

Thermal Noise Associated with Silicate Bonding

Senior Thesis By

Joshua Ryan Smith

Syracuse University Department of Physics

Syracuse, New York

Submitted May 9, 2002

Advisor: Professor Saulson

Acknowledgement

There are many people who have made my four years at Syracuse University a joy, and deserve to be thanked. I would like to thank Steve Penn, Gregg Harry, Andri Gretarsson, and Scott Kittelberger for their patience in teaching me about working in the lab, and research physics in general. Thanks also to my family, especially Connie and Pat for their support and encouragement, my grandfather for writing a chapter of my thesis that was unfortunately lost in the mail, and Aunt Mary for her computer expertise. Finally I would like to thank my thesis advisor, Professor Saulson, for encouraging me to attend Syracuse, encouraging me to stick with physics when I was considering changing majors, giving me a job in his lab, advising me on an internship at LIGO, giving me my own experiment and helping me to writing a thesis about it.

Abstract: This thesis represents the work I have done over the past two years on the silicate bonding experiment at Syracuse University. The sensitivity of an interferometric gravitational wave detector is limited by many sources of noise, among them thermal noise. In an effort to reduce thermal noise, some detectors including GEO600 and LIGO will switch from test masses suspended by steel wires to monolithic suspensions. Silicate bonding allows for the creation of monolithic suspensions of fused silica, but it is not well known to what degree silicate bonding will affect the thermal noise. In this thesis I compare the quality factor of vibration of unbonded fused silica half-rods with that of bonded fused silica rods composed of two half-rods bonded together. The measurements show a significant reduction of the quality factor due to mechanical loss associated with the silicate bond. The effects are scaled to estimate the effect on LIGO monolithic suspensions. These estimates predict that silicate bonding should not be the limiting thermal noise source for Advanced LIGO test masses.

Table of Contents

<u>Acknowledgement</u>	1
<u>Abstract</u>	2
<u>Table of Contents</u>	3
<u>List of Tables</u>	4
<u>List of Figures</u>	5
<u>Chapter I: Introduction</u>	6
1.1 Gravitational Wave Detection	6
1.2 Thermal Noise and Q	10
1.3 Interferometer Test Mass Suspensions	13
1.4 Silicate Bonding	17
<u>Chapter II: Experiment</u>	17
2.1 Sample Characterization	17
2.2 Experimental Setup	20
2.3 Experimental Procedure	23
2.4 Data Analysis Procedure	26
<u>Chapter III: Results</u>	30
3.1 Q Measurements	30
3.2 Scaled results	36
<u>Bibliography</u>	39

List of Tables

Table 2.3.1: Algor-predicted frequencies vs. measured frequencies

Table 2.3.2: Frequency in air vs. in vacuum

List of Figures

Figure 1.1.1: Gravitational wave strain on an interferometer

Figure 1.1.2: Noise curves for initial LIGO (Modified from figure by James Kent Blackburn, 4/11/1998)

Figure 1.1.3: Noise curves for Advanced LIGO

Figure 1.2.1: Thermal Noise vs. Frequency for three values of dissipation

Figure 1.3.1: LIGO Fused Silica test mass supported by metal wire sling

Figure 1.3.2: GEO600 monolithic suspension (Modified from figure by (IGR) Professor Norna Robertson, 3/16/2001)

Figure 1.3.3-4: GEO600 ears bonded to mirror surface

Figure 2.1.1: SilicabondA1

Figure 2.1.2: SilicabondB1

Figure 2.1.3: SilicabondC1

Figure 2.1.4: D1

Figure 2.2.1: Experimental Setup

Figure 2.4.1: Typical ringdown data

Figure 2.4.2: Qfit fit of the data with residual in red

Figure 2.4.3: Frequency and phase shift throughout the ringdown

Figure 2.4.4: Quick Q fit of the data

Figure 3.1.1: Semi-log graph of Q versus frequency for all samples

Figure 3.1.2: Algor-predicted mode shapes

Figure 3.1.3: Q with errorbars representing error as calculated by Qfit.m versus frequency

Figure 3.1.4: Mean Q versus frequency

Figure 3.1.5 Mean Q versus bond area

Chapter I: Introduction

1.1 Gravitational Wave Detection

Einstein's General Theory of Relativity forbids any information from traveling faster than the speed of light. This idea challenges Newton's notion of instantaneous action at a distance, and thus has consequences for the behavior of gravity. According to Relativity, if there is a change in the gravitational field, due for instance, to a planet being swallowed by a black hole, distant objects become aware of it and readjust their positions a finite time later when the news from the source reaches them. The gravitational information propagates as a wave at the speed of light, deforming space-time and readjusting the gravitational field. This phenomenon is called a gravitational wave.

Gravity, unlike electromagnetism, has only one charge, namely mass. This together with momentum conservation implies that an isolated gravitational system cannot have a dipole moment with a nonzero second time derivative. Thus, to produce gravitational waves a system must change its quadrupole moment.


The deformation resulting from a gravitational wave passing through space-time may be seen in the miniscule motions of freely falling bodies therein. If one imagines a field of equally spaced freely falling bodies, the effect of a gravitational wave passing through it at a given time is the stretching of the distances between the test masses in one direction, and the simultaneous shrinking of the separation in the perpendicular direction

The distance changes between these bodies represent a strain, $h = \frac{\Delta L}{L}$, in space-time,

imparted by the gravitational wave [see Figure 1.1.1]. . According to General Relativity,

and our knowledge of astronomy, this strain is incredibly small at large distances from the source, even when emanating from astronomical sources.

Figure 1.1.1: Gravitational wave strain on an interferometer



Albert Einstein predicted the existence of gravitational waves in his general theory of relativity in 1916, but because of their weak nature, it is only recently, with the help of new ideas and technology, that their direct detection has become a realizable goal. The experimental search for gravitational waves began in the 1960s with Joseph Weber, and the operation of his resonant bar detectors. Weber claimed in 1969 [1] that his bars registered coincident gravitational wave signals, however scientists who replicated his experiment found no signs of gravitational waves, and the physics community decided against his assertion, later claiming that the bars had not achieved the sensitivity necessary to expose such signals.

Gravitational waves have not, to this day, been directly detected, but there is little doubt of their existence. In 1974 Joseph Taylor and Russell Hulse of Princeton discovered a pulsar in a binary orbit with another star [2]. Taylor and Hulse carefully

observed the system and showed that it was inspiraling at precisely the rate prescribed by General Relativity for a system losing energy through emanation of gravitational waves. Their work earned them the Nobel Prize in Physics in 1993. The direct detection of gravitational waves would be a scientific milestone, serving as a powerful check of relativity, a new medium with which to observe black holes and other astronomical systems, and perhaps a glimpse at the formation of our universe.

A gravitational wave detector is an extremely sensitive instrument designed to detect the infinitesimal time varying strains in space-time caused by gravitational waves. There are at least two types of gravitational wave detectors: resonant mass detectors and laser interferometric detectors. A number of interferometric gravitational wave observatories are being built around the world [3-6], and most of these should be operational within the next few years. These first generation interferometric gravitational wave detectors are designed to have sensitivity several orders of magnitude better, and over a much wider frequency range, than that of the most sensitive resonant bars. Thus, our current hopes for the direct detection of gravitational waves lie, literally, in the arms of interferometric gravitational wave detectors.

Though these interferometers should be the most sensitive gravitational wave detectors ever constructed, it is still not clear if the first generation detectors will attain the sensitivity required to detect gravitational waves. Therefore, plans for the next generation of interferometers are now being developed, and research is under way to ensure that these interferometers will have the sensitivity necessary to detect several events per year. The sensitivities of these interferometers are limited by fundamental noise sources inherent in the instruments. Among these are seismic noise, shot noise, and

thermal noise, each of which limits the sensitivity over certain frequency bands [see Fig.1.1.2]. In Advanced LIGO, thermal noise from the internal degrees of freedom of the interferometer test masses is expected to be the limiting noise source in the interferometers' most sensitive frequency band (~30 – 300 Hz) [see Fig 1.1.3]. Thus thermal noise, such as thermal noise associated with silicate bonding, will help determine the sensitivity, and subsequently the number of events that advanced LIGO can detect.

Figure 1.1.2: Noise curves for initial LIGO (Modified from figure by J.K. Blackburn, 4/11/1998)

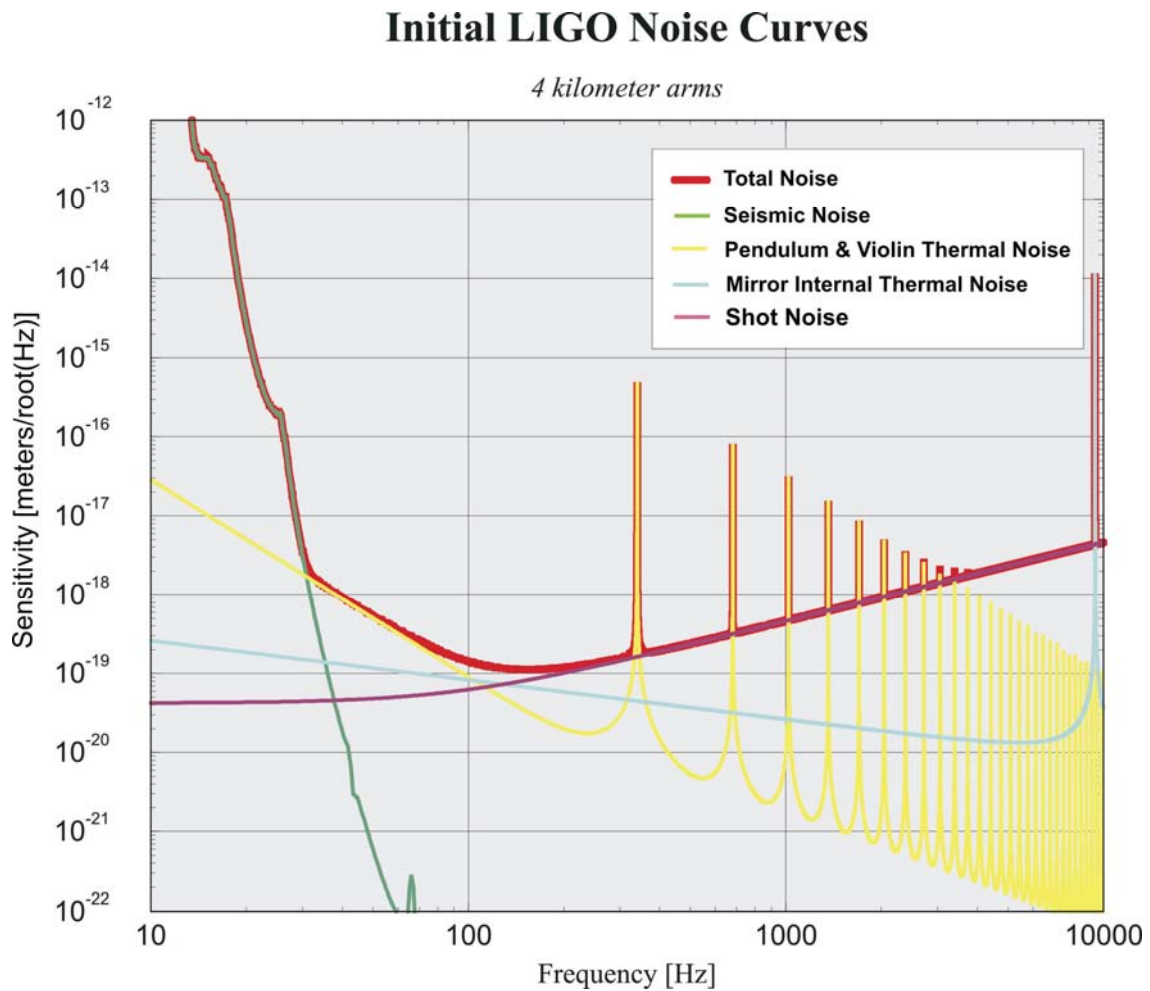
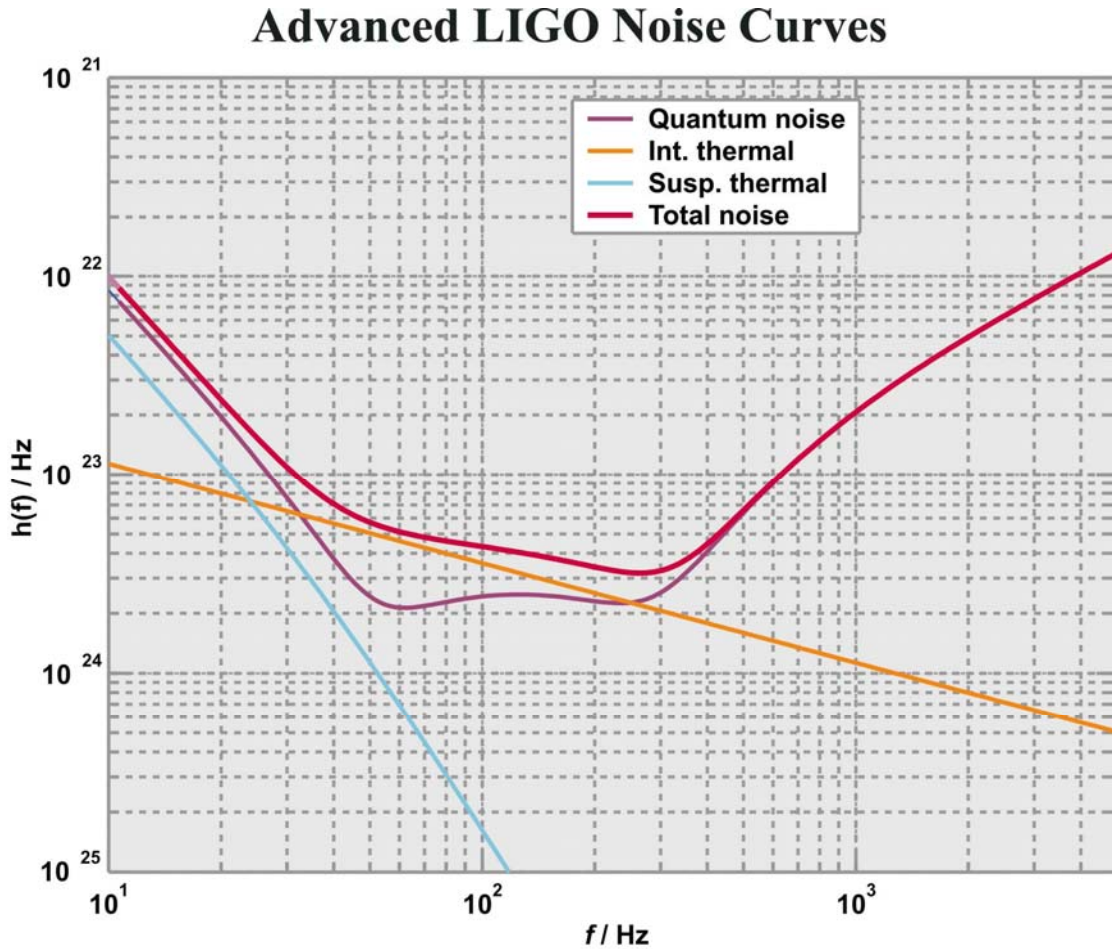


Figure 1.1.3: Noise curves for Advanced LIGO



1.2 Thermal Noise and Q

In the LIGO gravitational wave interferometers, thermal noise will most likely determine the limiting sensitivity between around 30 to 100 Hz, a band that is predicted to be nearly the most sensitive across all frequencies [refer back to Figure 1.1.2]. So, what is thermal noise, and how can we minimize it?

Thermal noise is a generalization of Brownian motion, and a consequence of the equipartition theorem, which states that the each degree of freedom of individual atoms

has thermal vibration with energy $\frac{k_B T}{2}$ on average. This idea is threatening for

gravitational wave detection, because it means that there is a fundamental limit to the

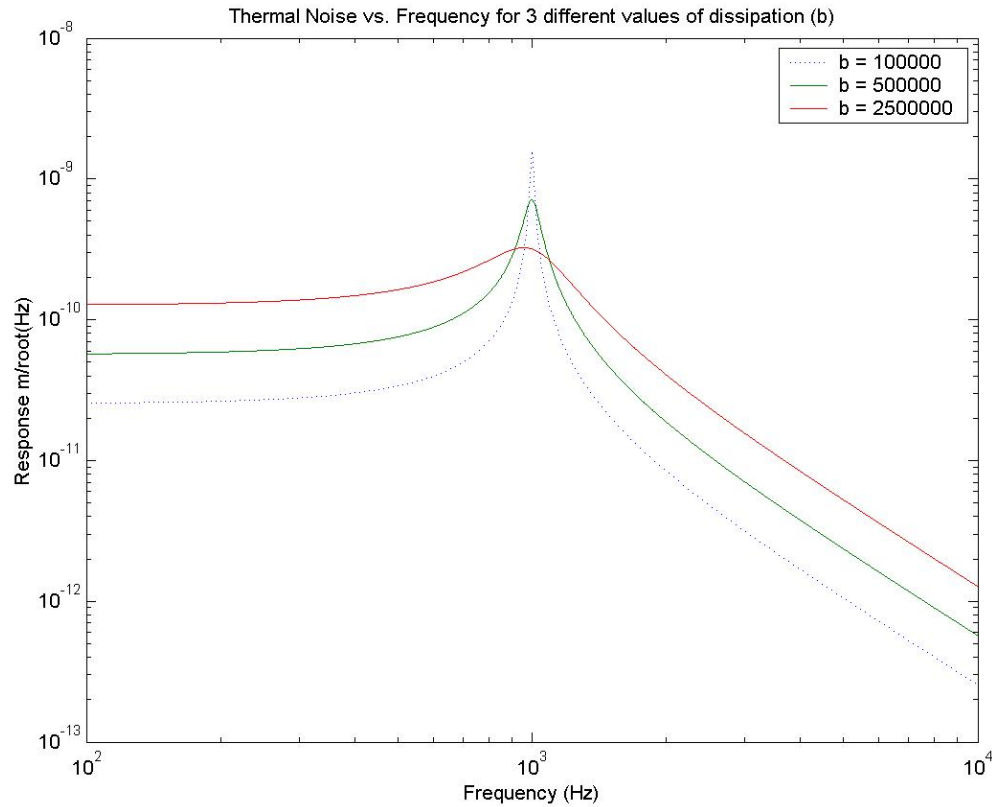
extent to which a test mass can remain at rest. In order to achieve the sensitivity necessary to detect gravitational waves, test mass thermal noise will have to be minimized. One way to accomplish this would be to go to very low temperatures, but this method creates a host of other technical difficulties, and should be avoided if possible. LIGO and most other interferometric gravitational wave detectors have decided instead to aggressively minimize thermal noise through careful choice of materials and suspension design.

In 1951-2 H.B. Callen *et al.* established a general relationship between fluctuation and dissipation known as the Fluctuation-Dissipation Theorem [7]. Using the Fluctuation-Dissipation Theorem, the thermal noise power spectrum for a pendulum can be written as,

$$x_{therm}^2(f) = \frac{k_B T b}{\pi^2 f^2 (b^2 + (2\pi f m - \frac{k}{2\pi f})^2)}$$

where b is the dissipation (i.e. the coefficient in the friction law $\vec{F}_F = -b\vec{v}$), m the mass of the pendulum, and k the effective spring constant for the pendulum. Thus the Fluctuation-Dissipation Theorem shows us that to reduce the level of dissipation is to reduce the level of off resonance thermal noise. To illustrate this consider figure 1.2.1.

Figure 1.2.1: Thermal Noise vs. Frequency for three values of dissipation



It may be seen referring to the equation or figure 1.2.1 that far from the resonance, the shape of the power spectrum is relatively constant, depending mainly on k at frequencies much lower than the resonance, and mainly on m at frequencies well above the resonance, and that the overall level of noise depends directly on b . The shape of the power spectrum near the resonance is primarily determined by the dissipation b , with low dissipation corresponding to a large resonance peak. Because of this, and equipartition theorem, which says that the integral of the thermal noise power spectrum over all frequencies is independent of dissipation, a lower dissipation translates into lower off resonance thermal noise. Thus to reduce off-resonance thermal noise

displacement spectral density in interferometer suspensions we must reduce the amount of dissipation.

To experimentally measure dissipation it is useful to introduce the concept of Q , a dimensionless measure of how small the dissipation is at the resonance frequency. Q is defined as the width of resonance,

$$Q \equiv \frac{f_0}{\Delta f},$$

where f_0 is the resonant frequency and Δf is the width of the resonance peak in the frequency response of the system measured at the level of half of the maximum power. A convenient way to measure Q , and the way that I have done it in this experiment, is to link it to the ringdown time τ by

$$Q = \pi f_0 \tau,$$

where τ is the time required for the amplitude to decrease by $\frac{1}{e}$.

Using these techniques we may measure the Q of various materials. Since Q is inversely proportional to dissipation, the higher the Q , the lower the thermal noise. Therefore, in order to improve the thermal noise of gravitational wave interferometers, we must install test masses and suspensions of the highest Q .


1.3 Interferometer Test Mass Suspensions

For interferometric gravitational wave detectors, the test masses must also act as mirrors. Therefore, it is important that the test mass material is chosen to have high optical quality as well as low thermal noise. Fused Silica, a high quality optical material,

has been shown to have a very high Q [8], and will therefore exhibit very low off resonance thermal noise. Because of this, and the fact that large, high quality fused silica optics are commercially available, fused silica has been chosen as the material for the optics in the first generation interferometers. Initial LIGO uses fused silica for the interferometer test masses, beam splitter, and other optics. Fused silica is also under continuing study as a fallback material for advanced LIGO should problems arise in the development of sapphire optics.

First generation LIGO uses fused silica test-masses suspended by steel wires [see Figure 1.3.1]. However, because fused silica has much lower thermal noise than steel, other detectors, such as GEO600, have switched to monolithic suspensions of fused silica in an attempt to minimize thermal noise.

Figure 1.3.1: LIGO Fused Silica test mass supported by metal wire sling



To create a monolithic suspension, it is necessary to attach fused silica suspension fibers to the test masses. Fused silica wires are very strong in tension, but their strength diminishes quickly with bending, so they cannot be fashioned into a sling like the metal wires. Welding them directly would damage the mirror surface or transfer heat to the bulk of the glass causing it to distort the light passing through, and thus must be avoided. A method is needed that will allow the suspension fibers to be attached without damaging the mirror surface, or introducing unnecessary thermal noise. Silicate bonding has been proposed as a possible method for advanced LIGO, and has already been implemented in the monolithic suspensions of GEO600 [see Figure 1.3.2]. In these suspensions, fused silica ears with sufficiently thin long and thin tips to allow welding and prevent the flow of heat into the bulk of the test mass are bonded to the side surfaces of the test masses [figure 1.3.3-4]. The effect of silicate bonding on thermal noise is now (here) being studied.

Figure 1.3.2: GEO600 monolithic suspension (Modified from figure by Professor Norna Robertson, 3/16/2001)





Figure 1.3.3-4: GEO600 ears bonded to mirror surface



1.4 Silicate Bonding

Silicate bonding is a bonding technique that was developed by Jason Gwo, of Stanford University, for assembling the fused quartz optical components of the Gravity Probe-B science instrument [9]. The technique is also applicable to fused silica and sapphire. It is a room temperature chemical bonding technique based on hydroxide catalysis. The bond is formed by bringing two polished surfaces in contact with one another in the presence of a hydroxide catalyst such as KOH or NaOH. The strength of the bond is comparable to that of fused silica, and has shown no degradation due to aging. Mechanical loss associated with silicate bonding has previously been shown to be low [10]. Because of this, as mentioned above, silicate bonding has already been implemented in the monolithic suspensions of GEO600, and is under consideration as a technique for Advanced LIGO. For a chemical understanding of silicate bonding refer to Gwo's paper [8].

Chapter II: Experiment


2.1 Sample Characterization

All of the samples were created at Glasgow University. They consist of 5mm fused silica rods, cut in half lengthwise, and polished to provide a flat along their length, with edges chamfered. The bonded samples were made by putting several drops of an aqueous solution of either KOH or NaOH on the flat surface of one D-rod, and then placing a second D-rod on top of the first such that the flat surfaces lined up with an overlap lengthwise. The samples were given sufficient time to allow the bonds to set.

They were stored at Syracuse University, individually wrapped in cloth, in a plastic container on the lab shelf.

Each sample was carefully measured, weighed, and scrutinized for scratches, areas where the bond did not take, and other defects that might contribute to thermal noise. During the experiment care was taken not to touch the samples, or unnecessarily damage them in any way. It was necessary to weld a thin fiber to each sample, but I tried to do this quickly and cleanly so as to minimize the damage done to the sample as well as the heat transferred to it. It should be noted that none of the samples were without defects, and that one of the bonded samples, silicabondC1, was bonded by a different method than the other bonded samples. The rest of this section consists of digital photos of the actual samples, and their physical characteristics.

Figure 2.1.1: SilicabondA1:



Total Length: 11.1 cm

Rod1 Length: 10.1 cm

Rod2 Length: 9.8 cm

Diameter: 0.5 cm

Bond Area: 4.4 cm²

Bond Catalyst: Potassium Hydroxide

Bonded on: 6/14/00 Glasgow University

Features: When seen from the side (as shown above) the bond is visible for its entire length, with some diffraction visible in the area of the bond. On each end there are small (0.2 cm long) V shaped scratches.

Figure 2.1.2: SilicabondB1:



Total Length: 11.7 cm

Rod1 Length: 10.1 cm

Rod2 Length: 9.8 cm

Diameter: 0.5 cm


Bond Area: 4.1 cm²

Bond Catalyst: Potassium Hydroxide

Bonded on: 6/15/00 Glasgow University

Features: When seen from the side (as shown above) the bond is visible for its entire length. There is a region approximately 1cm long, and 0.5cm wide near one end of the bonded area, that diffracts the light and appears mirrored at some angles. This may be a region where the bond did not take.

Figure 2.1.3: SilicabondC1:



Total Length: 11.6 cm

Rod1 Length: 10.0 cm

Rod2 Length: 10.0 cm

Diameter: 0.5 cm

Bond Area: 4.25 cm²

Bond Catalyst: Sodium silicate

Bonded on: (Date unknown, but later than '99) Glasgow University. It was transported from Glasgow University via pen case.

Features: There are several regions along the bonded area that diffract light and look like mirrored surfaces (one such region visible in the photograph about 6.5 cm from the left). These could be areas where the bond did not take. This sample has more visible defects than any of the other bonded samples.

SilicabondD1: (photo not available)

Total Length: 12.0 cm

Rod1 Length: 10.1 cm

Rod2 Length: 9.8 cm


Diameter: 0.5 cm

Bond Area: 4 cm²

Bond Type: Potassium Hydroxide

Bonded on: 6/16/00 Glasgow University

Figure 2.1.4: D1: Unbonded D-shaped rod of fused silica.



Total Length: 10.1 cm

Diameter: 0.5 cm

Bond Type: Unbonded

Features: There are some small dark scratches on each end surface.


2.2 Experimental Setup

This experiment aims to detect loss due to the internal friction of the samples.

Because of this, and in an effort to reduce systematic errors in the measurements, many steps were taken to reduce sources of extrinsic loss, or losses not due to the sample itself.

To reduce loss due to recoil of the support structure, the sample was suspended from a monolithic fused silica suspension consisting of a massive top bob, followed by a thin fused silica fiber welded to a massive intermediate bob, then another thin fiber connected by welding to the sample itself. This suspension functions as a series of harmonic oscillators, isolating the sample's vibration from the support. It is known as the double bob suspension [see Fig. 2.2.1], and has been employed as a recoil isolator in many thermal noise experiments at Syracuse University [8,11]. To further reduce recoil damping, this suspension was firmly clamped in a steel collet, which was attached to a massive aluminum plate supported by four aluminum legs, which were firmly bolted to a massive aluminum base plate that rested on and was bolted to the top of a large and heavy laboratory bench. Care was taken not to bump this table, or to, in any unintended way, excite the sample during measurements.

Figure 2.2.1: Experimental Setup



To reduce loss due to gas damping, The experiment was conducted inside a vacuum bell jar [see Fig. 2.2.2], and maintained the pressure at less than 1×10^{-5} torr for all measurements. A turbo molecular pump was used to attain vacuum, and pressure was measured using an ion gauge. Occasionally it was useful to break vacuum or increase the pressure significantly in order to use gas damping to damp out the pendulum mode of the sample. In these instances, care was taken to always allow the pressure to return to a reasonable vacuum before resuming measurements.

A comb capacitor [12], a device useful for exciting the normal modes, and thus hereafter referred to as the “excitor” was placed alongside the bottom of the sample for all measurements. Alternating wires of this excitor were given 500V DC voltage from a high voltage amplifier to induce a polarization in the sample, while the other wires were held at ground, to create a gradient. An AC voltage at the resonant frequency of the sample was then added to the DC voltage, to create a time varying force in order to excite the corresponding mode. Once the mode had been excited, or “rung up”, to an amplitude where it could be seen clearly above the noise, the excitor was grounded, using a switch box located on the lab bench, removing both AC and DC voltages, allowing the sample to ring down freely. It should be noted that throughout the experiment I used two different excitors, the one depicted in figure 2.2.1, and another consisting of wires wrapped horizontally around a vertical cylinder of the machinable ceramic called macor.

A laser and split-photodetector combination was used to monitor the amplitude of excitation of the sample. The laser was mounted on an x-y stage outside the bell jar, its beam shining through a focusing lens, then passing through the bell jar, and onto the top area of the sample, projecting a shadow on the photodetector mounted on a horizontal

translation stage outside the opposite side of the jar. As the sample resonates, the shadow moves back and forth over the two halves of the photodetector at the same frequency. The currents from each half of the photodetector are then compared using a differential current-to-voltage amplifier. The output voltage from this amplifier is channeled into a lock in amplifier, which shifts the rapid variations of the signal (typically $\sim 2\text{kHz}$) down to low frequencies, by comparing the input with a reference frequency, while preserving the time varying amplitude. When taking data, the output voltage from the lock-in amplifier is sent to a data acquisition board on a nearby computer, where it is recorded on hard disk. Throughout the experiment I used an oscilloscope to monitor both the voltage out of the current to voltage amplifier, and that out of the lock-in amplifier.

2.3 Experimental Procedure

The first step of the experiment was to find several normal modes of vibration above the $\sim 100\text{Hz}$ seismic wall, and below the $\sim 5\text{kHz}$ limit of the high voltage amplifier, for each sample. My original method for doing this was to use the frequencies calculated by Algor models. Algor is a finite element analysis program that is useful as a model, and capable of accurately predicting mode frequencies and mode shapes. I modeled each sample as a solid piece of glass, neglecting the bond, and used the Algor predicted frequencies and mode shapes as an estimate during the experiment. For a comparison of Algor predicted frequencies with the actual measured values see table 2.3.1. [see Table 2.3.1] For the first sample, I used the Algor-predicted frequencies as a ‘ballpark’ estimate, and made large and very time-consuming spectrum analyzer sweeps around

those frequencies. This method had such a poor success rate that I had to abandon it for a more expedient method called ‘feedback’.

Table 2.3.1: Algor-predicted frequencies vs. measured frequencies

Sample	Mode	Algor f (Hz)	Measured f (Hz)	% diff.
SilicabondA1	1 perp.	2444.7	2303.1	6.15
	1 par.	2454.6	2455.8	0.05
SilicabondB1	1 perp.	2188.9	2127.3	2.90
	1 par.	2239.4	2202.7	1.67
SilicabondC1	1 perp.	2312.62	2170.9	6.53
	1 par.	2338.46	2351.6	0.56
SilicabondD1	1 perp.	2100.32		
	1 par.	2174.9		
D1	1	1303.9	1310	0.47

Feedback is an acoustical method of exciting the normal modes of vibration, and therefore must be done outside of vacuum. The general idea is to use the output of the laser-split photodiode (after it passes through the current to voltage amplifier) as an input to some sort of noise generating device (in my case a radio with small speakers attached to it), and then acting on the sample with the sound generated from this system. If the noise excites any normal modes, its input is changed to have more power at the frequency of that mode, and thus the feedback loop continues driving the resonant mode to higher and higher amplitudes so that the frequency of the mode can be measured accurately using an FFT spectrum analyzer. When this method fails, it can be augmented with a technique we refer to as “priming the pump”, or simply giving the sample a gentle tap to excite a mode (usually evidenced by audible ringing), and then letting the feedback method take over and drive the resonance.

Experiment has shown [see Table 2.3.2] that for samples of similar geometry as mine and frequencies on order of $1kHz$, the frequency shift from air to vacuum is usually

less than $5Hz$, and thus after feedback one can do a quick $10Hz$ sweep to re-discover the mode in vacuum.

Table 2.3.2: Frequency in air vs. in vacuum

Sample	Mode	Air f (Hz)	Measured f (Hz)	% diff.
SilicabondA1	1 perp.	2299	2303.1	0.18
	1 par.	2453.7	2455.8	0.09
SilicabondB1	1 perp.	2118.8	2127.3	0.40
	1 par.	2200.5	2202.7	0.10
SilicabondC1	1 perp.	2170	2170.9	0.04
	1 par.	NA	2351.6	NA
SilicabondD1	1 perp.	2012.3	NA	NA
	1 par.	2051.1	NA	NA
D1	1	1309.25	1310.1	0.06

After the mode frequencies were known in air, the next step was to pump experiment down to a pressure below $1 \times 10^{-5} torr$. Once this was accomplished, acoustic feedback was out of the question. It was necessary to turn on the excitor, and sweep for modes using spectrum analyzer. All sweeps were conducted with the spectrum analyzer on swept sine mode, such that the spectrum analyzer supplied the high voltage amplifier with an AC signal, which was then sent to the excitor with an amplitude of $4.5V$. The sweeps were typically conducted over $10Hz$ regions centered on the frequency value as measured in air. Upon completion, the frequency response of the sweeps were examined for signs of resonance.

Once the frequency of the modes in vacuum were known, I would set the spectrum analyzer source frequency to the value of the peak from the sweep frequency response, and excite the sample using the electrostatic excitor. Typically the frequency would have shifted some since the time of the sweep due to temperature drifts, so after exciting the mode I would ground the excitor, and observe the FFT to pinpoint the exact peak of the resonance, then set the drive frequency to that, and excite the sample again, and so on,

until I could get close enough to the resonance frequency to excite the sample to a high enough amplitude for measurements to be taken (typically $>500\mu\text{V}$ at the output of the lock-in amplifier).

Once the sample was excited to a high enough amplitude, I removed the excitation. A grounding switch box was located on the lab bench so that with the flick of a switch I could ground the excitor, removing both AC and DC signals from it, allowing the sample to ring-down freely.

Immediately after the sample was grounded, the voltage output from the lock-in amplifier was sent to a data acquisition port on a nearby computer, and stored as to hard disk. For each measurement the number of data points, sampling rate, maximum amplitude at output of lock in amplifier, gain, vacuum pressure, and room temperature were recorded in the lab book.

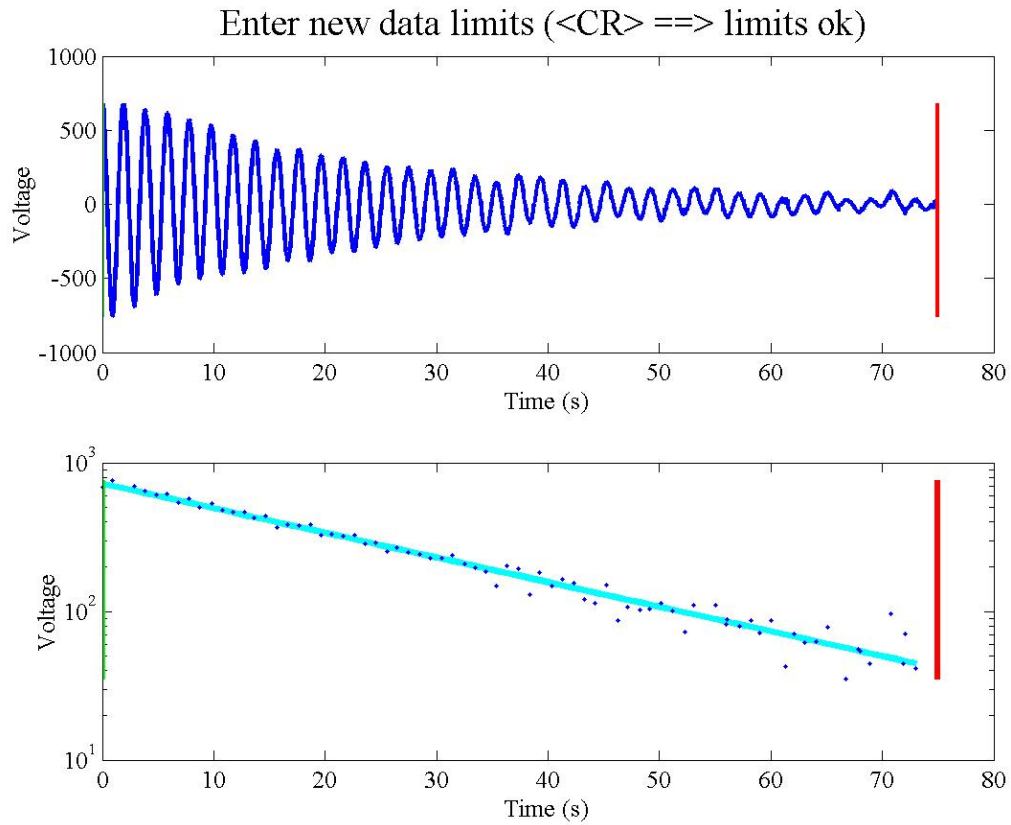
2.4 Data Analysis Procedures:

All ringdown data was fit using a Matlab program called Qfit, a revamped version of quickQ, both developed by Steven Penn of the Syracuse University experimental gravitation group. Qfit is designed to analyze data from a resonant ringdown and extract the time constant, τ , from the damped sinusoid $\text{Sin}(\omega t)e^{\frac{-t}{\tau}}$. The quality factor, Q , is then calculated from $Q = \pi f_{\text{mode}} \tau$.

Qfit allows the user to examine the raw data before fitting, and specify the time interval over which the data is fit. The procedure for analyzing a typical ringdown thus began by opening the ringdown voltage data in matlab [see Fig 2.4.1], and looking for spikes or other spurious activity in the data. Cuts were made to the data such that the fit

was applied to the cleanest portion of the data that still allowed the at least 2τ of time necessary for an accurate fit.


Figure 2.4.1: Typical ringdown data



After the initial fit the next step was to examine the residual [see Fig. 2.4.2] to get a qualitative idea of the accuracy of the fit. Large values for the residual indicate large differences between the fit and the data itself. When possible, further cuts to the data were made to ensure that the residual was small.

Figure 2.4.2: Qfit fit of the data with residual in red

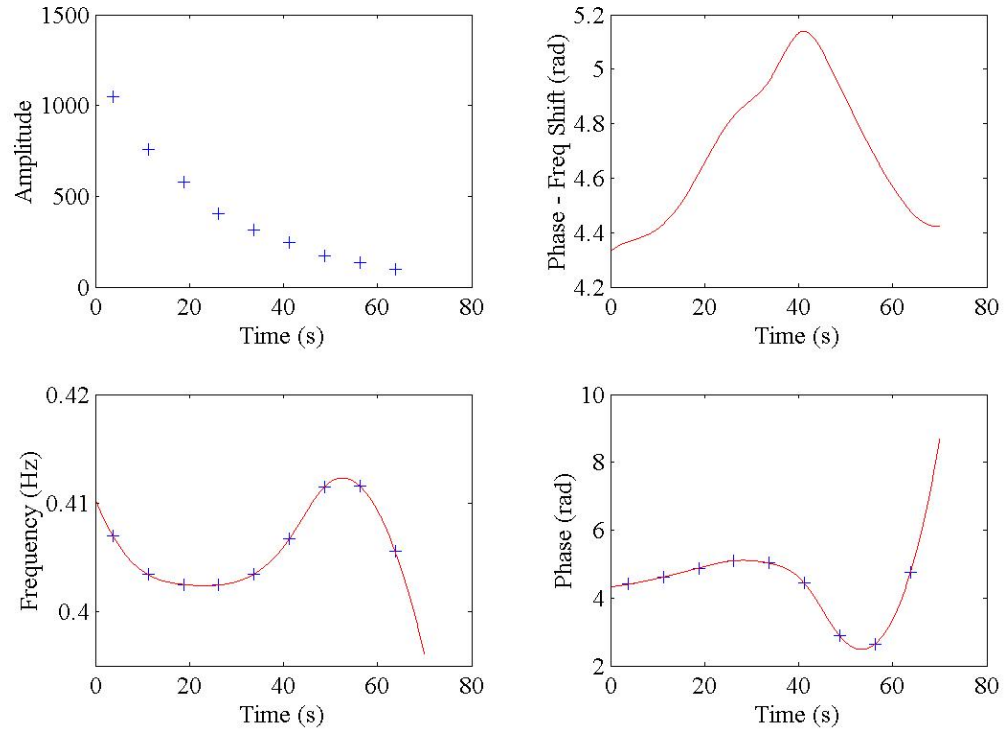
File: 021401a.dat, Freq = 2455.312Hz, Tau = 25.052s, $Q = 1.932e+005 \pm 1.528e+003$



At this point steps were also taken to ensure that the fit frequency was as accurate as possible to the resonant frequency of the sample. Qfit allows the user to examine the fit frequency change [see Fig. 2.4.3] and determine whether the fit frequency should be allowed to vary smoothly or be fixed. For most data, fits were made in both ways, and Q s were comparable. However, for some data, typically those of low Q , the fits were more accurate when the fit frequency was fixed, while for other data, typically those of higher Q , it was more accurate when allowed to vary smoothly. This could be due to temperature drifts that act to change the sample's mode frequencies over the long time duration of high Q measurements, but have little or no effect on the mode frequencies of short measurements.

Figure 2.4.3: Frequency and phase shift throughout the ringdown


Fit Constants vs Time: File = 021401a.dat, Freq = 2455.312Hz



After making the appropriate cuts and determining whether to allow the frequency to vary, the final fit is made, and the program prints out two values of Q . The first Q value is a rough estimate, called “Quick Q ”, generated by fitting to the envelope [see Fig. 2.4.4], and the second, which I call “ Q fit”, is generated by fitting a damped sinusoid to the data [refer back to Fig. 2.4.2]. It was useful to examine both Quick Q and Q fit to check that they were comparable. All Q s in this thesis are those generated by Q fit.

Figure 2.4.4: Quick Q fit of the data

File: 021401a.dat, Freq = 2455.312Hz, Tau = 24.256s, Q = 1.871e+005 ± 3.147e+003



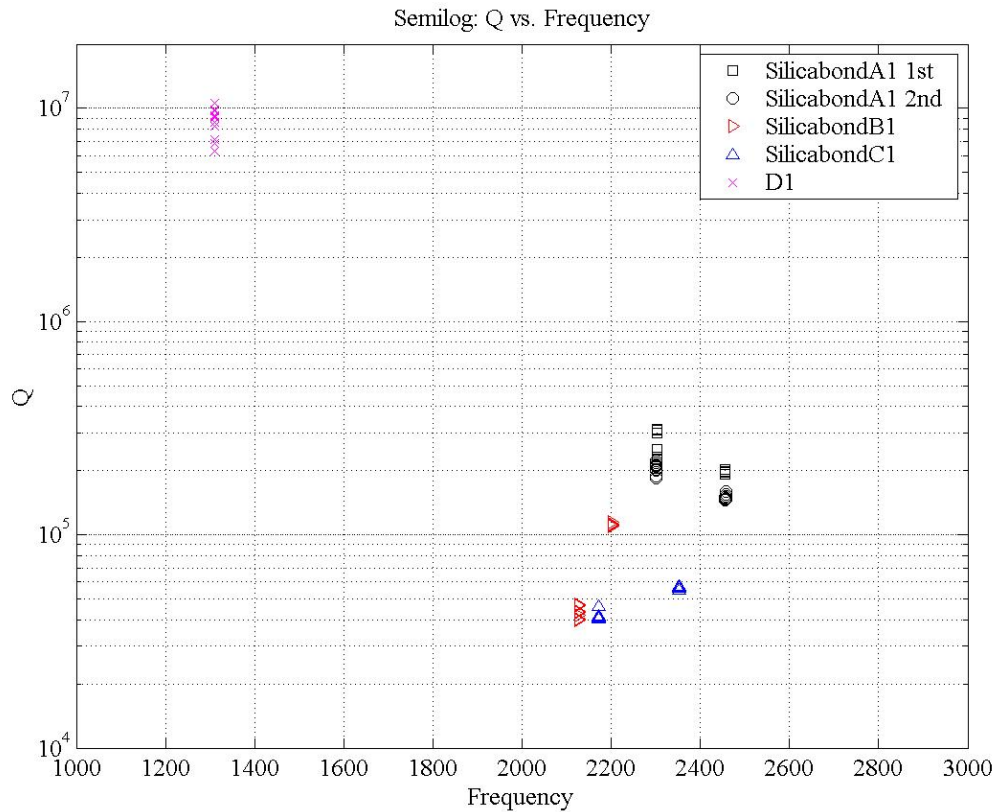
Chapter III: Results

3.1 Q Measurements

Q measurements were taken on the two lowest frequency modes of vibration of each bonded sample, and on the single lowest frequency mode of vibration for the unbonded sample. The measurements show a significant reduction of Q in the bonded samples when compared to the unbonded sample. Since excess loss always acts to reduce the measured Q , the best indicator of the true internal friction of a sample is the quality factor of the highest Q mode over all modes and all hangings. The unbonded sample, D1, has a high Q of $1.055 \times 10^7 \pm 4.051 \times 10^3$, while the highest Q recorded in a

bonded sample was $3.098 \times 10^5 \pm 5.743 \times 10^3$, a difference of nearly two orders of magnitude. Figure 3.1.1 is a plot of the data from all samples with Q on a log scale plotted against frequency. Each point on the graph represents the Q calculated by Qfit.m for a given measurement.


Figure 3.1.1: Semi-log graph of Q versus frequency for all samples



On the graph there are two different frequencies represented for each of the bonded samples. This is a result of the fact that the first two modes of the samples, which would actually be degenerate in the case of a rod, have different frequencies. According to the Algor mode shape predictions, the lower frequency mode for a given sample is the mode with displacement occurring across or perpendicular to the plane of the bond, with the higher frequency mode corresponding with mode displacement parallel to the plane of

the bond [see Fig. 3.1.2]. Hereafter I will refer to these mode shapes as simply perpendicular and parallel respectively. The displacement of the first mode of the unbonded sample is perpendicular to the flat surface.

Figure 3.1.2: Algor-predicted mode shapes (from left to right: undisplaced bonded sample, Parallel mode, perpendicular mode, first mode for D-rod)




The results of this experiment show a surprising similarity in Q between parallel and perpendicular modes. We had expected the Q of the perpendicular mode to be much lower than that of the parallel with the assumption that the bonded region would be a

neutral plane during motion parallel to its surface. Research is currently underway to determine how energy is distributed in the modes in this geometry, and in particular, how involved the bonded region is in determining the loss in each mode.

During the course of the experiment some attempts were made, with little success, to verify whether the displacement of the rods was truly parallel or perpendicular to the bond surface. The samples were always hung with their bond surfaces facing either parallel to or perpendicular to the laser beam. It was possible to excite both modes in the same orientation, at times one better than the other, but not consistently so.

Figure 3.1.3 is the same plot as figure 3.1.1 except focused on a much narrower linear scale of Q to give a better perspective of the spread between the Q s of the bonded samples. There are also errorbars (although I admit they are not easy to see on this plot) representing the error calculated by the Q fit program. The spread of Q s was typically small, but it is interesting to note that not all measurements on a given frequency of a given sample fit nicely within each other's errorbars. Note also the drop in Q of SilicabondA1 from the first hanging to the second hanging. This may be due to the welding procedure.

Figure 3.1.3: Q with errorbars representing error as calculated by Qfit.m versus frequency



To get another idea of the error associated with these measurements I plotted the mean value for a set of measurements with errorbars corresponding to the standard deviation of the Q within that set [see Fig. 3.1.4]. The standard deviation was typically small with the exception of the first hanging of the 2303Hz perpendicular mode of SilicabondA1. By comparison of figure 3.1.3 and figure 3.1.4 it may be seen that the standard deviation among a set of measurements is typically larger than the error calculated by Qfit for a given measurement. This may be because of factors, such as temperature or different amplitudes of excitation, that change with time between measurements, making the Q s within a given set less self consistent.

Figure 3.1.4: Mean Q versus frequency

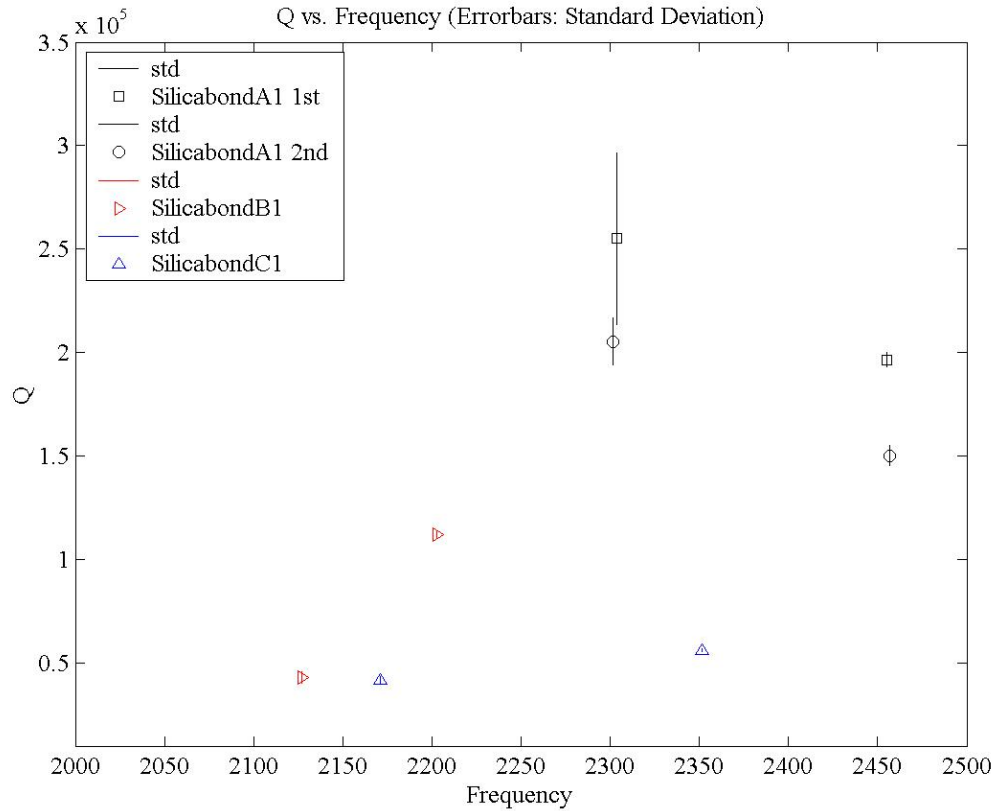
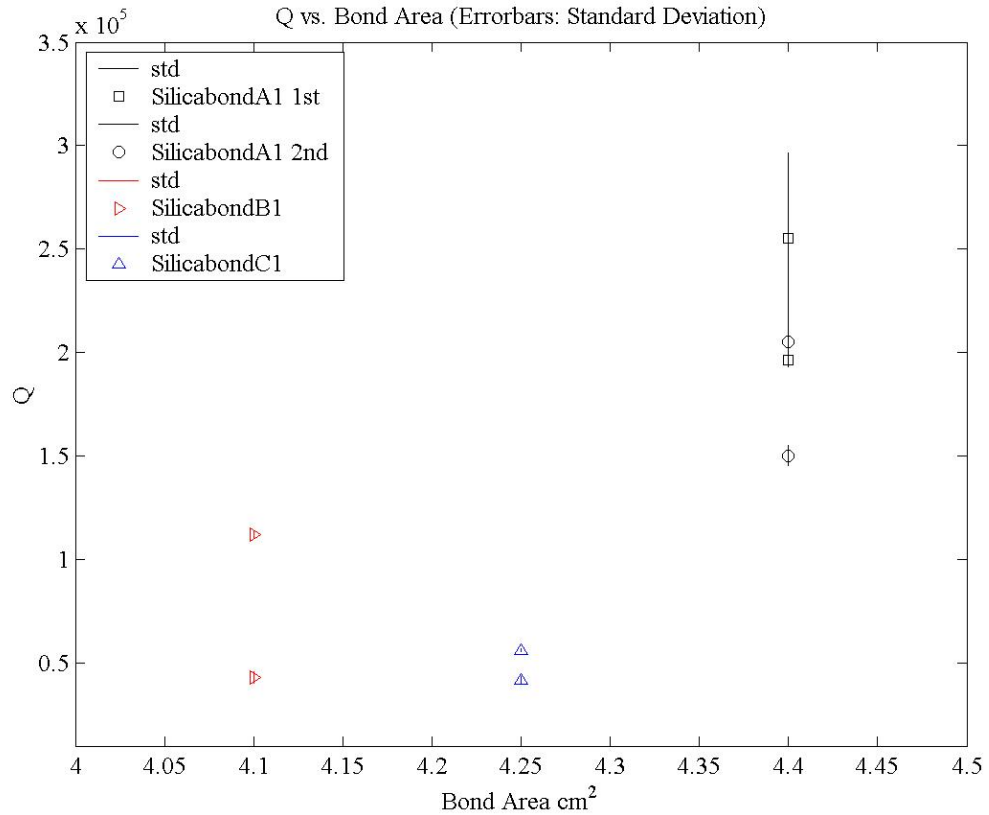


Figure 3.1.5 plots the mean Q with standard deviation errorbars again, but this time versus bond area. This plot shows no clear correspondence between bond area and Q for the bonded samples. This is not too surprising, considering that the differences in bond areas among the samples are quite small, ranging only from $4.1cm^2$ to $4.4cm^2$. It should also be noted that the samples were all unique in more ways than bond area. Each sample had some defects that also might have contributed to the Q depreciation, and one sample was actually bonded using a different chemical.

Figure 3.1.5 Mean Q versus bond area



The most interesting Q to bond area comparison is that of the Q s from the bonded samples to that of the unbonded D-rod. In the next section this comparison is scaled to estimate for the effect of silicate bonding ears to the sides of LIGO sized test masses.

3.2 Scaled Results

For the purposes of scaling, I rely on the idea that all sources of dissipation act to reduce Q , and thus focus my calculations on the highest Q measured for both the bonded and unbonded case. The highest Q measured in the unbonded D-rod was

1.0×10^7 , and the highest Q measured in a bonded sample was 3.1×10^5 measured in SilicabondA1.

The fact that the Q of the unbonded sample is much greater than that of the bonded sample indicates that the dominant source of loss in the bonded case is the bond itself. Let us assume that the bond is the dominant source of loss, and scale the loss from the small bonded samples measured here to predict the Q of the test masses that will be used in advanced LIGO.

The Q of the bonded sample is,

$$Q_{\text{Sample}}^{\text{Bonded}} = 3.1 \times 10^5 \quad \text{when} \quad \left(\frac{V}{A}\right)_{\text{Sample}}^{\text{Bonded}} = \frac{2.0 \text{ cm}^3}{4.4 \text{ cm}^2}.$$

Here V is the volume of the bonded sample (i.e. the composite volume of the two D-rods), and A is the area of the bond. For a 40kg Advanced LIGO test mass made of fused silica ($\rho = 2.202 \frac{\text{g}}{\text{cm}^3}$),

$$\left(\frac{V}{A}\right)_{\text{AdvLIGO}} = \frac{1.8 \times 10^4 \text{ cm}^3}{10 \text{ cm}^2},$$

where V is the volume of the LIGO mirror as calculated by $\frac{m}{\rho}$, and A is an approximate value for the total bond area from four ears bonded to the sides of the mirror. Thus we would predict that an Advanced LIGO mirror with loss associated with the bond as its dominant source of loss to have Q of

$$Q_{\text{Bond}}^{\text{AdvLIGO}} = Q_{\text{Sample}}^{\text{Bonded}} \frac{\left(\frac{V}{A}\right)_{\text{AdvLIGO}}}{\left(\frac{V}{A}\right)_{\text{Sample}}^{\text{Bonded}}} = 1.2 \times 10^9.$$

This Q is nearly two orders of magnitude greater than the measured bulk Q of fused silica test masses, $Q_{bulk} = 2.0 \times 10^7$ [Q mes. Pdf]. If we were to then predict the Q of a fused silica test mass with losses associated with contributions from both bulk loss and loss associated with the bond,

$$\frac{1}{Q_{AdvLIGO}} = \frac{1}{Q_{bulk}} + \frac{1}{Q_{AdvLIGO}^{Bond}},$$

It is easy to see that the role of the bonds in determining the loss is negligible compared to that of the bulk.

These estimates predict that silicate bonding should not be the limiting thermal noise source for Advanced LIGO test masses. The influence on thermal noise should be even lower than predicted here because the bonds will be on the sides of the mirrors, far away from the mirrored surfaces where thermal noise is most important. Therefore, silicate bonding should be a suitable technique for the creation of monolithic suspensions for Advanced LIGO.

Bibliography

- [1] J. Weber, *Phys. Rev. Lett.* **22** (1969) 1320.
- [2] R. A. Hulse and J. H. Taylor, *Ap. J. Lett.* **195** (1975) L51.
- [3] A. Abramovici *et al.*, *Science* **256** (1992) 325.
- [4] A. Giazotto *et al.*, *Nucl. Instr. Meth. Phys. Res.* **A289** (1990) 518.
- [5] K. Danzmann *et al.*, GEO600: Proposal for a 600 m laser interferometric gravitational wave antenna, Max-Planck-Institut fuer Quantenoptik Report 190, Garching, Germany (1994).
- [6] K. Tsubono, *Gravitational Wave Experiments*, Proceedings of the First Edoardo Amaldi Conference (World Scientific, Singapore, 1995), p. 112.
- [7] H. B. Callen and T. A. Welton, *Phys. Rev.* **83** (1951); H. B. Callen and R. F. Greene, *Phys. Rev.* **86** (1952) 702.
- [8] S. D. Penn, G. M. Harry, A. M. Gretarsson, S. E. Kittelberger, P. R. Saulson, J. J. Schiller, J. R. Smith, S. O. Swords, *Rev. Sci. Instrum.* **72** (2001) 3670-3673.
- [9] Dz-Hung Gwo, Ultra-precision bonding for cryogenic fused-silica optics, Stanford (1998).
- [10] S. Rowan, S. M. Twyford, J. Hough, D.-H. Gwo, R. Route, *Phys. Lett. A* **246** (1998) 471-478.
- [11] A. M. Gretarsson and G. M. Harry, *Rev. Sci. Instrum.* **70** (1999) 4081-4082
- [12] A. Cadez and A. Abramovici, *J. Phys.* **E21** (1998) 453.
- [13] P. R. Saulson, *Fundamentals of Interferometric Gravitational Waved Detectors* (World Scientific, Singapore, 1994).