Chapter 7

In search of the perfect mirror

7.1 New mirrors for a new cavity

In October 2004, the decision was made to replace the lab 1 cavity built by Kevin Birnbaum and Theresa Lynn with a new, single-sided cavity. Central to the cavity project was the need to re-establish connections with the manufacturers of our high-finesse mirrors. Collaborations between our group and Ramin Lalezari of PMS Electro-Optics in Boulder, Colorado (later Research Electro-Optics) had produced record-low mirror losses in 1991 [8]. These "supermirrors" had total losses (transmission plus scatter and absorption) of just 1.6 parts per million, corresponding to a cavity finesse of 1.9×10^6. A few years later, Quentin Turchette and postdoc Michael Chapman worked with Research Electro-Optics (REO) to develop the tapered 3 mm / 1 mm mirrors which are now used in several labs worldwide. Christina Hood then traveled to REO in order to develop improved mirror cleaning and handling techniques [21].

The state-of-the-art mirrors which our lab relies on for strong-coupling cavity QED were thus a product of extensive back-and-forth dialogue with industry, but by 2004, this dialogue had lapsed for several years. Meanwhile, Ramin had left REO to start his own company, Advanced Thin Films (ATF), in Longmont, Colorado. We hoped that by placing orders with both REO and ATF for new mirror coating runs, we could encourage the two companies to push the scatter and absorption of their mirrors to new lows.

In March 2005, I drove to Colorado with a breadboard of mirror-testing equipment, since neither company had the ability to characterize such low-loss mirrors at 850 nm. I spent several weeks there that spring and summer and have made two subsequent trips to measure more recent ATF coating runs. In this chapter, I summarize that experience, focusing both on techniques to characterize mirrors efficiently and on what we've learned about the present limits of mirror coating technology. Furthermore, I discuss the implications of current mirror technology for proposed cavity QED experiments in our lab.

7.2 The nuts and bolts of mirror testing

We can define a mirror at a given wavelength by its reflection, R, its transmission, T, and its scatter and absorption losses, $S + A = l$, where

$$R + T + l = 1. \tag{7.1}$$

In theory, one could determine the values of R, T, and l by placing the mirror in a laser path, measuring the fractions of the beam transmitted and reflected, and attributing the rest to losses. In practice, when we need to discern transmissions and losses on the order of 10^{-6} or 10^{-7}, detector nonlinearity and scattered light into the detectors present serious problems for this method. A more reliable approach is to construct an optical cavity from two identical mirrors and then to characterize the cavity.

Christina Hood and Jun Ye outline their procedure for characterization of cavity mirrors in Refs. [21] and [42]; specifically, this procedure is based on their 1999 investigation of mirrors from REO coating run no. T95, the source for the current lab 11 cavity and past lab 1 cavities. First, a measurement of cavity finesse $\mathcal{F}$ determines total losses $\mathcal{L} = T + l$, since $\mathcal{F} = \frac{FSR}{FWHM} = \frac{FSR}{2(\kappa/2\pi)} = \frac{2\pi}{\mathcal{L}}$ in the low-loss cavity limit [9]. Here FSR is the cavity's free spectral range, the spacing between longitudinal modes, which can be determined from cavity length or (for a very short cavity) with a wavemeter; κ is the HWHM linewidth of the cavity's TEM_{00} mode. If the mirrors' losses are low enough and the cavity is long enough, κ can be measured directly via cavity ringdown in order to determine finesse. Alternatively, frequency sidebands applied to a probe laser can provide a meterstick for κ as the cavity length is scanned with a piezo. (More sophisticated methods of measuring κ for short cavities are available when the cavity length can be actively locked; Section 4.3 of Ref. [29] provides details.) Next, in order to partition total losses into T and l, cavity transmission and reflection on resonance are measured simultaneously and compared with the input power to the cavity.

This technique presupposes the cavity mirrors to be identical, a reasonable assumption if they are from the same coating run and appear defect-free under microscope inspection. It is also possible to characterize the two mirrors independently by repeating the transmission/reflection partitioning described above with light incident from the opposite side of the cavity [117].

7.2.1 A portable testing apparatus

When we discussed a new coating run with REO in the winter of 2004–5, they had recently completed a mirror coating/coning process for Dieter Meschede's group at the University of Bonn. REO had initially been unable to meet the specifications of $\mathcal{F} \sim 500,000$ until a student arrived from Bonn with equipment to quantify mirror losses at the company. The advantage of having feedback within a few hours about the results of a coating run is that the same ion beam sputtering (IBS) machine can be used again right away for a second coating, without any changes to the machine's settings except a few tweaks indicated by the measurement, thus insuring repeatability. It was decided that I would travel to Boulder to measure the results of an initial test run and all subsequent attempts until the mirrors were found to be consistent with our target values of T and l. I would also be able to re-measure the mirrors after they were coned by REO opticians in order to document any new losses. (From the Bonn run, about 50% of the mirrors that entered the coning process were "catastrophically damaged" and another 20% had their finesse reduced by the process [118]).

In order to replicate Christina's and Jun's measurements in Colorado, I assembled a breadboard of optics equipment which could be transported by car. An external-cavity diode laser (ECDL) at 852 nm would be used to characterize the mirrors. The laser passed through a prism pair, an isolator, and an acousto-optic modulator (AOM) to switch off the beam for cavity ringdown. Coupling through optical fiber provided spatial cleaning and allowed for easy replacement of the 852 laser with another fiber-coupled laser. A HeNe laser, for example, was often used to align the cavity mirrors, since the mirrors have much higher transmission at visible wavelengths. At the fiber output, the beam was telescoped and then mode-matched to the TEM_{00} cavity waist with a lens on a translation stage. Before entering the cavity, the beam passed through a polarizing beam splitter (PBS) cube and a quarter waveplate. Reflected light from the cavity, after a second pass through the quarter waveplate, was then reflected by the PBS and focused onto a New Focus 1801 125 MHz detector; transmitted cavity light was collimated and then focused onto an identical detector. Before the breadboard left for Boulder, undergraduate Yat Shan Au replicated it in the lab 1 "clean hood" so that we would have our own testing and cavity construction setup available [105].

In order to mount the cavity mirrors for measurement purposes, I used two miniature v-groove setups (for 7.75 mm and for 3 mm diameter mirrors), both with nylon-tipped set-screw clamps from above, as shown in Figure 7.1. Each v-groove was originally machined from a single aluminum piece; the groove was then cut in half, with one half mounted in a fixed position while the other was attached to a miniature piezo-driven translation stage from Physik Instrumente (Karlsruhe, Germany). This allowed the cavity length to be adjusted over several millimeters (with the translation stage micrometer) and to be scanned over much smaller distances with a voltage input to the piezo from either a function generator or a battery box.

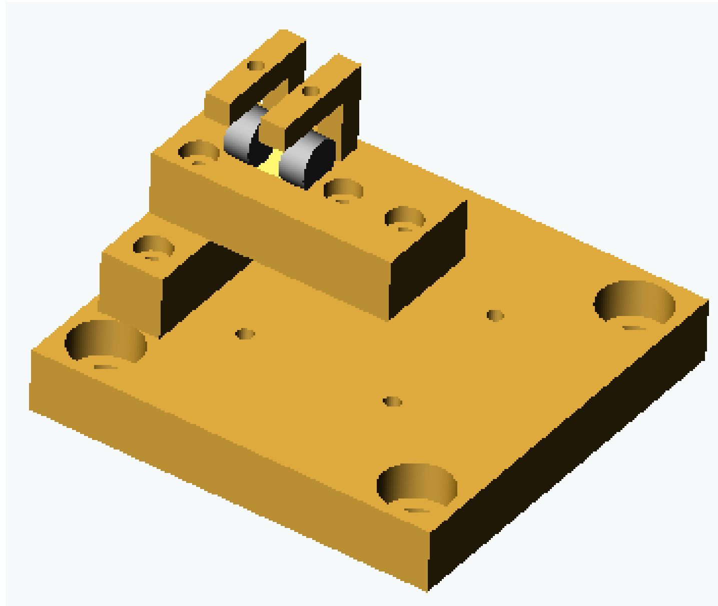

Figure 7.1: Mirror mount assembly for test cavities. Two 7.75 mm mirrors sit in a v-groove (which was later cut in half) and are clamped from above with nylon-tipped set screws. The right half of the v-groove was mounted on a piezo-driven translation stage from Physik Instrumente (not pictured), attached with screws to the base plate. This v-groove block could be quickly exchanged for a smaller one, machined for 3 mm mirrors.

7.2.2 Cavity construction and alignment

When two mirrors were ready to be tested, they were placed into the grooves of the mirror holder blocks using Teflon-tipped tweezers. The rear face of each mirror substrate was gently placed in contact with a shallow notch the end of the groove to ensure reproducible cavity length, and the mirrors were then secured by set screws. Initial alignment was done with all coupling lenses removed and with a HeNe input to the fiber coupler; the mirrors in the beam path before the cavity were then adjusted so that (a) the beam intersected both mirrors near their centers, (b) the reflections from the back of the first mirror and the front of the second mirror overlapped the incoming beam, and (c) the reflected and transmitted beams were incident on their respective detectors. The coupling lenses were then replaced and positioned so as to preserve the beam path. At this point, scanning the applied voltage on the cavity piezo should produce HeNe fringes, visible either on a white card (for slow scan rates, $\sim$ a few Hz) or at the detectors. Simply by switching the fiber-coupled input back to the ECDL laser source, transmission peaks and reflection dips at 852 nm should be visible on the detectors and can then be improved by mirror alignment. Triggering the AOM to switch off once a transmission threshold is crossed provides the sought-after ringdown signal.

There is a rather steep learning curve associated with cavity construction and measurement, as I discovered before traveling to Boulder (and as Dal Wilson has discovered since). Here Christina Hood's alignment procedure outlined in 5.2.2 of Ref. [21] proved invaluable. I would like to add a few points of my own in the hope that they may be helpful to future cavity-builders.

7.2.2.1 Reflected spots

It is worth noting that the rear face of a mirror substrate (either unconed or coned) is not perpendicular to its cylindrical surface. The rear face has been cut at an angle (wedge) in order to prevent problems associated with secondary reflection. (Additionally, the rear face is anti-reflection coated.) This is certainly true for the most recent coating runs, though I am unsure about earlier runs. As a result, when positioning cavity mirrors by aligning their reflected spots, one should expect that the spots reflected from the two faces of a single substrate will not be in alignment, and that their relative position is just a function of the rotation angle of the mirror. Only the spots reflected from the mirror faces themselves need to be aligned.

7.2.2.2 Mirror cleaning

From the REO technician who was most successful at cleaning the coned mirrors, I learned that she used lint-free cotton swabs (Huby-340, distributed by Sanborn Co.) which have a fairly dense tip. After dark-field microscope inspection, she applied a small amount of spectrophotometric-grade solvent to a swab tip. She then made a gentle, quick contact with the mirror surface, rotating her wrist outwards as she brought the swab towards her in order to avoid dragging particles across the surface.

Mirrors are now stored and transported in Gel-Pak containers, mirror face up; the rear surface is attached to the sticky Gel-Pak surface and can be cleaned with solvent upon removal if necessary.

For cleaning the 7.75 mm mirrors, technicians at both REO and ATF used a wafer spin cleaner (PM80, Headway Research). The mirror is set in a chuck, held in place with suction, and then spun via a foot pedal. A solvent is first applied directly to the mirror from a squirt bottle in order to wet it. (Again, the solvent should be spectrophotometric grade. When an empty bottle is filled, it should be first cleaned with compressed nitrogen to remove water vapor and should be filled on the same day as the cleaning.) A glue-free cotton swab is then also wet with solvent. While the squirt bottle is used to keep the mirror continuously wet, the swab is touched to the center of the optic. About ten seconds of gentle pressure is applied as the swab is gradually moved to the mirror edge. The piece is then spun for another 10 to 20 seconds without swab or solvent until it is dry. The appropriate sequence of solvents to use when removing varnish is a) water, b) acetone, and c) isopropyl alcohol. When just removing dust from a mirror, isopropyl alcohol alone is sufficient. Methanol is not well-suited for spin cleaning because it dries too fast. We have subsequently used a LIGO spin cleaner in the East Bridge sub-basement to clean 7.75 mm mirrors and have found this to be a faster, more consistent process than cleaning by hand.

7.2.2.3 Measurement laser

Most previous cavity measurements in our lab by Christina, Theresa, and others were carried out using a Ti:sapphire laser. The large power from the Ti:sapphire was not helpful for these measurements and in fact had to be heavily filtered (see Section 7.2.2.4), but the tunable range of the laser made it possible to map out mirror coating curves [42] and to measure the free spectral range of a small cavity [4]. Only after a series of cavity measurements using ECDLs did I appreciate a secondary benefit of the Ti:sapphire: its narrow linewidth. The supermirror cavities we construct for physics experiments have short lengths ($\sim 40\ \mu$m) and thus relatively broad linewidths ($\kappa \sim 2\pi \times 4$ MHz). But for testing purposes, it is convenient to build longer cavities whose ringdown time $\tau = \frac{1}{2\kappa}$ is much longer than the detector response time or the time to turn off the cavity input field ($\sim$ 25 ns for a good AOM). For example, a ringdown time of 1 μs might correspond to a 5 mm, 80 kHz cavity, still broader than our Ti:sapphire linewidth, but narrower than our ECDLs, which have linewidths of roughly 500 kHz. So even a lossless test cavity on resonance with an input ECDL would only be able to accept and transmit a small percentage of its input power. For mirrors where $T \ll l$, the transmission efficiency is of course much worse, and a signal may be hard to discern from noise. One solution to this difficulty is presented in Section 7.2.3.

7.2.2.4 Cavity input power

Since the cavity input power is multiplied by a factor of $\mathcal{F}$ to generate the intracavity power, and since the waist of the intracavity field is quite small, the power per unit area incident on the cavity mirrors can be quite astonishing, even for small input powers. In lab 11, "lab lore" in the past has been that the cavity input power at 852 nm should never be greater than a few μW to avoid damage to the cavity coatings. However, when searching for a transmission signal while aligning test cavity mirrors, one would like to use as much input light as possible. When I discussed the question of a damage threshold with Sam Richman, the metrology lab manager at REO, he estimated that "continuous wave damage of these IBS coatings is governed by a thermal process that is directly related to the absorption... We don't make any measurements of this continuous wave damage threshold, but it is probably in the neighborhood of 10^8 W/cm^2" [119].

Assuming a radius of $\sim$ 20 μm for the beam spot size on the mirrors, a finesse of 10^6, and perfect mode-matching of a narrow-linewidth laser, this would suggest that a damage threshold might correspond to $\sim$ 1 mW of input power. Since our mode-matching is in practice far from perfect, it seems safe in the future to limit cavity input powers to $\leq 100\mu$m.

7.2.2.5 Loss partitioning

In the measurements of the cavity described in Ref. [42] (the same one still under vacuum in the lab 11 experiment), the cavity length was locked on resonance and resonant powers in transmission and reflection P_t and P_r were recorded. With the additional knowledge of the cavity input power P_i, one can solve equations (2.1) and (2.4) of [42] to find

$$T = \frac{2r(\pi/\mathcal{F})}{1+r}, \tag{7.2}$$

where $r = \frac{P_t}{P_i - P_r}$. Scattering and absorption losses are then given by

$$l = \frac{\pi}{\mathcal{F}} - T. \tag{7.3}$$

Instead of locking the cavity, one can also compare the relative heights of the transmission peak (P_t) and the reflection dip ($P_i - P_r$) on identical detectors while scanning the cavity. Theresa Lynn and I used this technique in early 2003 to characterize some initial coating runs that Ramin had done at his new company. For the technique to be accurate, it is important that the scan time across the peaks be much greater than the cavity ringdown time, so cavity lengths should be relatively short. (A symptom of this problem is a visible asymmetry in the appearance of the cavity scans when the ringdown time is nonnegligible.) If a voltage scan of the cavity piezo is too noisy, one solution is to turn off the applied voltage and let the cavity drift passively across the resonance, capturing images of the transmission peak and reflection dip with a digital scope.

Note that when mode-matching to the cavity is bad, $P_r \gg P_t$, $P_i \approx P_r$, and r depends very sensitively on an accurate measurement of P_r. This difficulty is compounded for a low-transmission coating in which P_r would be greater than P_t even with perfect mode-matching. At REO in Boulder, because of diode laser inefficiencies (7.2.2.4), I was only able to do a very rough partitioning of low-transmission mirror losses.

7.2.2.6 Design of cavity mounts

I had hoped that the machining of the mirror grooves (Figure 7.1) would be accurate enough to define a cavity without adjustment, and that this geometry would be reproducible when one mirror was exchanged for another. Neither of these assumptions held in practice, due in part to centration error on the mirrors from the machining process; that is, after a mirror substrate has been coned down, the angle at which light reflects off the mirror will change as the mirror is rotated around its cylindrical axis. In order to form a cavity, I had to resort to sliding pieces of tape under the second mirror block in order to compensate for the tilt angle, tapping the mirror to rotate it within the v-block, and loosening the screws attaching it to the piezo stage for left/right adjustment.

For the future, I would recommend a design in which the second mirror holder is replaced with a small Lees mirror mount attached to the piezo stage, modified with adapters so that it can accept either mirror diameter.

7.2.3 A simpler solution: self-locking cavity ringdown

After characterizing unconed and coned mirrors at REO in April 2005, I drove the breadboard apparatus 15 miles north to the ATF facilities in Longmont. In August 2005, I returned to ATF to provide quick feedback as they attempted a series of coating runs. During this visit, Ramin suggested a simpler measurement method for cavity ringdown based on optical feedback to the laser from the test cavity.

In ordinary circumstances, optical feedback is the bane of the experimentalist. We use current, temperature, and grating position to carefully select the frequency of our semiconductor laser diodes, and unwanted reflections back into the diode can seriously disrupt that frequency. However, if we instead set up our test cavity to feed back into the laser diode on purpose, the result will be that a) the diode frequency attempts to lock to a cavity mode, with a subsequent line-narrowing [120], and b) because of the resonant frequency and narrowed linewidth, there is a substantial power buildup within the test cavity [121].

Note that for most lab applications, we require a very specific laser frequency (referenced to cesium) that can only be achieved by tuning an external grating. In this case, however, we only need to know about the mirrors' behavior to within a few nm, since that is the scale on which the coating curve varies. We can start with a diode which lases at approximately at 852 nm and then allow the test cavity to do the rest of the work. In doing so, we have discarded the need for isolators, piezos to scan the length of the test cavity in order to match it to the laser wavelength, and the construction of a temperature-stabilized external cavity around the diode.

This approach is particularly appealing for testing low-transmission mirrors, where $T \sim 0.1$–0.5 ppm $\ll l \sim 2$ ppm. The small ratio of T/l means that on-resonance cavity transmission may be reduced by a factor of 100. When coupled with already low efficiencies because of the narrow cavity linewidth with respect to the diode laser's linewidth, the result is a ringdown signal that can be difficult to measure. Here the fact that optical feedback allows the laser to adapt to the cavity becomes a tremendous advantage; Dahmani reports feedback-induced linewidth narrowing by a factor of 1000 [120].

7.2.3.1 A breadboard setup for cavity-locked lasers

At ATF, Ramin had already set up a prototype system consisting of a HeNe laser, a lens, two test mirrors, and a fast detector built by Mark Notcutt of JILA. The detector had a Schmitt trigger that would send a signal to turn off the laser whenever the cavity output exceeded an adjustable threshold, thus triggering a cavity ringdown. When the HeNe was locked to the cavity, the cavity buildup field (or rather, its scatter from air particles) was visible to the naked eye. At the end of my August stay, we adapted the 852 nm ECDL as the basis for a similar setup, with satisfying results: despite the external grating which mitigated the cavity feedback effect, we observed much more substantial cavity buildup than I had previously achieved. As a result, I could fit cavity ringdowns with higher confidence and with greatly reduced preparation time.

Upon returning to Caltech, I constructed and tested a similar arrangement, which Dal and I later set up at ATF in January 2007. As depicted in Figure 7.2, a bare

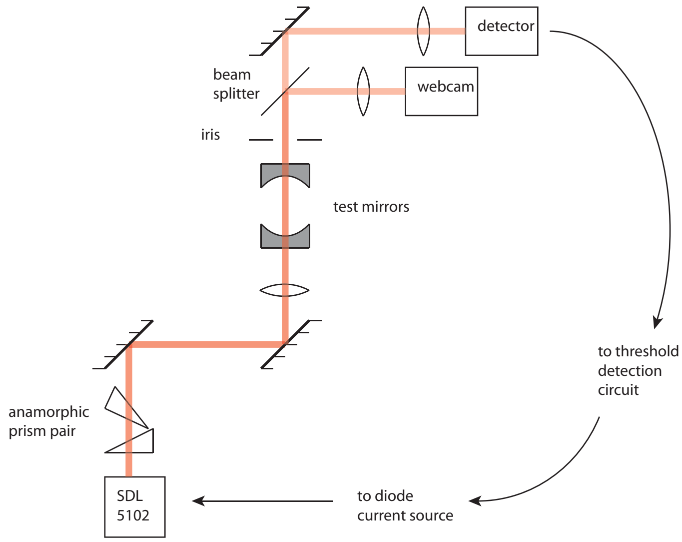

Figure 7.2: A simple breadboard setup for self-locking cavity ringdown, implemented both at Advanced Thin Films and in our lab for testing mirror coatings

SDL-5102 diode is driven by a commercial diode laser current supply (TMD-219 from Power Technology) which can be switched off via a TTL signal. Again, the testing components are simply the diode, a prism pair, two lenses, a detector (Thorlabs PDA10A), and the cavity mirrors. The cavity mirrors are mounted in Newport mirror mounts with a 7.75 mm to 1 inch adapter machined at ATF. The detector signal goes to a variable gain amplifier which sets the turn-off threshold, followed by a monostable multivibrator circuit (adapted from one built by former graduate student James Chou) which provides the TTL signal for the current supply.

Alignment of the system can be done very quickly and reliably. The diode output is first attenuated and then aligned into the detector. The second cavity mirror is placed in the center of the beam path and adjusted so as to align its reflection with the incoming light. (Using an IR viewer, one can sometimes see the diode output light “flash” when feedback is achieved. Alternatively, the very small amount of transmitted light visible on the output detector — with the attenuation now removed — becomes suddenly noisier.) Finally, the first cavity mirror is placed in the beam path, and its retroreflection is also aligned into the diode. Once the retroreflection is nearly aligned, transmission peaks become visible on the output detector, and both cavity mirrors can then be adjusted so as to maximize transmission. Sometimes it is necessary to tap the optical table or the mirror mounts in order for the cavity to find a laser mode to which it can lock; at ATF, this source of vibration was conveniently provided by the somewhat noisy HEPA filter on the optical table.

For low transmission mirrors, an iris after the second cavity mirror can eliminate scattered light into the detector. It is also useful to pick off some cavity output light with a beam splitter and focus it onto a camera (in our case, a USB-connected webcam from which we removed the lens) in order to image the cavity mode. The cavity mirrors can then be tweaked so as to maximize the TEM_{00} mode rather than higher-order modes, which may have different ringdown times as they sample larger areas of the mirror surface. The mirror adapters can be rotated within their mounts in order to sample the mirror coating in different locations and thus characterize its uniformity. It may also be useful to place one of the mounts on a translation stage for this purpose.

One note about reflected light from the cavity: this method relies on that reflection as feedback to lock the diode laser, but as the laser comes into resonance with the cavity, the reflected power and thus the feedback signal drops, creating an unstable situation. Researchers concerned with stable locking have employed techniques including multiple cavity modes [120], waveplates and polarizers [121], and spatial filtering [122] in order to address this issue. For our purposes, we only require the laser to lock to the cavity for long enough to trigger the cavity ringdown, so this level of sophistication is unnecessary. However, in the future, one might consider using a beam splitter to bring some of the reflected cavity light to a second detector and then implementing a loss partitioning scheme similar to that in Section 7.2.2.5.

7.2.3.2 Scattering losses in air

Finally, the test cavities we have built using this self-locking scheme have been several cm long, convenient for accurately measuring cavity length with calipers. In this situation, however, air loss becomes a small but non-negligible issue. Here I am grateful for the help of David Robichaud, a graduate student in Caltech's Okimura group who uses cavity ringdown spectroscopy to study atmospheric free radical chemistry. In Robichaud's air loss calculations, he considers only Rayleigh scattering in the atmosphere, i.e., by N_2 and O_2 (other less significant sources of atmospheric attenuation could be due to Mie scattering of aerosols and to absorption). Rayleigh scattering is the limit of Mie theory in the case where particle size is much smaller than the optical wavelength. The Rayleigh extinction coefficient due to scattering is given by

$$a_s(\lambda) = (\frac{2\pi}{\lambda})^4 \frac{(n^2(\lambda) - 1)^2}{6\pi N}, \tag{7.4}$$

where λ is the wavelength of light, $n(\lambda)$ is the index of refraction, and N is the density of scatterers [123]. Note that $a_s(\lambda) = N\sigma(\lambda)$, where $\sigma(\lambda)$ is the scattering cross-section.

We can calculate N in units of m^{-3} from the ideal gas law: if $T = 296$ K and P = 1 atm, then $N = 2.48 \times 10^{25}/m^3$. The refractive index of air at 852 nm is about 1.000269 [124]. Using equation (7.4), we find $a_s = 1.83 \times 10^{-8}$. The attenuation due to Rayleigh scattering is then given by $I(x) = I(0)e^{-a_s x}$, where $I(x)$ is the intensity of light at distance x from the source. So for each mirror bounce, i.e., each trip of distance d within a cavity of length d, the light is attenuated by a factor of $e^{-a_s d} \simeq (1 - a_s d) = (1 - l_{Rayleigh})$. If d is 10 cm, then $l_{Rayleigh} = 0.18$ ppm, and $l_{Rayleigh} + l + T = \mathcal{L}_{total}$, where $\mathcal{L}_{total}$ is the per-mirror loss inferred from cavity ringdown.

7.3 Results of the coating runs

7.3.1 REO

Our agreement with REO was for "best effort" fabrication, coating, and coning of two sets of mirrors. Each set consisted of ten mirrors with 5 cm radius of curvature (ROC) and five mirrors with 10 cm ROC. In both cases, the target $S + A$ losses were < 2 ppm; for the first (high transmission) set, the target transmission was 15 ± 2.5

ppm at 852 nm, while for the second (low transmission) set, the goal was $T = 0.1 - 0.5$ ppm. An additional specification was for $R > 99.9\%$ at 936 nm, so that the cavity mirrors could support a magic-wavelength FORT as in lab 11. (In order to achieve this for the high transmission mirrors, the center wavelength of the coating would in fact be higher than 852 nm, though the transmission at 852 nm would still meet our target.)

Once I had arrived in Boulder, REO did a test run of high and low transmission mirrors. The initial high transmission mirrors were found to have $T =$ 8–10 ppm at 852 nm, so the coating engineers shifted the design parameters to aim for a higher transmission in a second test run, which then met specifications. Sam Richman was perplexed by a $\sim 25\%$ discrepancy between my ringdown data and transmission measurements done on the REO spectrophotometer, but after the same discrepancy was found to hold for some 15 ppm mirrors which he characterized at 633 nm, we concluded that the spectrophotometer calibration was at fault, further evidence for the necessity of on-site ringdown measurements. Meanwhile, the test run for low transmission mirrors was consistent with our targets for T and l, and so the final coating run replicated this design.

After the completion of the final coating runs, I characterized the mirrors before they were coned. One unfortunate fact of the final high transmission coating was an unusual number of large defects at 200x magnification, noticed by the technicians who first inspected the mirrors. The result was that while my best measurements indicated $T =$ 15.5–16 ppm and $l =$ 1.5-2 ppm, I also observed l as high as 5 ppm. On two different pairs of mirrors, I did a series of four or five different ringdown measurements, each time adjusting the beam position on the mirrors. Variation of loss across a single mirror pair was consistent with the 17–21 ppm variation I observed across different pairs, indicating that mirrors weren't wholly "good" or "bad"; rather, all mirrors had defects that needed to be avoided. REO technicians planned to use microscope inspection in order to cull the mirrors with the fewest number of defects for coning. In contrast, measured losses of the low transmission run were very consistent, in agreement with their defect-free appearance under the microscope. Measured total

coating run	superpolishing run	measured T, l
RN# L3-2029,L3-2030	5 cm: LT# IK838 10 cm: LT# IK716	$T = 15.5 - 16$ ppm $l = 1.4 - 1.7$ ppm
RN# L3-2034,L3-2039	5 cm: LT# IK839 10 cm: LT# IK840	$T = 0.1 - 0.2$ ppm $l = 1.4 - 1.7$ ppm

Table 7.1: Measured transmission and losses for coating and superpolishing runs at REO, April 2005

losses of 1.6–1.8 ppm along with inferred transmission $T = 0.1$–0.2 ppm (consistent with spectrophotometer data) suggested that $S + A = 1.4$–1.7 ppm.

The REO machinist struggled with the coning process, and only about a third of the mirrors weren't visibly damaged afterwards. In her thesis, Christina Hood speculated about possible damage to the mirror coatings from the machining process [21], but I found that the surviving mirrors had not experienced any measurable degradation in finesse. Moreover, for the high transmission set, the large number of mirrors initially coated allowed enough coned ones to be produced with no major defects in the central millimeter. I was able to select ten high- and low-transmission 5 cm ROC mirrors which met our ringdown standards, and the 10 cm mirrors were then chosen by microscope inspection. Table 7.1 lists the coating and superpolishing run numbers for these mirrors.

During my 2005 visits, I emphasized our hope to push the scattering and absorption losses to new lows. There was certainly interest in this project among the metrology staff and coating engineers, who were curious to understand whether minimum losses were dominated by scattering or by absorption. In their opinion, the primary suspect was absorption: while scattering decreases with better superpolishing of the substrates, they doubted that there was further room for improvement in that direction. Absorption, on the other hand, may be due to impurities in the metal targets inside the ion beam sputtering chamber, or sputtering of the chamber surfaces at the edges of the ion beam target.

In the following months, Sam Richman undertook a series of experiments at REO to explore the possibility of loss reduction. Test parameters included the use of different coating materials, sputtering chambers, and deposition settings. Mirrors were coated at both 1064 nm and 633 nm, two wavelengths at which in-house characterization is straightforward. If the results of these experiments were promising, he hoped to proceed with an 852 nm coating run which we could then characterize ourselves. However, Sam was unfortunately not able to see any significant improvements in mirror losses.

7.3.2 ATF

In addition to our order for mirror coating and coning, we were also able to purchase superpolished substrates from REO which we could provide to Ramin Lalezari at ATF. The 5 cm and 10 cm ROC substrates were from the same batches (IK716, IK838, IK839, IK840) as those coated at REO. Ramin would then undertake a series of five coating runs using these substrates, again hoping to minimize scattering and absorption. This project included constructing new tooling for holding the substrates, which he hoped would minimize point defects during coating.

In early August of 2005, I returned to Colorado for ringdown measurements at ATF. The first test run had been completed just before my arrival and had a target transmission of 5 ppm, though in fact total losses were found to be only 5–5.5 ppm and rough partitioning of transmission and reflection suggested $T \sim 2.9$ ppm, $S + A \sim 2.3$ ppm. The second coating run took place while I was there and had a target transmission of 0.6 ppm, but total losses per mirror were 2.9–3.2 ppm, and partitioning suggested $T \sim 0.3$ ppm, $S + A \sim 2.6$ ppm. (The measured transmissions for the two runs were self-consistent; that is, if T were in fact 3 ppm instead of 5 for the first run, then we would expect the transmission for the second run also to be around half of its target value.)

A third coating run in September 2005 included superpolished substrates purchased from General Optics to compare with those from REO. However, both substrate types from this run had total losses of approximately 3 ppm, no better than the mirrors from the second run.

As Ramin and his colleagues had exhausted their ideas for reducing losses, the project remained on hold until late 2006. At that time, ATF had made recent progress on mirror losses at other wavelengths by reducing contaminants in the coating film due to parts of the chamber being sputtered. Ramin hoped that they could apply what they had learned to a new coating run, which took place in January 2007. Dal and I both traveled to Longmont, where we set up the new test apparatus (Section 7.2.3) and discussed possible approaches to mirror coning with Ramin and Jeff. The mirrors from this run were found to have $L \approx 2.7$ ppm, or $l = 2.2$ ppm incorporating estimated transmission and air losses. These mirrors were later re-annealed at 500 C, but with no measurable improvements.

Finally, a fifth run in March 2007 used a higher purity SiO_2 target within the coating chamber. This batch was noticeably cleaner than past runs under microscope inspection, and both Ramin's initial measurements and subsequent tests here at Caltech confirmed that l = 1.7–1.9 ppm. For the first time, the ATF coatings demonstrated losses as low as those measured at REO. However, spectrophotometer data indicates that the mirrors from this run are not as reflective at 936 nm as we would need to support a FORT, so in order to obtain suitable cavity mirrors, a new low-transmission coating run would be required, as well as a corresponding high-transmission run.

ATF does not have the capabilities to cone these mirrors down to the tapered 3 mm / 1 mm dimensions which we have used in the past. With assistance from Ramin, Dal has been pursuing some promising options for having this machining done by a third party.

7.4 Single-sided cavity calculations

Given the effort put into mirror development, it is worth considering quantitatively what impact these new mirrors could have on our experiments. One important benchmark is the probability that a photon generated within a cavity will be transmitted by the “open” (high-transmission) mirror. This is given by

$$P_{trans} = \frac{T_o}{T_c + T_o + 2l}, \tag{7.5}$$

where T_o and T_c are transmission through the “open” and “closed” mirrors, respectively. For the current lab 11 cavity, $T_o = T_c = 4.3$ ppm, $l = 2.9$ ppm, and $P_{trans} = 0.30$ [42]. This constitutes the most significant loss in our output detection path. If, instead, $T_o = 15.5$ ppm, $T_c = 0.1$ ppm, and $l = 1.5$ ppm, based on the REO coating run measurements, then $P_{trans} = 0.83$, a nearly threefold improvement. This transmission probability also governs the reverse process, that is, the acceptance of a resonant photon into the cavity by a STIRAP process [31].[1] Thus, in order to improve the efficiency of the coherent state transfer process described in Chapter 3, and especially for future experiments in which a photon extracted from one cavity is mapped into another (or the same) cavity, a single-sided cavity using currently available mirrors would be very useful.

More problematic is the question of experiments which rely on reflecting a photon back from the cavity input. For example, Ref. [44] describes a reflection-based scheme for QND measurement of photon number: an atom with ground states $|a\rangle, |b\rangle$ and excited state $|e\rangle$ within an optical cavity is prepared in a superposition of the two ground states, and the cavity is tuned to the $|b\rangle \to |e\rangle$ transition. If a photon resonant with this transition is sent to the cavity, then it can be detected afterwards by a rotation on the atom followed by measurement in the basis $\{|a\rangle, |b\rangle\}$. This can be seen by considering the two ground state cases separately; the photon causes a phase flip only if the atom is in $|a\rangle$:

$$|a\rangle \to -|a\rangle \tag{7.6}$$

$$|b\rangle \to |b\rangle, \tag{7.7}$$

[1] One might imagine that only the ratio of T to l at the photon input port need be considered. However, by a time-reversal symmetry argument, both ports must be taken into account, since the cavity can emit photons through either mirror. Thus, in order to map photons efficiently into a symmetric cavity, both mirrors would ideally be used as inputs. A single-sided cavity allows us instead to use just one mirror.

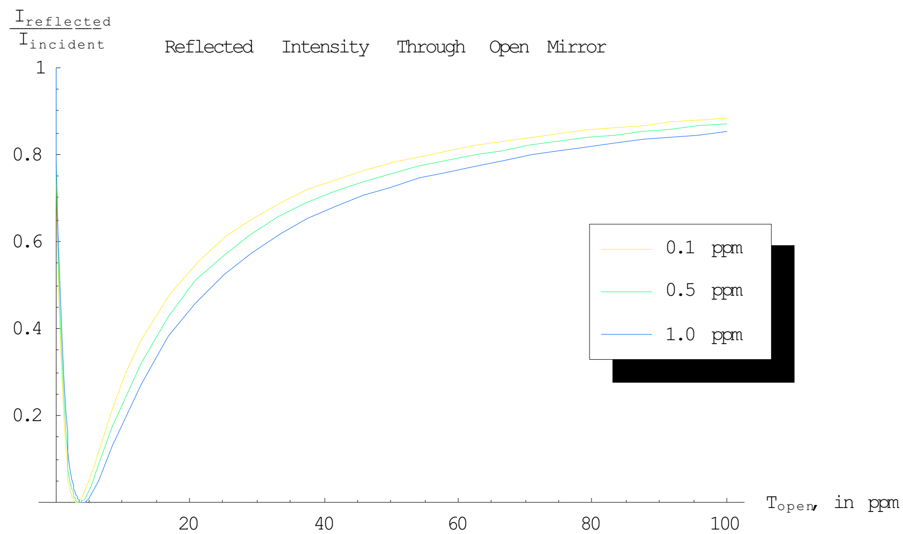

Figure 7.3: Fraction of incident intensity that is reflected from the transmissive mirror of an asymmetric cavity. Here both mirrors are assigned scattering and absorption losses $l = 1.5$ ppm, and T_o is plotted from 0 to 100 ppm for three values of T_c: 0.1 ppm, 0.5 ppm, and 1.0 ppm.

so that $|b\rangle + |a\rangle$ is mapped to $|b\rangle - |a\rangle$ only in the presence of a photon. A $\pi/2$ rotation subsequently maps

$$|b\rangle + |a\rangle \longrightarrow |b\rangle \tag{7.8}$$

$$|b\rangle - |a\rangle \longrightarrow |a\rangle, \tag{7.9}$$

two orthogonal states which can be distinguished quickly and efficiently via state detection with a probe laser. [30]

For lossy mirrors, however, the probability that a resonant photon is reflected from an empty cavity (that is, when the atom is in $|0\rangle$, dark to the cavity transition) is given by

$$P_{refl} = (\frac{T_c - T_o + 2l}{T_c + T_o + 2l})^2. \tag{7.10}$$

This function has a minimum at zero when the impedance-matching condition $T_o = T_c + 2l$ is satisfied. It is plotted in Figure 7.4 as a function of open mirror transmission for $l = 1.5$ ppm, $T_c = \{0.1, 0.5, 1\}$ ppm. We see that for our target transmission,

$T_o = 15$ ppm, only about 40% of the incident light is reflected. In the QND scheme, a photon would always reflect when the atom was in $|b\rangle$, but would be lost more than half the time when the atom was in $|a\rangle$.

We could, of course, increase the transmission of the open mirror to improve the reflection efficiency. The problem is the corresponding decline in cavity finesse and increase in κ; we gradually move out of the strong coupling regime. The QND protocol requires strong coupling so that photons coupled to the $|b\rangle \to |e\rangle$ transition are completely off-resonant when an atom in $|b\rangle$ is present in the cavity. In this case, the loss of strong coupling would mean a state-dependent phase shift less than π and nonorthogonal final atomic states.

7.5 Outlook

Unfortunately, neither REO nor ATF have been able to produce mirrors with losses smaller than those of the 1991 supermirror run [8]. However, at this point, both companies have the capability to produce mirrors with $l < 2$ ppm and a more thorough understanding of what contributes to loss in their facilities. When the coning process is successful, it seems to have no measurable effect on cavity finesse, and by outsourcing the coning to a skilled glass machinist, we may be able to avoid losing mirrors in the process. Using a simple new technique for cavity ringdown, we are now able to characterize mirrors quickly and reliably, with possible extensions to loss partitioning.