, Russ Miller suggested that we include a residual gas analyzer (RGA200, SRS) in order to diagnose any leaks we encountered and a titanium sublimation pump (Lesker) to reduce the final system pressure by an order of magnitude. As we were concerned about the possibility of titanium sputtering near our mirrors, there is no line of sight between this pump and the cavity.

The previous chamber had used a cesium ampoule to load the upper chamber MOT from a reservoir of background gas. We decided instead to implement the newer technique of "getters," alkali metal dispensers which emit a vapor when they are resistively heated [112, 113]. Using getters is simpler, more compact and avoids problems associated with successfully breaking the glass ampoule under vacuum. It also permits a lower background pressure in the chamber, especially if the getters are operated in a pulsed configuration. However, the metal in each dispenser eventually becomes depleted; how long this takes depends on the current at which the dispenser is operated, as well as whether it has been subjected to short bursts of high current. We have seen getters in use in the atomic ensembles experiment (lab 2) last only several months. In comparison, the amount of cesium provided by the lab 1 ampoule has been sufficient for the past ten years. Cindy Regal reports that at JILA, getters are used for glass chambers, but in stainless steel chambers (where alkali atoms are absorbed by the walls), ampoules are still used. Our getter design replaced a 2.75" window on the upper chamber with a five-port 2.75" multiplexer (Kimball Physics); the center port was for a MOT beam, while the others could be used for getters. With two getters per port, we could in principle stock our chamber with eight getters and hope for several years of operation before they were all depleted. One encouraging factor is our chamber geometry, which allows us to place the getters $\sim$ 2 cm from the upper MOT. Thus, we are able to form a bright MOT with lower currents than have been necessary in lab 2 and, more recently, in the microtoroid experiment: Dal and summer student Jie Wu measured $N_{atoms} \sim 10^8$ with 3.25 A through the getters, well below the threshold current of $\sim$ 4 A where the response becomes nonlinear [114]; labs 2 and 11a operate at or above this threshold. (It is, however, necessary to

Vacuum Components	°C
Stainless steel chambers, upper and lower	450
Stainless nipples and flanges	450
Stainless tees, crosses, multiplexers	450
Differential pumping tube	450
Getter feedthroughs	450
PZT feedthrough	450
Valves for roughing pumps	450 open
	350 closed
Viewports, without AR coating	400
Ion pumps without magnets	350
Titanium sublimation pump cartridge	350
Residual gas analyzer (RGA)	300

Table 6.2: Maximum baking temperatures for vacuum chamber components

operate the getters briefly above threshold during bake-out as part of an "outgassing" process to remove impurities from the surface. This process should be repeated after each instance in which the getters are exposed to air [115].)

In Table 6.2, I list the maximum baking temperatures for our vacuum chamber components, also reproduced in Yat's thesis. In the past, we have obtained UHV viewports from Larsen with anti-reflection (AR) coatings applied by Guernsey. The uncoated viewports have a baking temperature of 400 °C limited by the glass-to-metal transition, but according to Guernsey engineers, the coatings could not be baked above 250 °C. As this would potentially limit our bake-out temperature, we arranged for a coating run by Advanced Thin Films (Longmont, Colorado) on bare Larsen viewports. These coatings reflect less than 0.5% of light between 800 and 950 nm and no longer limit the viewport baking temperature.

In order to determine our target bake-out temperature, Yat was able to find measured thermal desorption spectra for 316LN stainless steel. [116] These indicated a desorption peak for water molecules at around 300 °C. Thus, even if our cavity mount assembly required lower temperatures, we would at least plan to pre-bake the chamber itself above 300 °C for several days. We purchased a large custom oven from Milmetco which could reach temperatures of 600 °F (316 °C). We threaded a 48 inch braided bellows through a hole in the side of the oven so that we could bake the

chamber under vacuum. Inside the oven, one end of the bellows attached to a tee from which two smaller bellows connected to the upper and lower all-metal valves of the chamber. Outside the oven, the other end of the bellows allowed us to pump either with a turbo pumping station (Pfeiffer TSU071) or, at lower pressures, with an ion pump (VacIon Plus 40 Starcell). The chamber, of course, has its own ion pumps attached, but as their magnets have to be removed temporarily for baking, we are not able to use them inside the oven.

6.2.6 Current project status

We assembled the two-chamber vacuum system, absent the cavity or the cavity mount apparatus, in the fall and winter of 2005. We successfully demonstrated a pressure of $\leq 5 \times 10^{-12}$ torr in the empty chamber (the resolution limit of the nude ion gauge) one month after firing the titanium sublimation pump, with blanks substituted for the viewports during this initial test. Dal has also explored the option of placing an unbaked cavity (with the cavity mount components baked separately before assembly) inside of a pre-baked chamber and vibration isolation system, venting it to nitrogen during the $\sim$ 15 minute transfer process, and he has confirmed that the chamber with cavity returns to its initial pressure after several months. (In this instance, the pressure was $\sim 10^{-10}$ torr in order to expedite the process. As the use of the titanium sublimation pump should improve the pressure by an order of magnitude, this result suggests that we may not need in-situ baking of the cavity mirrors.)

At the moment, the chamber is simply waiting for a cavity. As I have discussed in this chapter, some questions about how to improve the cavity design, such as incorporating control over birefringence, have not yet been resolved. The primary cause for delay, however, has been our quest for better cavity mirrors. In the following chapter, I describe our goals in improving mirror losses, the process of mirror characterization, and the results of this effort to date.