

Measurement of Thermo-Optic Properties of Thin Film Dielectric Coatings

Thesis by

Gregory H. Ogin

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2013

(Defended August 10, 2012)

Acknowledgements

First and foremost, I would like to thank my parents for spending the last 30 years reinforcing the idea that my education is a top priority for them (and in turn instilling it as a top priority for me).

Thanks to all the great teachers I've had, from elementary school through graduate school, for pushing me forward and putting up with my antics.

Thanks to my great advisers at the University of St. Thomas, Los Alamos National Laboratory, and Caltech, who kept me going through the tough parts of research and truly set me up to succeed in the world of science.

Thanks to my friends and loved ones, both in school and out, who kept me sane throughout the process.

And last but not least, thanks to Caltech and the National Science Foundation (award PHY-0757058) for supporting me financially through the past seven years of graduate school and enabling this thesis.

Abstract

We have carried out an experiment to test the theory of the thermo-optic response of a dielectric stack mirror coating and to measure parameters of interest in calculating thermo-optic noise. Specifically, we measured the coefficient of thermal expansion and the change of index of refraction with temperature (dn/dT) for thin film silica (SiO_2) and tantala (Ta_2O_5) in mirror coatings. These measurements were achieved by driving thermal fluctuations in such mirrors in one arm of a small Michelson interferometer. We report on the results of that experiment along with its potential implications for future gravitational wave detectors, and suggest next steps for this important line of investigation.

Contents

Acknowledgements	iii
Abstract	iv
1 Introduction	1
1.1 Gravitational Waves	1
1.1.1 Wave Solutions to Einstein's Equations	1
1.1.2 Effects on Matter	2
1.2 Signal Strength / Expected Sources	2
1.3 LIGO Interferometers	3
1.4 The Advanced LIGO Noise Curve	4
2 Coating Thermo-Optic Noise Theory and Background	8
2.1 Calculating Temperature Fluctuations	8
2.2 Thermo-Elastic Expansion of the Coating	10
2.3 Thermo-Refractive Coefficient of the Coating	11
2.3.1 Quarter-Wave Full Derivation	11
2.3.2 1/8-3/8 Calculation	18
2.4 Previously Published Parameters	20
2.5 Our Measurement	23
3 The Experiment	27
3.1 Introduction	27
3.2 Theory	28

3.2.1	Mirror Geometry	29
3.2.2	1-D Approximation	30
3.2.2.1	Temperature Profile	30
3.2.2.2	Thermo-Refractive Response	31
3.2.2.3	Substrate Thermo-Elastic Response	32
3.2.3	Effects of Penetrating Radiation in 1-D Approximation	32
3.2.3.1	Solving the Heat Equation	33
3.2.3.2	Results	37
3.2.3.3	Conclusions	39
3.2.4	Low Frequency Green's Function Calculation	40
3.2.5	Combined Signal	44
3.3	Simulation	44
3.3.1	Poisson Ratio Correction	47
3.3.2	Boundary Conditions	52
3.3.2.1	Finite Boundaries	52
3.3.2.2	Clamping Pressure	53
3.4	Materials and Methods	54
3.4.1	Mirrors and Coatings	54
3.4.2	Experimental Layout	55
3.4.2.1	Optical Probe Layout	55
3.4.2.2	Optical Pump Layout	57
3.4.2.3	Data Readout Electronics	59
3.4.2.4	Locking Electronics	60
3.4.2.5	The Servo Loop	60
3.4.3	Rough Calibration	63
3.4.3.1	Michelson Readout	65
3.4.3.2	Pump Power Calibration	67
3.4.3.3	Pump Beam Spot Size	74
3.4.4	Sapphire Substrate Response Calibration	77
3.4.4.1	Intensity via LF Substrate Response	77

3.4.4.2	Sapphire Response Calibration Function	77
3.4.5	Data Acquisition	80
3.4.5.1	Laser Alignment	80
3.4.5.2	Warming Up and Pumping Down	81
3.4.5.3	Pre-Run Data	82
3.4.5.4	Data Acquisition Routine	82
3.5	Results	83
3.5.1	Data Processing	83
3.5.2	Combined TE/TR	85
3.5.3	TE Alone	87
3.5.4	Error Analysis	89
3.6	Poisson Factor Correction	93
4	Parameter Extraction	96
4.1	The Measurement Matrix	96
4.2	Error Analysis	98
4.2.1	Statistical Error	98
4.2.2	Known Unknowns	99
4.3	Results	100
5	Conclusions and Future Directions	102
5.1	Conclusions	102
5.1.1	Comparison with Published Values	102
5.1.2	Possible Implications for Future Gravity Wave Detectors	103
5.2	Future Work	107
5.2.1	Decrease Noise in Gold-Coated Measurements	107
5.2.2	Coating Design for Increased Signal	108
5.2.3	Boundary Condition and Poisson Ratio Effects	109
5.2.4	Advanced LIGO Coating Measurements	109
5.2.5	Further Coating Studies	110

Chapter 1

Introduction

1.1 Gravitational Waves

1.1.1 Wave Solutions to Einstein's Equations

In 1916 Albert Einstein gave us our current theory of gravity, General Relativity (GR). In contrast to the previous Newtonian theory, which involved instantaneous action (force) at a distance and took place on a background of flat Euclidian space, Einstein's theory dispensed with the idea of force and instead described the geometry of a four-dimensional spacetime. In Einstein's theory, objects acting solely under the influence of gravity move in freely-falling trajectories through a curved spacetime, and the spacetime takes its curvature due to the presence of energy (the most concentrated form people are familiar with being mass). Given a distribution of energy and momentum, the curvature of spacetime and the metric are determined by Einstein's field equations

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi GT_{\mu\nu} \quad (1.1)$$

where $T_{\mu\nu}$ on the right side is the *stress-energy tensor*, and the left side contains the spacetime metric g and the Ricci tensor and scalar which are contractions of the spacetime curvature. For more details see [7].

In the weak-field regime, we can write the metric as a background flat metric η plus a small perturbation:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.2)$$

where we will assume that h is small enough that terms of order h^2 and higher can be neglected. Using this definition along with the choice of a convenient gauge (the so-called *transverse traceless gauge*) we find that solutions to equation 1.1 in free space (with $T_{\mu\nu} = 0$) obey the differential equation

$$\square h_{\mu\nu} = 0 \tag{1.3}$$

which is recognized as the wave equation, supporting transverse waves in the metric perturbation which propagate at the speed of light.

1.1.2 Effects on Matter

By altering the spacetime metric, these waves impart an oscillating transverse quadrupole strain on space as they go by. The fact that the motion takes the form of a strain means that objects separated by a certain distance L will appear to move apart or together by a fraction h of that distance so that the displacement observed is hL , where h is the perturbation mentioned in equations 1.2 and 1.3. The quadrupole nature of this radiation means that as space expands maximally in one direction perpendicular to the direction of propagation, it will be contracting maximally in the other direction perpendicular to the direction of propagation, and vice versa in the following half cycle.

Due to the constraints imposed by the transverse traceless gauge and the field equations, $h_{\mu\nu}$ only has two linearly independent traveling modes, which are referred to as “plus” and “cross” modes. If we align our transverse axes so that the plus modes stretch and squeeze along the \hat{x} and \hat{y} directions in the manner described above, then the cross mode will stretch and squeeze in the $\hat{x} + \hat{y}$ and $\hat{x} - \hat{y}$ directions.

1.2 Signal Strength / Expected Sources

The source of gravitational waves is a mass/energy distribution with a time-changing quadrupole moment. The canonical example of a time-changing quadrupole moment

is the spinning bar-bell, or two massive objects orbiting each other.

For a pair of masses with a quadrupole moment Q that changes with time, the observed effects of emitted gravitational waves at a distance r away from the source will be on the order of

$$h \sim \frac{G}{c^4} \frac{\ddot{Q}}{r} \quad (1.4)$$

Unfortunately, since Newton's constant G is a small number and the speed of light c is a large number, gravitational wave strains are extremely small. There are no sources on Earth or in our solar system that will create waves that we will detect any time soon. Our best bet is looking for waves caused by neutron star and black hole inspirals and collisions.

LIGO's standard candle is the few solar mass neutron star binary system, which in its final merger phase (when the gravitational waves are at a frequency of many 10s of Hz to a few kHz) give gravitational wave strains on Earth of the order $h \simeq 10^{-21}$. This is the level of sensitivity LIGO has achieved, and Advanced LIGO (AdLIGO) hopes to beat.

1.3 LIGO Interferometers

The quadrupole radiation pattern that gravitational waves assume is perfect for detection with a traditional Michelson interferometer, which uses the interference of light to measure the path length difference in two perpendicular arms. Since one arm is lengthened while the other is shortened, the difference in the arm lengths is twice the distance by which either arm is changed. If we have such an interferometer where both arms are initially the same length L , then the measured length change (presuming of course that the arms are aligned along the polarization of the incoming wave) is

$$L'_2 - L'_1 = L(1 + h) - L(1 - h) = 2Lh \quad (1.5)$$

The LIGO detectors are, at their heart, Michelson interferometers. The arms contain Fabry-Perot cavities, and there are power recycling and resonant signal extraction techniques, but these are essentially techniques for more sensitively measuring the distance to the mirrors at the ends of the 4 km arms of the interferometer.

With arm lengths of 4 km, the instruments need to be sensitive to mirror motion on the order of 10^{-19} m, or $1/10,000^{\text{th}}$ the width of a proton. Acoustic vibrations move objects in air many orders of magnitude more than this, so the mirrors are held in vacuum. Seismic motion of the ground moves objects many orders of magnitude more than this, so there is aggressive passive and active isolation of the mirrors. Microscopic random motion of the suspensions that hold the mirrors in place can be devastating if the supports aren't carefully designed and materials carefully selected and pure. A huge number of noise sources must be carefully accounted for, controlled, and minimized. Some of the most challenging sources arise from the very surfaces of the mirrors whose position is being measured.

1.4 The Advanced LIGO Noise Curve

As one can see in Figure 1.1, the AdLIGO interferometers will be pushing down to low enough strain sensitivities that we believe the first direct detection of gravitational waves is likely. Also, the interferometers' main noise source will be quantum in origin over a wide range of frequencies, a very theoretically interesting region comprised of radiation pressure at low frequencies and shot noise at high frequencies. Observation of quantum behavior in ~ 40 kg test mass mirrors is truly a spectacular achievement on its own. The quantum noise floor is spoiled by Brownian motion of the mirror's surface in the interferometer's most sensitive frequency range. This noise source is of pressing concern to LIGO but will not be addressed here.

The noise source we address in this thesis is the thermo-optic noise, shown on a slightly less cluttered noise curve in Figure 1.2 with two different theoretical treatments in brown dashed line and green solid line. The entire purpose of the paper in which this graph appeared [23], on which I was a co-author, was to introduce this

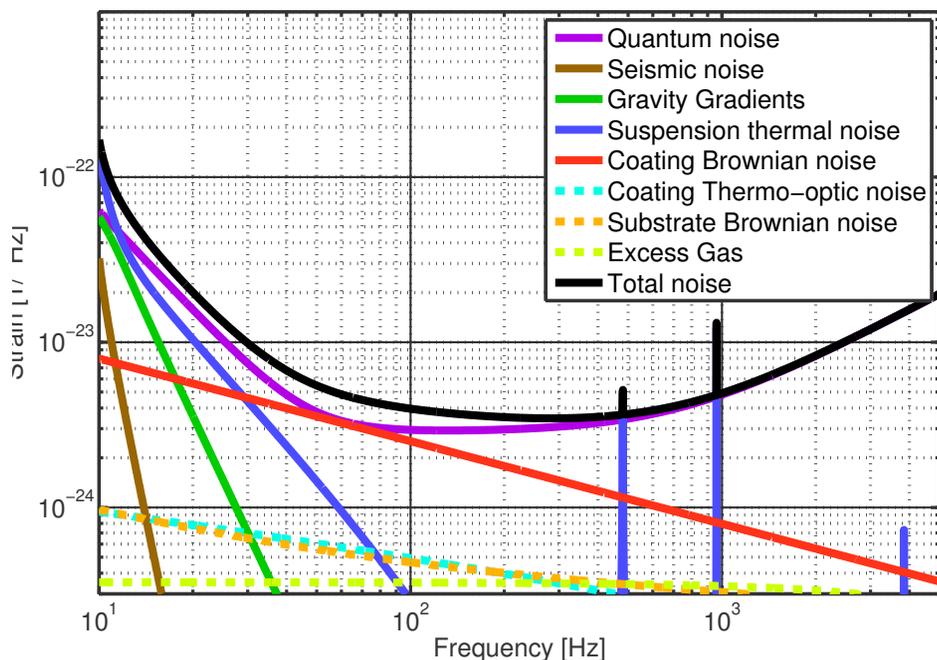
AdvLIGO Noise Curve: $P_{\text{in}} = 125.0 \text{ W}$ 

Figure 1.1: An Advanced LIGO Predicted Noise Curve

new theory to supplant the old theory. Thermo-optic noise is the result of local temperature fluctuations (or more accurately fluctuations of local energy density) in the mirror coating, and is comprised of two components. These are typically broken up as “thermo-elastic” (TE), or the expansion of the coating with temperature change, and “thermo-refractive” (TR), which is the change in the optical properties of the coating with temperature, which in turn change the reflection coefficient of the multilayer mirror coating.

As a side note, this nomenclature is slightly unfortunate in the following way. The term thermo-refractive was chosen because as the index of refractions change the layer reflection coefficients change and the layers’ optical path lengths change — this index of refraction dependence is why the effect was called thermo-refractive, but it must be emphasized that the optical path length of a layer depends on *both* its index of refraction and its thickness. The reflection properties of the stack will change when layer thicknesses change, *in addition to* the fact that the surface of the coating will move. Thus we will keep the TE and TR nomenclature, but note that there is a

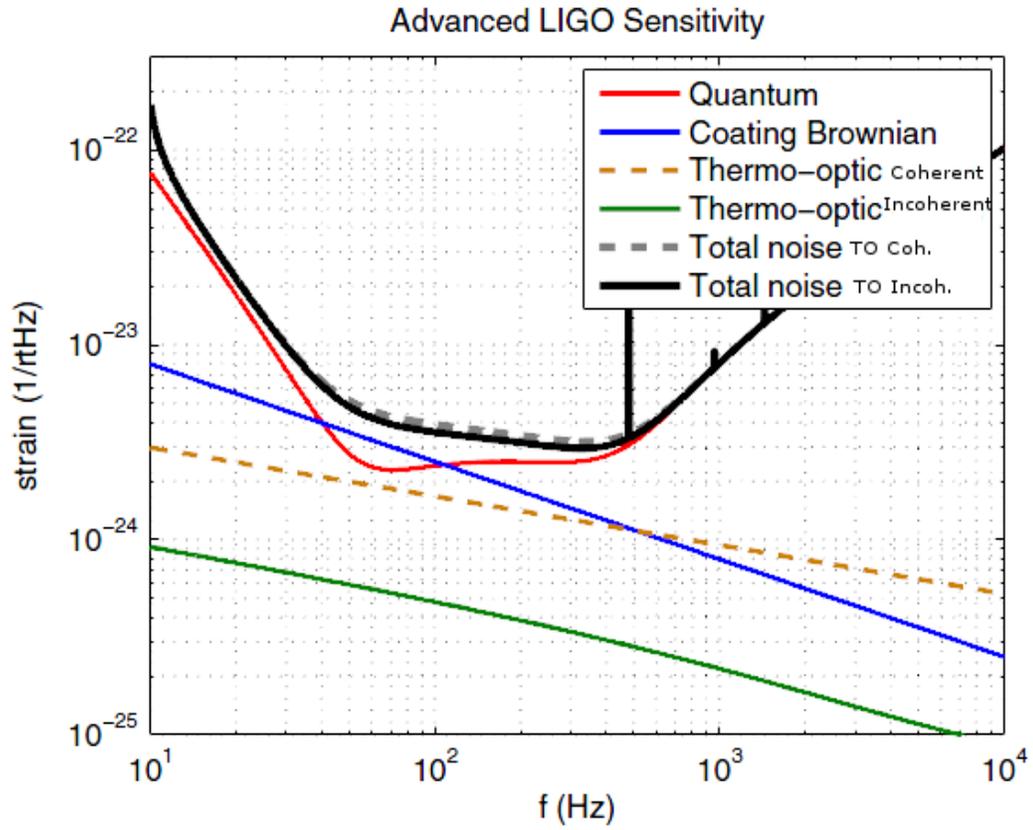


Figure 1.2: An Advanced LIGO predicted noise curve from Evans et al[23] with thermo-optic, Brownian, and quantum noise highlighted.

”TE-like” component to the TR effect.

In previous analyses, noise power spectral densities for TE and TR effects had been calculated separately and then added as incoherent noise sources. We argued in [23] that these effects were both driven by the same underlying random temperature fluctuations and thus needed to be added coherently. What’s more, given the expectation that the coefficients involved (the coefficients of thermal expansion α and the changes of index of refraction with temperature β) would be positive, we showed theoretically that they should partially cancel. As the temperature increased, the surface of the mirror would expand into the oncoming beam, making it look like the light from the measurement laser had traveled a shorter distance, but the complex reflection coefficient of the coating would change in such a way that it would appear on reflection that the light had traveled a longer distance. This is the difference in the OLD versus NEW thermo-optic noise curves in figure 1.2.

This result was good for the LIGO community; it meant that a noise source we had previously kept in consideration (especially given the fact that the exact parameters were not well known) was now likely to be an order of magnitude smaller than we would care about. However, we observed that a) this theory had not been experimentally tested, and b) measured values for the parameters in question showed great variation. One paper even had one of them listed as negative rather than positive - if this were the case, there was an outside chance that the TE and TR effects would *add* coherently rather than partially cancel. This would be bad news since it would indicate the NEW noise would actually be larger than the OLD noise.

Given this discrepancy and uncertainty, we set out to test the theory and measure the parameters in an environment as close to the LIGO mirror coatings as we could manage. The theory, our experiment, and its results are described in the rest of this thesis.

Chapter 2

Coating Thermo-Optic Noise Theory and Background

The dielectric mirror coatings in the LIGO experiment add their own noise to the interferometer due to statistical fluctuations. These fluctuations' influence on the gravitational wave signal are most commonly broken down in two categories: mechanical fluctuations in density/volume of the coating, commonly referred to as "Brownian Noise", and motion or apparent motion that is the result of fluctuations in the local effective temperature of the coating, commonly referred to as "Thermo-Optic Noise". This thesis will be concerned with the latter, and measurement of the parameters that govern the translation of thermal fluctuations into motion (and apparent motion) of the mirror.

2.1 Calculating Temperature Fluctuations

The Fluctuation-Dissipation Theorem of Callen and Welton [6] provides a means for calculating the fluctuation of a coordinate given knowledge of the loss induced by applying a force that will drive that coordinate. Following Levin [19], we consider the injection of entropy, which is the thermodynamic conjugate that drives temperature, with a profile weighted according to the amount of optical power we see at that part of the mirror.

$$\frac{\delta s}{dV} = \frac{1}{T} \frac{\delta Q}{dV} = F_0 \cos(2\pi ft) q(\vec{r}) \quad (2.1)$$

where F_0 is the magnitude of a dummy force that will be divided out, f is the frequency of the noise component, and $q(\vec{r})$ is a form factor that in a sense aggregates the fluctuations of different parts of the mirror into their effects on the single parameter we are measuring and want to know the fluctuations of. Braginsky [25] identified the coordinate of interest (temperature weighted by the power profile of the measurement beam) to be

$$\delta\hat{T} = \frac{1}{\pi r_0^2 l} \int_{-\infty}^{\infty} dx dy \int_0^{\infty} dz \delta T(\vec{r}, t) e^{-(x^2+y^2)/r_0^2} e^{-z/l} \quad (2.2)$$

with r_0 the radius of the measurement beam and l the effective coating thickness. From this we deduce (Levin [19] equation 14) that the appropriate form factor q will be

$$q(\vec{r}) = \frac{1}{\pi r_0^2 l} e^{-(x^2+y^2)/r_0^2} e^{-z/l}. \quad (2.3)$$

Injecting entropy in the form of equation 2.1, we can solve the heat equation for the temperature profile and use the fact ([18], [20], [19]) that the dissipated energy is given by

$$W_{diss} = \int d^3r \frac{\kappa}{T} \langle (\delta T)^2 \rangle \quad (2.4)$$

with κ the thermal conductivity of the medium, and $\langle \dots \rangle$ indicating a time-average over one full period of the injection oscillation.

Performing this calculation and inserting the result into the appropriate formulation of the fluctuation-dissipation theorem gives the resulting power spectral density of temperature-coordinate fluctuations [19] [25]:

$$S_{\delta\hat{T}}(f) = \frac{\sqrt{2} k_B T^2}{\pi r_0^2 \sqrt{\omega C \rho \kappa}}. \quad (2.5)$$

Equation 2.5 gives us an appropriately weighted temperature variation of the surface of the mirror for use in calculating the effects of these random fluctuations on the apparent position of the mirror. In the next section, we will describe the exact

Thermo-optic Noise: TO = TE + TR

- **Thermo-Elastic (TE):** Mirror's surface expands into probe beam. By convention, negative $d\phi/dT$

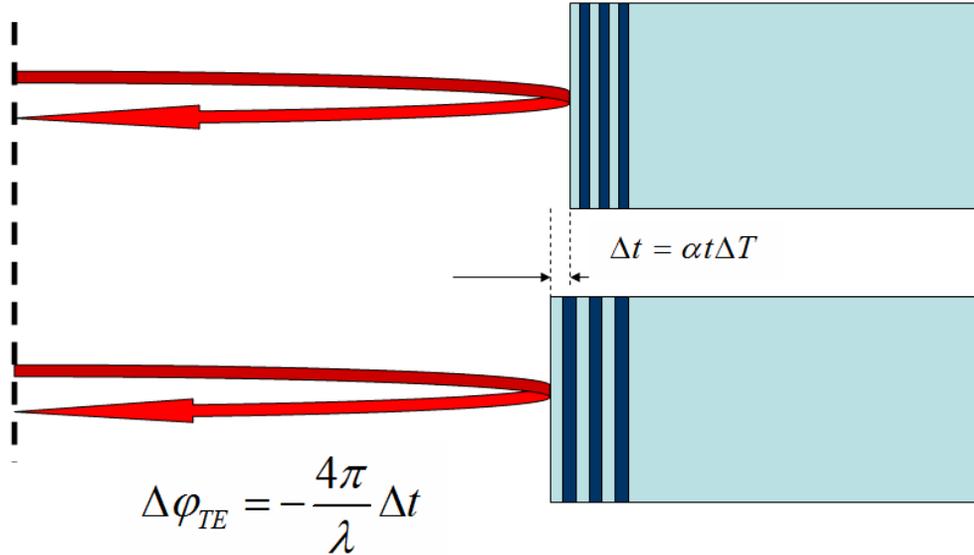


Figure 2.1: Illustration of thermo-elastic effect. We define this part of the effect to include only the motion of the mirror surface

mechanisms by which a change in temperature causes both an actual motion of the surface, as well as an observed change in the position of the surface.

2.2 Thermo-Elastic Expansion of the Coating

In the frequency regime we're working in, it is safe to assume that the coating is at a uniform temperature (in other words $\lambda_{thermal}$, the length scale over which thermal dynamics take place is much greater than the thickness of the coating). The first effect we need to calculate is the simple expansion of the coating's surface into the beam, as illustrated in Figure 2.1.

When the temperature increases by an amount ΔT , each of the layers will expand by an amount

$$\delta l = \alpha t_{layer} \Delta T$$

where α is the coefficient of thermal expansion of the layer (indeed, this is the definition of that parameter) and t_{layer} is the thickness of the layer. To find the total motion of the surface of the mirror, we add up the coefficients for the layers, and take the negative of the result to calculate the change in the path length of the laser:

$$\delta l = - \sum_{layers} \alpha_i t_i \Delta T. \quad (2.6)$$

This is the thermo-elastic (TE) part of the thermo-optic response.

2.3 Thermo-Refractive Coefficient of the Coating

Next, we need to calculate the rate of change in the optical phase of the reflected laser beam with respect to temperature, as illustrated in Figure 2.2.

2.3.1 Quarter-Wave Full Derivation

We will start out by doing the full calculation for a so-called quarter-wavelength (QWL) stack. This means our coating will be made of doublets of high- and low-index materials, in which each layer has an optical path length equal to one quarter of the wavelength of the light. In our case the mirrors are also capped with a full wavelength of low-index material. Following Fowles [13], we start with a plane wave of the form

$$\vec{E}(z, t) = E_0 \hat{x} e^{i(kz - \omega t)}$$

and on reflection from a multilayer coating the field will pick up a complex reflection coefficient given by

$$r = \frac{A + Bn_s - C - Dn_s}{A + Bn_s + C + Dn_s} \quad (2.7)$$

Thermo-optic Noise: TO = TE + TR

- **Thermo-Refractive (TR):** Coating layers deviate from $\lambda/4$ condition – due to both physical expansion and change in index of refraction. To first order, this manifests as a change in the phase of the reflected beam.

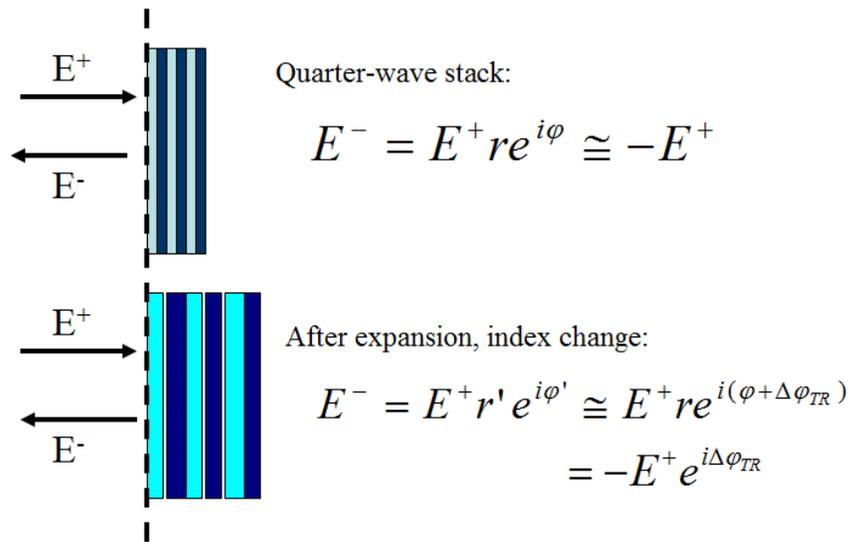


Figure 2.2: Illustration of thermo-refractive effect. We define this part of the effect to include only the change of phase that occurs promptly on reflection from the surface, and not surface motion

where n_s is the index of refraction of the substrate, and we have assumed the index of refraction of the incident medium (air/vacuum) is equal to 1. The A, B, C, D coefficients are entries of the coating matrix

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = M = M_1 M_2 M_3 \dots M_N. \quad (2.8)$$

Here the constituent matrices are for individual layers (layer 1 corresponding to the top layer facing the vacuum, and layer N corresponding to the layer adjacent to the substrate) and take the form

$$M_i = \begin{bmatrix} \cos(kl) & -i/n \sin(kl) \\ -in \sin(kl) & \cos(kl) \end{bmatrix} \quad (2.9)$$

with kl being the phase accumulated by the light in a one-way pass through a layer and n being the index of refraction of that layer.

For traditional quarter-wave stacks, kl for all layers is equal to $\pi/2$ (one quarter of a wave), and we have a simplified version of the M s

$$M_{i,\lambda/4} = \begin{bmatrix} 0 & -i/n \\ -in & 0 \end{bmatrix}$$

and combinations for a layer of high index followed by a layer of low index give

$$M_H M_L = \begin{bmatrix} -n_L/n_H & 0 \\ 0 & -n_H/n_L \end{bmatrix}$$

so that the overall transfer matrix for N doublets of this sort (for a total of $2N$ layers) is

$$M_{tot,\lambda/4} = \begin{bmatrix} (-n_L/n_H)^N & 0 \\ 0 & (-n_H/n_L)^N \end{bmatrix}$$

and the reflection coefficient is

$$r = \frac{(-n_L/n_H)^N - (-n_H/n_L)^N n_s}{(-n_L/n_H)^N + (-n_H/n_L)^N n_s} = \frac{1 - (n_H/n_L)^{2N} n_s}{1 + (n_H/n_L)^{2N} n_s} \quad (2.10)$$

Many coatings have an additional half-wave low index layer on the side facing the vacuum, but this will not change the reflection coefficient in the ideal layer case.

What we will calculate now is the predicted change in the reflection coefficient as we add a small phase to each of the layers. I'll add in a change of

$$kl = \pi/2 \longrightarrow (kl)' = \pi/2 + \delta$$

This means our layer matrices will change, but as long as we operate in the regime where $\delta \ll 1$ we can linearize the result and drop terms higher than first order in δ . In this case, our layer matrices become

$$M_{i,\delta} = \begin{bmatrix} -\delta & -i/n \\ -in & -\delta \end{bmatrix}$$

And, again to first order in the small parameter, we have doublet matrices of the form

$$M_H M_L = \begin{bmatrix} -\delta_H & -i/n_H \\ -in_H & -\delta_H \end{bmatrix} \begin{bmatrix} -\delta_L & -i/n_L \\ -in_L & -\delta_L \end{bmatrix} = \begin{bmatrix} -n_L/n_H & i(\delta_H/n_L + \delta_L/n_H) \\ i(\delta_L n_H + \delta_H n_L) & -n_H/n_L \end{bmatrix}.$$

Note that as we take powers of this matrix, the on-diagonal elements will remain the same as they were in the original matrix without this δ shift since, as we perform the matrix multiplications, any contribution to diagonal elements from off-diagonal elements will be second order (or higher) in the small parameters.

For the crucial off-diagonal parameters, we will develop a recursive formula. First, to help keep track of the algebra, we'll rewrite the doublet matrix in the form

$$M_H M_L = \begin{bmatrix} A & ix_1 \\ ix_2 & 1/A \end{bmatrix} \quad \text{with} \quad A = -n_L/n_H,$$

$$x_1 = \delta_H/n_L + \delta_L/n_H \quad \text{and} \quad x_2 = \delta_L n_H + \delta_H n_L.$$

Also, given the above argument about the diagonal elements containing no powers of the small parameters (the x s in this case), we know that the n^{th} power of this doublet matrix has to take the form

$$(M_H M_L)^n = \begin{bmatrix} A^n & i x_1 f_{1,n}(A) \\ i x_2 f_{2,n}(A) & 1/A^n \end{bmatrix}$$

And, in fact, if you'll excuse the mixed notation (raised numbers on A 's are powers of the number, raised numbers on M coefficients are the power of the matrix in which the coefficient resides), we can come up with a recursion relation for the off-diagonal elements of the n^{th} power of this doublet matrix:

$$(M_H M_L)^n = (M_H M_L)^{n-1} M_H M_L = \begin{bmatrix} A^{n-1} & M_{12}^{n-1} \\ M_{21}^{n-1} & 1/A^{n-1} \end{bmatrix} \begin{bmatrix} A & i x_1 \\ i x_2 & 1/A \end{bmatrix}$$

$$\begin{aligned} M_{12}^n &= A^{n-1} i x_1 + (1/A) M_{12}^{n-1} \\ M_{21}^n &= (1/A^{n-1}) i x_2 + A M_{21}^{n-1}. \end{aligned}$$

It can be shown that these recursion equations are both satisfied by the same solution (with the only exception being x_1 for M_{12} elements and x_2 for M_{21} elements):

$$M_{12}^n = M_{21}^n = i x_{1,2} \frac{1}{A^{n-1}} \sum_{m=0}^{n-1} A^{2m}.$$

Since A is less than 1 in magnitude, we can, for large values of n , take the limit of the sum as the upper bound becomes infinite, and write for “many doublets”

$$M_{12}^n = M_{21}^n \longrightarrow i x_{1,2} \frac{1}{A^{n-1}} \frac{1}{1 - A^2}$$

or, replacing the indices of refraction and negative sign that were hidden inside the A s,

$$M_{12}^n = M_{21}^n \longrightarrow ix_{1,2} \frac{n_H n_L}{n_H^2 - n_L^2} (-n_H/n_L)^n.$$

Finally, to complete the calculation we add one more layer on the top (multiplied from the left) that is a half-wavelength of low index. Since this layer is twice as thick as the quarter-wavelength layers and we are operating around $kl = \pi$, its form will be

$$M_{cap} = \begin{bmatrix} -1 & i/n_L 2\delta_L \\ in_L 2\delta_L & -1 \end{bmatrix}.$$

Plugging these into equation 2.7, we arrive at the reflection coefficient for our doublet:

$$r = \frac{n_S - (A)^{2n} + i[2\delta_L(1 + n_L A^{2n}) - x_1 f_1 n_L A^n + x_2 f_2 A^n]}{-n_S - (A)^{2n} + i[2\delta_L(1 - n_L A^{2n}) - x_1 f_1 n_L A^n - x_2 f_2 A^n]}.$$

We may safely drop terms of second order in A^n , since $(1.43/2.065)^{11} = .018$, so the square will be 4 orders below unity, in which case we have

$$r = \frac{n_S + i[2\delta_L - x_1 f_1 n_L A^n + x_2 f_2 A^n]}{-n_S + i[2\delta_L - x_1 f_1 n_L A^n - x_2 f_2 A^n]}. \quad (2.11)$$

The magnitude of this number will be effectively determined by the n_s terms in the numerator and denominator, since they will be of order unity, and as expected this reduces to the value of equation 2.10 as $\delta \rightarrow 0$. To find the phase of this reflection coefficient, we take

$$\arg(r) = \tan^{-1} \left(\frac{\Im(top)}{\Re(top)} \right) - \tan^{-1} \left(\frac{\Im(bottom)}{\Re(bottom)} \right) \cong \frac{\Im(top)}{\Re(top)} - \frac{\Im(bottom)}{\Re(bottom)}$$

which after some algebra simplifies to

$$\arg(r) = 4 \frac{\delta_L}{n_L} + \frac{2n_H n_L}{n_H^2 - n_L^2} \left(\frac{\delta_H}{n_L} + \frac{\delta_L}{n_H} \right) \equiv \phi_c. \quad (2.12)$$

Now to attend to the details of the δ s in terms of temperature and the coefficients that we're interested in measuring. By definition, we have that δ is the deviation from an optical path length of $\pi/2$, or

$$\delta = (kl)' - \pi/2 = kl + l\delta k + k\delta l - \pi/2 = l\delta k + k\delta l$$

where we have ignored second-order terms in the small variations. Plugging in the definition of the wavevector k and the thermal expansion l , we get

$$\delta = \frac{2\pi}{\lambda_0} l(\beta + n\alpha)\Delta T = \frac{\pi}{2} \left(\frac{\beta}{n} + \alpha \right) \quad (2.13)$$

with β the change of index of refraction with temperature (dn/dT) for the layer, α the coefficient of thermal expansion ($l^{-1}dl/dT$) for the layer, and n the index of refraction of the layer.

As long as we're here, let's check the condition that the change in optical path length be much less than one. This would imply changes in temperature

$$\Delta T \ll \frac{2}{\pi(\beta/n + \alpha)} \approx 400^\circ C. \quad (2.14)$$

We expect temperature variations on the order of a few $^\circ C$, so the above analysis should hold. Finally, we can plug the formula 2.13 into equation 2.12 to obtain a formula for the effective β for our quarter-wavelength coating

$$\beta_{eff} = \frac{\partial \phi_c}{\partial T} = \frac{\pi}{n_H^2 - n_L^2} (\beta_H + \alpha_H n_H + \beta_L + \alpha_L n_L) + \frac{2\pi}{n_L} \left(\frac{\beta_L}{n_L} + \alpha_L \right). \quad (2.15)$$

It is worth noting that this entire calculation separates into the first term on the right side of 2.15, which is for the quarter-wave stack (and is the same with or without the half-wave cap), and the second term which is present only with (and due to) the half-wave cap.

When we designed our mirror coatings, we chose to include the half-wave cap because this enhanced thermo-refractive effect would provide for a larger signal. In

future studies, for instance those trying to determine whether layer thickness affects thermo-optical properties, the signal from this layer might actually make it more difficult to see the effect being studied.

2.3.2 1/8-3/8 Calculation

In our experiment, we will be using both the traditional quarter-wave coatings described above, and also coatings made of alternating 1/8th and 3/8th wavelength thicknesses. Such coatings, which have altered thickness ratios but still satisfy the condition that a doublet adds up to half of an optical wavelength, are known as Bragg reflectors and also provide for a maximum-reflectivity plateau at the wavelength of interest.

To calculate the effect of temperature change on these coatings, I performed essentially the same calculation as above, but let Mathematica keep track of the gory details. We start by defining layer matrices as in equation 2.9, the overall matrix of equation 2.8, and the reflection coefficient of equation 2.7, leaving the optical path lengths in the layers $\varphi_i \equiv k_i l_i$ as variables. Then I define my reflectance function

$$r = r(n_L, n_H, \varphi_L, \varphi_H, \varphi_{cap}, N).$$

I then have Mathematica carry out this calculation for a standard mirror

$$r_0 = r(1.43, 2.065, \pi/2, \pi/2, \pi, 11)$$

and also for a stack which has had its temperature raised by 1°C (which keeps us in the linear regime, see 2.14, and is a very convenient ΔT for calculating derivatives by difference approximation), leaving the α s and β s as variables in the full matrix calculation. The result will be very messy, but we will let Mathematica keep track of the details.

$$r' = r(n_i \rightarrow n_i + \beta_i, \varphi_i \rightarrow \varphi_i(1 + \alpha_i + \beta_i/n_i))$$

Now we can calculate the rate of change of the reflection coefficient with temperature by noting that

$$\frac{dr}{dT} \cong \frac{\Delta r}{\Delta T} = \frac{r' - r_0}{1^\circ C}.$$

Now that we have in Mathematica a full expression for the reflection coefficient with temperature, we can extract the pieces, i.e., the partial derivatives of r with respect to $\alpha_L, \alpha_H, \beta_L, \beta_H$ by evaluating the above expression with the other parameters set to zero. In other words, to first order in these parameters, we know that the complicated expression in Mathematica's memory must be reducible to

$$\frac{dr}{dT} = a\alpha_L + b\alpha_H + c\beta_L + d\beta_H \quad (2.16)$$

and we can extract a , for instance, by evaluating

$$a = \frac{1}{\alpha_L} \left. \frac{dr}{dT} \right|_{\alpha_H=\beta_L=\beta_H=0}$$

and discarding terms higher than first order in α_L . Doing similar evaluations for the other constants gives us a numerical formula for the change in reflection coefficient of our quarter-wavelength (QWL) coating:

$$\frac{dr}{dT_{QWL}} = -i(6.4\alpha_L + 2.9\alpha_H + 4.5\beta_L + 1.4\beta_H). \quad (2.17)$$

We note here that this formula matches exactly with the results of the previous section. Numerically this is clear from plugging numbers for indices of refraction into Equation 2.15, and the $-i$ coefficient out front is because we are expanding around a reflection coefficient that starts off at -1 ; at the -1 point in the complex plane, increasing the phase of the complex number without increasing the size of the number corresponds to motion in the $-i$ direction.

Now that we have verified this technique against the analytic results of the previous section, we move on to apply it to the other coating we are going to test. This coating differs from the previous one in the ratio of high-index to low-index optical path. For

this second round of coatings, we have $\varphi_L = 3\lambda/8$ and $\varphi_H = \lambda/8$. Otherwise the calculation remains exactly as described above, and evaluation leads to the following formula for the change in reflection coefficient for what I will refer to as the Bragg coating:

$$\frac{dr}{dT}_{Bragg} = e^{i243^\circ} (9.6\alpha_L + 2.18\alpha_H + 6.9\beta_L + 1.18\beta_H). \quad (2.18)$$

The strange angle on the change of the reflection coefficient reflects the fact that this coating without temperature modulation does not have a -1 reflection coefficient like the quarter wave stack. The baseline reflection coefficient for this 1/8-3/8 stack is

$$r_{0,Bragg} = 0.995e^{i153^\circ}. \quad (2.19)$$

Comparing 2.19 and 2.18, it is clear that the change in reflection coefficient with temperature is again at 90° from the original reflection coefficient, meaning the amplitude of the reflected field is unchanged, and an increased temperature corresponds to an increase in reflected phase.

2.4 Previously Published Parameters

Thin film material parameters are notorious for having different values than those of the same materials in bulk form. In this section we discuss the current state of knowledge of these parameters. Values from the literature and from Advanced LIGO (AdLIGO) software and documents are listed in Tables 2.1 through 2.4. We will comment on the measurements listed and why they prompt us to our own study of these parameters.

From Tables 2.1 and 2.2, we have Cetinorgu's values for the coefficients of thermal expansion for silica and tantala. This reference uses the curvature induced upon heating (due to different coefficients of expansion in the substrate and coating) to determine the coefficient of expansion of the coating. They also do a more extensive

Table 2.1: Previously published α silica measurements

Measured Value	Author	Notes
$2.1 \times 10^{-6} K^{-1}$	Cetinorgu et al. [11]	DIBS, 1.1- μ m-thick coating
$5.5 \times 10^{-7} K^{-1}$	Braginsky et al. [4]	LIGO ETM coatings, though they seem to have assumed bulk silica value

Table 2.2: Previously published α tantala measurements

Measured Value	Author	Notes
$4.4 \times 10^{-6} K^{-1}$	Cetinorgu et al. [11]	DIBS, 1.1- μ m-thick coating
$(5 \pm 1) \times 10^{-6} K^{-1}$	Braginsky et al [3]	LIGO ETM coatings, though they seem to have assumed bulk silica value
$-4.4 \times 10^{-6} K^{-1}$	Inci [17]	Fiber tip coated with 3.3 μ m tantala coating
$3.6 \times 10^{-6} K^{-1}$	Crooks et al. [8]	Unclear where this comes from — possibly private communication with REO

Table 2.3: Previously published β silica measurements

Measured Value	Author	Notes
$8 \times 10^{-6} K^{-1}$	Braginsky (GWINC v2)	Could not find specific reference

Table 2.4: Previously published β tantala measurements

Measured Value	Author	Notes
$1.21 \times 10^{-4} K^{-1}$	Inci [17]	Fiber tip coated with 3.3 μ m tantala coating
$6 \times 10^{-5} K^{-1}$	Gretarsson [15]	Silica tantala multilayer coating, needs assumed value of α for silica, tantala, β of silica

study of index of refraction versus ion beam energy, and Young’s modulus/Poisson ratio measurements using nanoindentation. The CTE is what we’re interested in though. These numbers are quite interesting, and we will look at comparisons between them and our results. There may or may not be significant differences though, due to their dual-ion-beam-sputtering deposition technology and the fact that they have a monolayer that is roughly ten times thicker than our individual layers.

A reference to Braginsky for the CTE of fused silica is listed in our internal programs for determining noise levels in LIGO interferometers (Bench and GWINC); however, I was unable to track down an actual publication of an experiment to measure it, and it is (suspiciously) the exact same value as that for bulk silica. Braginsky’s measurement for the CTE of tantala was measured using an actual LIGO end mirror coating on a silica cantilever, along with a bending technique similar to Cetinorgu’s. Using this technique, he came up with a value for the lumped CTE of the whole coating, from which he inferred the CTE of tantala — though I suspect he used that bulk fused silica number for the fused silica stack layers.

Inci performed a beautiful experiment that enabled him to determine both the CTE and dn/dT for a tantala layer in one measurement. He used a fiber “Michelson”-type interferometer where one fiber end was a “mirror” comprised of a thin layer of tantala. He illuminated the interferometer with a superluminescent diode that had wavelengths spanning multiple free spectral ranges of the etalon formed by the tantala layer. He monitored the output of the interferometer with an optical spectrum analyzer as he raised the temperature of the tantala fiber end, and by watching the interference as a function of wavelength and temperature was able to do a one-shot measurement of both material properties. While a truly elegant measurement, the resulting numbers are of questionable use in our application for two reasons. The first is that the tantala layer is, at 3.3 microns, quite thick (he needed 3.3 microns in order to get multiple free spectral ranges over the wavelength range of the diode), and if these properties depend on thickness his results could be in an entirely different regime from the 129 nm that constitutes $\lambda/4$ at a free space laser wavelength of 1064 nm. Second, a single coating deposited on the end of a glassy optical fiber is under

significantly different stresses and boundary conditions than a thin coating in a stack on top of a large glassy substrate.

Finally, for the last of the α measurements, Crooks et al. refer to a value of the CTE of tantala in their paper on Brownian noise, but we were not able to track down their reference — it was likely in a referenced private communication with the coating company, Research Electro-Optics (REO).

The change of index with temperature is even less well measured. For fused silica, the one value we found was quoted in GWINC (LIGO internal software), but again I was unable to track down a reference.

As mentioned before, the fiber interferometer measurement by Inci was able to give a value of dn/dT , but the previous objections to the applicability of his numbers apply here as well.

Finally, Andri Gretarsson made a measurement of dn/dT for tantala in a silica/tantala coating by using a laser on the edge of the mirror's reflectance plateau while heating the sample up in an oven. This seems like a quite applicable method for determining this parameter; our only caveat is that he needed to assume values for all the other coating parameters in order to extract a number for dn/dT of tantala; a troubling requirement given the state of knowledge above.

2.5 Our Measurement

We would like to measure these material parameters in a setting as close to that of their eventual use as possible. To that end, we have designed a pump-probe style experiment and a set of coating samples that will allow us to measure these parameters as they appear in a dielectric stack mirror coating.

A cartoon of the measurement idea is presented in Figure 2.3. We start with a mirror and a dielectric coating, and cyclically heat it with a chopped (or more precisely AOM-modulated) heating laser. The heating laser is a CO_2 laser with a wavelength of $10.6\mu\text{m}$ which is highly absorbed in the coating (in fact we put a layer of chromium under the coating to make sure that any absorption happens in the coating, not the

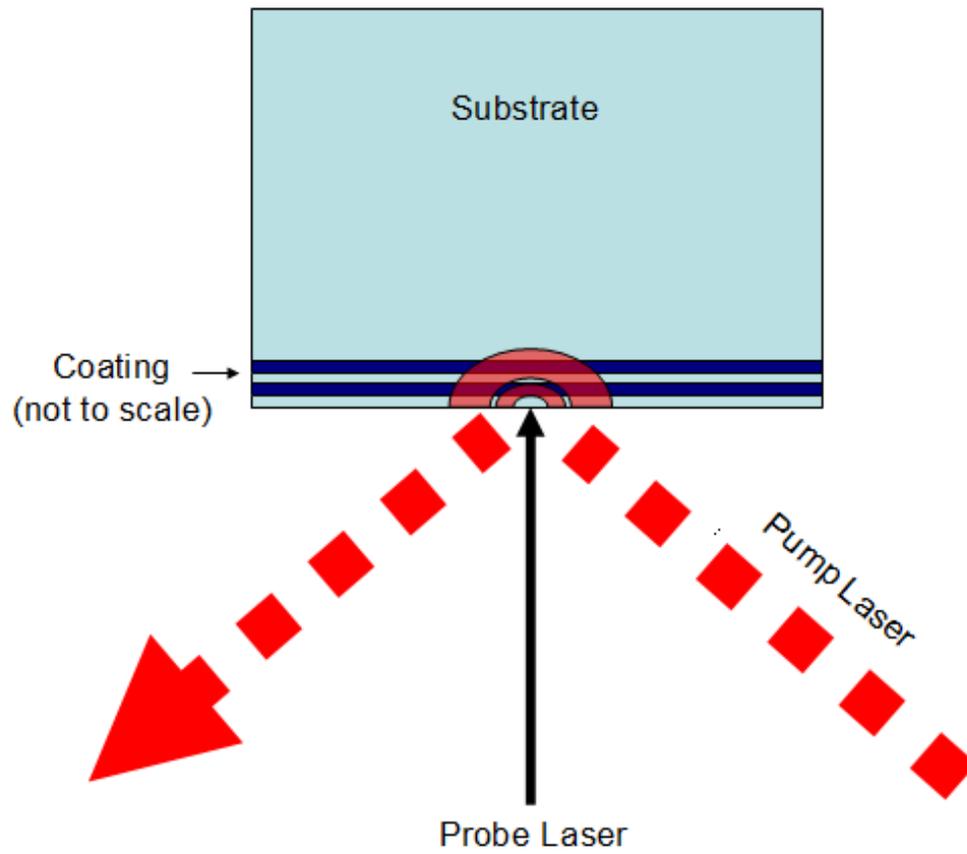


Figure 2.3: Measurement Idea — heat the coating with a chopped CO₂ laser, watch for changes in reflection coefficient and expansion of coating with 1064 nm probe laser

substrate). The oscillating temperature in the coating creates thermal waves that penetrate into the substrate a thermal distance scale that is approximately the same as the thermal wavelength, given by

$$\lambda_T = \sqrt{\frac{\kappa}{\rho C_P \pi f}} \quad (2.20)$$

where κ is the thermal conductivity, ρ the density, C_P the specific heat, and f the frequency of the heating beam modulation. As long as we operate in a regime where this thermal distance scale is large compared to the coating thickness (4 microns), all the parameters above are substrate parameters.

As we increase the modulation frequency of the heating beam, we decrease the length scale of the thermal variations, and we will switch from a regime where the effects measured by the probe beam are dominated by the expansion of the substrate to a regime where the effects measured by the probe beam are largely influenced by the changing temperature of the coating.

The response of the coating to these driven temperature changes is what we are measuring. It is determined by the coefficient of thermal expansion (α) of both coating materials, as well as the change of index of refraction with temperature (β) of the materials. In other words, when we take one measurement, we get a number we will refer to as β_{eff} for the whole coating, which will be a linear combination of α_L , α_H , β_L , and β_H .

To sort out the four parameters we are interested in, we make changes in the coatings and look for changes in the resulting β_{eff} . Our first parameter of variation is to use both a traditional “quarter wavelength” (QWL) coating in which each layer constitutes 1/4 of an optical wavelength, and also a so-called “Bragg” coating in which doublets of layers add up to 1/2 of an optical wavelength but the ratio of low-to high-index material is altered (in our case we will use 1/8 wavelength of high index tantala and 3/8 wavelength low-index silica).

These are the two coatings we use, but we still need two more independent measurements to sort out the four parameters of interest. For the final measurements,

we apply a thin layer of gold on the surface of both types of coatings above and re-take the measurements. The gold layer is thin enough that its expansion can be neglected, but thick enough to fully block both laser beams from penetrating even into the coating. Some small amount of CO_2 power will still be absorbed in the gold, and drive thermal oscillations, but the probe beam will measure *only* the expansion of the coating; when the probe does not penetrate into the coating at all, the indices of the layers will not play a role in the measured response.

These four measurements, the “combined” measurements where we see both expansion and index effects of the QWL and Bragg coatings, and the “thermo-elastic expansion (TE) alone” measurements of the QWL and Bragg coatings give us enough information to infer the four parameters of interest.

As a final detail, it turns out that the expansion of our chromium layer mentioned earlier is large enough that we need to measure that too. We add one more measurement, of the gold-coated type, with QWL coating where the layer of chromium is thin (too thin to block the CO_2 radiation, but that doesn’t matter under the gold layer), so that we can estimate its effects and subtract them out of the final measurements.

Chapter 3

The Experiment

3.1 Introduction

We designed our photothermal experiment to measure the photothermal coefficients (dn/dT) as well as the coefficients of expansion for layers in our dielectric stack mirror coating. This is a pump-probe style measurement using a low-noise 1064 nm NPRO as the probe laser and an amplitude-modulated 10.6 μm CO₂ pump laser to heat the device under test.

The CO₂ laser has benefits and drawbacks as pump. Its major benefits are that it is absorbed well in the coating (glasses have very high absorption coefficients at this wavelength), and CO₂ lasers are available with very high power for relatively low cost. Thus driving the heating of the coating under test is much easier than using a laser for which the coating materials have low absorption. This property also makes it almost impossible *not* to separate the pump and probe beams before the photodetectors, mitigating the problem of direct pump/probe crosstalk; any glass windows, lenses, or other optics in the probe path will completely block stray pump light.

The first drawback of this laser wavelength is the flip-side of its benefits. Because it is easily absorbed by most materials, any scattered light will be efficiently absorbed wherever it happens to land, and because the laser is relatively powerful it can create heating expansion large enough to drive mechanical resonances in, e.g., mounts of any mirrors it might land on. These driven mechanical resonances are particularly devious because they show up at the driving frequency of the heating beam, and thus,

while they are completely spurious they are impossible to separate from the optical signal we are trying to detect. The solution to this problem is to be very careful to avoid (to the extent possible) scattering of the CO₂ beam, e.g. by keeping the beam focused, avoiding clipping on optics, and shielding all the measurement optics from the CO₂ laser path. We used a wall of aluminum foil to sequester the optic under test and the heating beam from the rest of the interferometer. In the same vein, it was also important to dump the excess CO₂ power off of the interferometer table; such a dump by definition involves absorbing the entire remaining beam, which we found created enough motion in the dump to shake interferometrically sensitive optics and ring up resonances of their mounts and the table.

The second drawback to the 10.6 μ m wavelength is that although it is strongly absorbed, the absorption length (the depth at which the electric field is reduced to $1/e$ of its initial strength) is roughly 34 μ m, which is significantly longer than the depths we want the thermal waves to penetrate. Our solution to this problem, described below, is to include a layer of chromium underneath the dielectric stack coating. The chromium is highly reflective of the heating laser, preventing its electromagnetic radiation from penetrating far into the substrate, but being at the back of the substrate it does not (except through its expansion, also described below) significantly affect the probe beam reflection.

3.2 Theory

In this section, we will work out the expected signal from the interferometer as a function of the modulation frequency of the probe beam. To do this, we will need to work out the expected temperature profile of the mirrors in our experiment, both in a frequency regime where simplifications make the math particularly easy, and in a more complicated low frequency regime.

Table 3.1: Typical values of mirror geometry parameters

r_{mirror}	r_{pump}	r_{probe}	$\lambda_{thermal}$	$t_{coating}$
12.7 mm	2.0 mm	0.37 mm	10 μm	4 μm

3.2.1 Mirror Geometry

The essential geometry of our problem can be seen in Figure 2.3. The relevant details are the relative dimensions of a few important parameters:

$$r_{mirror} > r_{pump} > r_{probe} > \lambda_{thermal} > t_{coating}. \quad (3.1)$$

Table 3.1 shows values for these parameters in a typical run.

The mirror is being heated by a probe beam that is approximately half its diameter. This will potentially make edge effects important at the few percent level. We analyze these effects in COMSOL simulations described below.

The pump beam is 5-10 times larger than the probe beam, which simplifies the problem of figuring out the signal that the probe beam will read out. Rather than integrating the phase picked up by the probe beam at each transverse point, weighted by the Gaussian beam profile, we can just assume the phase is a constant across the probe beam, and equal to the value of that number at the center of the pump beam. The error in making this approximation is of order $\sim 1/5e$ in the tails of the probe Gaussian, so we estimate this also at less than a few percent.

The heating power is modulated, giving rise to oscillating changes in the temperature profile in the mirror. These effects generally do not travel forever, but penetrate only to a certain depth $\lambda_{thermal}$, beyond which temperature changes are negligibly small (and exponentially suppressed). For almost the entire range of our tests, the distance scale over which thermal effects are significant is much smaller than either the pump or probe beams, allowing us to treat temperature variations as plane waves propagating in from the surface.

3.2.2 1-D Approximation

In this section we will work out the response of the mirror under the simplest of assumptions; namely that 1) the mirror and heating beam are infinite in extent with a given power per unit area incident on the surface, and 2) the thermal wavelengths are much longer than the coating thickness so that it can be ignored in determining the temperature profile, and will respond only to the temperature at the very surface of the mirror.

3.2.2.1 Temperature Profile

In this regime, the only dimension along which the temperature could be expected to vary is the z -direction; into the substrate. With no variation in the transverse directions, the heat equation takes the form

$$\frac{\partial u}{\partial t} = a \frac{\partial^2 u}{\partial z^2} \quad (3.2)$$

where u is the temperature (or more precisely the small deviation about the mean temperature), z is distance into the substrate ($z = 0$ being the coating's location, $z < 0$ assumed to be thermally insulating vacuum), and the constant a being given by

$$a \equiv \frac{\kappa}{\rho C_p} \quad (3.3)$$

with κ being the thermal conductivity, ρ the density, and C_p the specific heat of the bulk substrate. Again, we are working in the limit where the coating is a surface layer thin enough that it does not contribute to the heat equation dynamics.

Since this problem statement has no heat sources in the substrate, we can just write out solutions to the homogeneous equation, and we will use the influx of power at the surface as our boundary condition. Solutions to the heat equation here take the form

$$u(t, z) = C e^{i(\omega t - kz)} \quad (3.4)$$

and the heat equation will give us a relationship between the temporal and spatial frequencies:

$$i\omega = -ak^2.$$

$$\implies k = e^{-i45^\circ} \sqrt{\omega/a} \quad (3.5)$$

We selected which root to use by demanding that the temperature variations go to zero as z goes to infinity (as opposed to an unphysical exponential rise).

Now to determine the constant C in equation 3.4 (which will be a function of frequency but constant in time and space), we will employ the power flux boundary condition at the surface:

$$-\kappa \frac{\partial u}{\partial z} \Big|_{z=0} = \frac{P_0}{A} e^{i\omega t} \quad (3.6)$$

where κ is the thermal conductivity of the substrate, P_0 is the incident power in the pump beam, and A is the area over which the pump beam is distributed.

Taking the appropriate derivative and solving for C gives us

$$C = \frac{P_0}{ik\kappa A} = \frac{P_0 \sqrt{a} e^{-i45^\circ}}{\kappa A \sqrt{\omega}}.$$

We now have our temperature profile throughout the substrate:

$$u(t, z) = \frac{P_0 \sqrt{a} e^{-i45^\circ}}{\kappa A \sqrt{\omega}} e^{i(\omega t - kz)}. \quad (3.7)$$

3.2.2.2 Thermo-Refractive Response

With this temperature profile in hand, we can go ahead and calculate the response of the mirror as seen by the probe beam. The thermo-refractive effects are simply

the temperature at the surface of this temperature profile multiplied by the effective β for the coating, calculated in the previous section.

$$\delta l_{TR} = \frac{\lambda}{4\pi} \beta_{eff} \frac{P_0 \sqrt{a}}{\kappa A} \frac{1}{\sqrt{\omega}} e^{i(\omega t - 45^\circ)} \quad (3.8)$$

The prefactor of $\lambda/(4\pi)$ is to convert the phase change calculated in the previous section into an equivalent length so as to compare it to the other length changes that are observed.

3.2.2.3 Substrate Thermo-Elastic Response

In addition to the thermo-refractive effect, we will also observe the expansion of the substrate. To calculate this effect, we integrate the temperature profile along z , multiplied by the coefficient of thermal expansion of the substrate (α).

$$\begin{aligned} \delta l_{TE} &= -\alpha \int_0^\infty u(t, z) dz = -\alpha \frac{P_0 \sqrt{a}}{\kappa A} \frac{e^{-i45^\circ}}{\sqrt{\omega}} e^{i\omega t} \int_0^\infty e^{-ikz} dz \\ &= \alpha \frac{P_0 a}{\kappa A} \frac{1}{\omega} e^{i\omega t} e^{i90^\circ} \end{aligned} \quad (3.9)$$

The negative sign in this first equation reflects the fact that in the readout of our interferometer an increase in the size of the substrate will be read as a decrease in the corresponding arm length. Also note that the negative sign is responsible for changing the expected 90 degree lag into a 90 degree lead on the driving signal (or if you prefer, a 270 degree lag, but the lock-in amplifier we used returns phases in the range $\{-180^\circ, 180^\circ\}$).

3.2.3 Effects of Penetrating Radiation in 1-D Approximation

During the course of setting up and debugging the experiment, I discovered that my first run of mirrors had a chromium layer that was far too thin to prevent radiation from penetrating into the substrate. A second set of mirrors was ordered to ensure this wouldn't be a problem, but it was of interest to calculate how any penetrating

CO₂ radiation would affect the results of the experiment.

3.2.3.1 Solving the Heat Equation

Once again we will solve the heat equation in the regime where the transverse spot size is much greater than the thermal length scale and the probe beam length scale, but this time we will apply two uniform volumetric heat injections - one in the coating and another in the first $\sim 30\mu\text{m}$ of the substrate. This is realistic for the coating, and although we would expect the substrate power to be injected with a profile that exponentially decays with distance from the surface, this should be a decent approximation to observe the effects we are concerned with. The plan will be to inject a total power equal to the amount we observe being absorbed, and then change the fraction of that power that is absorbed in the coating vs. substrate to see if we should expect an observable difference in the response of the interferometer.

The heat equation with a forcing function becomes

$$\frac{\partial u}{\partial t} = a \frac{\partial^2 u}{\partial z^2} + F \quad (3.10)$$

where our forcing function will be defined to be

$$F = \begin{cases} e^{i\omega t} \text{frac} P_0 / (\rho_C C_C A z_1) \equiv f_1 e^{i\omega t} & \text{for } 0 \leq z < z_1 \\ e^{i\omega t} (1 - \text{frac}) P_0 / (\rho_S C_S A l) \equiv f_2 e^{i\omega t} & \text{for } z_1 \leq z < z_2 \\ 0 & \text{for } z_2 < z \end{cases} \quad (3.11)$$

Region I from $z = 0$ to $z = z_1$ is the coating, with $z_1 = 4 \mu\text{m}$. Regions II and III are both substrate, the distinction being that Region II will have a uniform power injection (this will be the “absorption layer”) and Region III will not. I am going to take the total power I observe being absorbed by my mirror, P_0 , and inject some fraction frac uniformly in the coating (Region I) and inject the remaining power uniformly into Region II which has a length l (so that $x_2 - x_1 = l$). We will look for solutions that vary sinusoidally with time at the same frequency as the driving force:

$$u(t, z) = u(x)e^{i\omega t}.$$

This gives the heat equation the inhomogeneous form

$$i\omega u = a \frac{\partial^2 u}{\partial z^2} + f_i. \quad (3.12)$$

The homogeneous solutions to this equation are the same as in the previous section, equations 3.4 and 3.5, and are reprinted here

$$u(t, z) = Ce^{i(\omega t - kz)} \quad \text{with} \quad k = e^{-i45^\circ} \sqrt{\omega/a}.$$

Next we need a particular solution to the equation with the forcing function in each of the regions where the forcing function is nonzero. Fortunately, this particular solution is easy to deduce and takes on the simple form

$$u_p(z, t) = \frac{-if_i}{\omega} e^{i\omega t}. \quad (3.13)$$

So our solutions in the free region will be exponentially decaying thermal waves, and in the other two regions they will be linear combinations of exponentially decaying and growing thermal waves, combined with uniform oscillation at the driving frequency. Or in other words, my temperature profile will be

$$u(z) = \begin{cases} A'e^{ik_1z} + A''e^{-ik_1z} - \frac{if_1}{\omega} & \text{Region I} \\ Be^{ik_2z} + Ce^{-ik_2z} - \frac{if_2}{\omega} & \text{Region II} \\ De^{-ik_2z} & \text{Region III} \end{cases}.$$

My first boundary condition will be an easy one — I'll assume no heat flows into or out of the mirror surface. Previous calculations have had the driving force flow in through the surface, but in this case we're modeling the driving thermal force as being sourced in the bulk. Aside from that, the only heat that would enter or leave the surface is radiative heat transfer to other parts of the inside of the vacuum chamber. Since the temperature differences in question are small (a few degrees at most), this

will be negligible.

The effect of assuming no heat flow at $z = 0$ is that the temperature profile must be flat there, which in turn will imply $A' = A''$. We can therefore rewrite that first temperature profile as

$$u_1(z) = A'(e^{ik_1z} + e^{-ik_1z}) - \frac{if_1}{\omega} = A \cos(k_1z) - \frac{if_1}{\omega}.$$

We now have four coefficients to be determined (A, B, C, D), and to solve for them I will impose boundary conditions of 1) temperature continuity and 2) power flux continuity at each of the two interfaces between layers.

For the first boundary, we have from temperature continuity

$$A \cos(k_1z_1) - \frac{if_1}{\omega} = B e^{ik_2z_1} + C e^{-ik_2z_1} - \frac{if_2}{\omega} \quad (3.14)$$

and from power flux continuity

$$-\kappa_1 k_1 A \sin(k_1z_1) = \kappa_2 (ik_2 B e^{-k_2z_1} - ik_2 C e^{-ik_2z_1}) \quad (3.15)$$

For the second boundary, we have from temperature continuity

$$B e^{ik_2z_2} + C e^{-ik_2z_2} - \frac{if_2}{\omega} = D e^{-ik_2z_2} \quad (3.16)$$

and from power flux continuity

$$\kappa_2 (ik_2 B e^{-k_2z_2} - ik_2 C e^{-ik_2z_2}) = -\kappa_2 ik_2 D e^{-ik_2z_2} \quad (3.17)$$

These equations can be solved to determine the four unknowns. The algebra becomes cumbersome, so I will spare the reader some of the details and in addition write C in terms of A , and D in terms of B and C . Thus they will all be determined with minimal clutter, and from here on we will use Mathematica to keep track of algebraic details.

$$A = \frac{i}{\omega} \frac{f_1 - f_2 + f_2 e^{-ik_2 l}}{\cos(k_1 z_1) + ir \sin(k_1 z_1)} \quad (3.18)$$

$$B = \frac{if_2}{2\omega} e^{-ik_2 z_2} \quad (3.19)$$

$$C = e^{ik_2 z_1} \left[-irA \sin(k_1 z_1) + \frac{if_2}{2\omega} e^{-ik_2 l} \right] \quad (3.20)$$

$$D = e^{ik_2 z_2} [-B e^{ik_2 z_2} + C e^{-ik_2 z_2}] \quad (3.21)$$

and we have defined the ratio

$$r = \frac{\kappa_1 k_1}{\kappa_2 k_2} = \frac{\kappa_1}{\kappa_2} \sqrt{\frac{a_2}{a_1}}. \quad (3.22)$$

We have now written down the full solution for $u(z, t)$, or alternatively, in phasor notation, $u(z, \omega)$. To determine the mirror response in a Michelson interferometer, we need to add up the TE and TR pieces. To find the TE response, we integrate the temperature profile over z multiplied by α in the different regions.

$$\delta l_{TE} = -\alpha_c \int_0^{z_1} u_1(z, \omega) dz - \alpha_S \int_{z_1}^{z_2} u_2(z, \omega) dz - \alpha_S \int_{z_2}^{\infty} u_3(z, \omega) dz. \quad (3.23)$$

In practice the third integral will only be taken to the back of the substrate, but with the exponential decay of waves the distinction is one we don't have to worry about.

The other piece of the mirror response is the thermo-refractive response of the coating. As in previous sections, we will assume the coating is approximately at the same temperature throughout, which is a good approximation as long as the thermal wavelength is longer than the coating thickness. Thus, we will use the temperature at $z = 0$ as the temperature of the entire coating, and simply multiply by the effective β calculated in Section 2.3.

$$\delta l_{TR} = \frac{\lambda}{4\pi} \beta_{eff} u(0, \omega) \quad (3.24)$$

where the prefactor has been added to convert the phase change with temperature into an equivalent observed displacement (a displacement of one wavelength in a Michelson arm will change the phase of the returning light by 4π).

3.2.3.2 Results

Plotting the above solutions in Mathematica, we find two major results worth noting. The first result is for a sapphire substrate, where the entire response is dominated by the TE expansion.

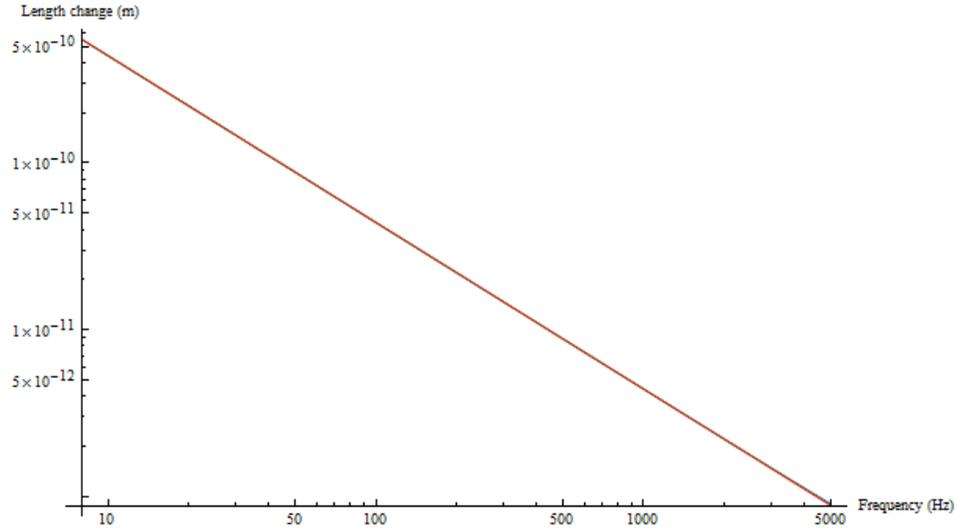


Figure 3.1: Sapphire response magnitude — location of absorption makes no difference in this frequency range

We show the magnitude response in Figure 3.1. For sapphire in the frequency range of interest (10 Hz – 5000 Hz), there is no appreciable difference in response magnitude whether the power is absorbed in the coating or in the substrate absorption layer. This value agrees extremely well with the calculation done in Sections 3.2.2 and 3.2.4. Looking at the phase in Figure 3.2, we see that there is a slight difference in the phase at the high end of the frequency range, where the thermal wavelengths are just starting to get small enough to make a difference in the balance between TE

and TR effects.

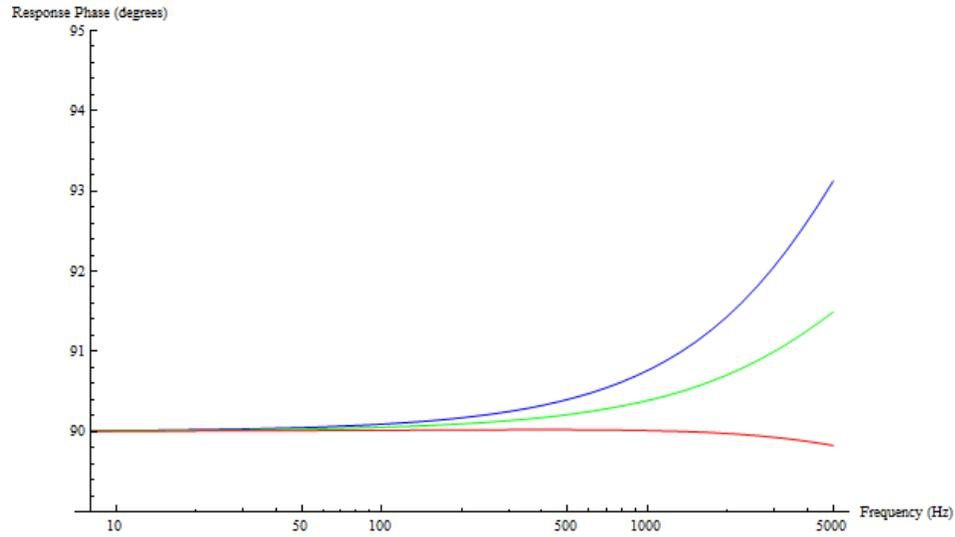


Figure 3.2: Sapphire response phase for fraction absorbed in coating = 100% (blue), 50% (green), 0% (red)

The next application of these formulae that is instructive to look at is the effects of penetrating radiation in the case of a Zerodur substrate, where the entire response of the mirror is dominated by the coating response. The simple theory discussed in Section 3.2.2 suggests that the surface temperature (and thus the coating response magnitude) will fall as $f^{-1/2}$. The traditional thermal expansion that is proportional to f^{-1} is due to the integration of the thermal profile (which brings down an additional factor of $f^{-1/2}$). In Zerodur, while the thermal wavelengths are similar to those of fused silica or Infrasil, the coefficient of expansion is so small that the effects of the integrated temperature profile in the substrate do not contribute significantly.

The interesting feature to note in Figure 3.3 is that the more radiation that is absorbed in the substrate, the more the magnitude response deviates from the expected $f^{-1/2}$ frequency dependence (the observed deviation even with 100% absorbed in the coating is due to some residual mixing with a non-zero substrate expansion as well as the beginnings of the onset of higher frequency dynamics which will tend to shift the TR response to f^{-1} as well).

The trend observed in the magnitude of our Zerodur response, namely that absorption in the substrate layer leads to significant deviation from predicted frequency

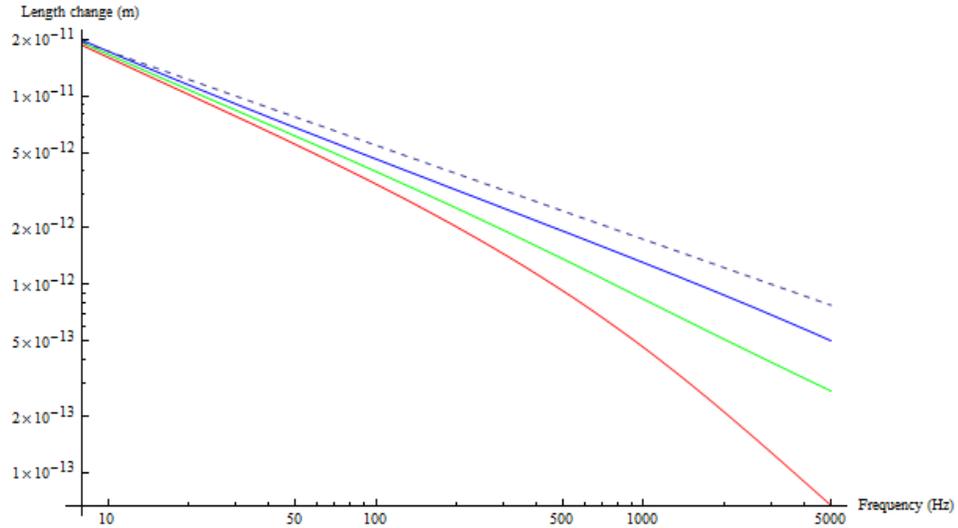


Figure 3.3: Zerodur response magnitude for fraction absorbed in coating = 100% (blue), 50% (green), 0% (red). Dashed line is arbitrary $f^{-1/2}$ for visual aid.

dependence, is even more pronounced in the phase response. The simple prediction with no substrate expansion (just TE/TR coating response) would be a flat 135° phase response. As can be seen in Figure 3.4, it is indeed the case that absorption in the coating is closest to this (with a move towards 90° at low frequency as non-zero substrate expansion starts to become important), while significant deviation is in evidence with more substrate absorption.

3.2.3.3 Conclusions

We take two conclusions away from this calculation with radiation penetrating into the substrate. First of all, the Sapphire sample should show almost no deviation with the effects of penetrating radiation, which means it will still be a useful calibration tool even if this effect is taking place. Second, and more importantly, we can use the response of samples on a Zerodur substrate as a check to see that our chromium layer is doing its job; to the extent that the Zerodur samples show the expected frequency response (magnitude and phase), we can be assured that the heating power is indeed being absorbed solely in the coating.

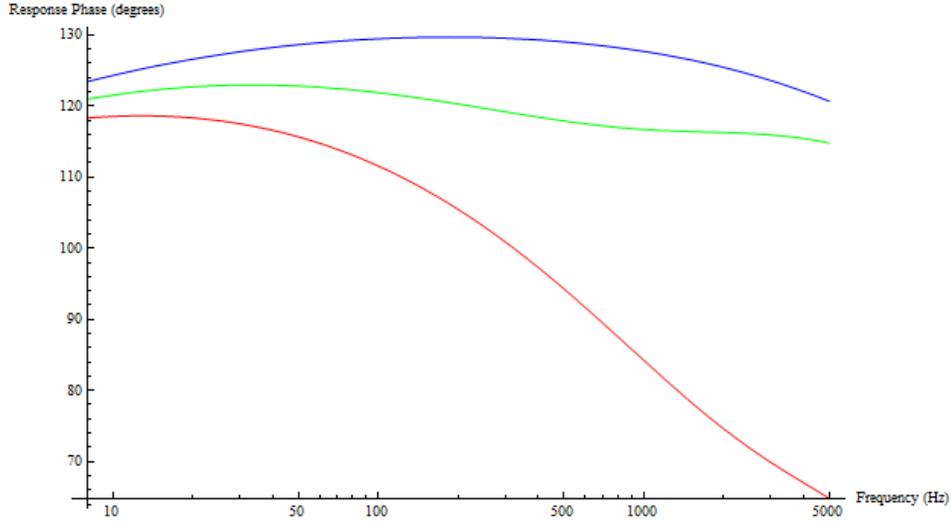


Figure 3.4: Zerodur response phase for fraction absorbed in coating = 100% (blue), 50% (green), 0% (red)

3.2.4 Low Frequency Green's Function Calculation

As the last important bit of theory, I will calculate the expected low frequency response of the substrate, as the thermal wavelength approaches the spot size and the 1-D approximation begins to break down. With the exception of Sapphire, most of our experiments do not reach the regime where the thermal wavelength is on the order of the heating spot size, but there are observable corrections that we can make at the low end of our frequency response even so.

As usual, we start with the heat equation with a forcing function

$$\frac{\partial u}{\partial t} - a\nabla^2 u = F \quad (3.25)$$

And again, we will be working in the regime where thermal wavelengths are much larger than the coating thickness, but we will relax the condition that the thermal wavelengths are larger than the heating spot size. This increases the dimensionality of the problem to 3-D (well technically 2-D since we still have axial symmetry, but we will solve it in 3-D).

We will take the heat injected by the CO₂ laser to be a Gaussian intensity profile centered at the center of the mirror and probe beam, at $x = y = z = 0$, and to be

nonzero (and constant as a function of z) only in the first few microns of the substrate (the coating). Our sinusoidal heat injection (normalized to a total absorbed power P_0) is

$$F(x, y, z, t) = \frac{P_0}{\rho C_P \pi r_0^2 z_0} e^{-\frac{x^2+y^2}{r_0^2}} e^{i\omega t} \quad \text{for } 0 < z < z_0. \quad (3.26)$$

The heat equation (3.25) has the Green's function solution for point sources in an infinite homogeneous space given by

$$G_\infty(t, \vec{x}, \tau, \vec{x}') = \left[\frac{1}{4\pi a(t-\tau)} \right]^{3/2} e^{-\frac{|\vec{x}-\vec{x}'|^2}{4a(t-\tau)}}. \quad (3.27)$$

We will assume for this calculation that the Green's function for point sources on the surface of an infinite half-space will be twice this, since point sources on the surface see essentially the same symmetry with the only exception that the energy that would have flowed into the negative half-space now also flows into the positive half-space. Technically for heat injection that is not directly on the surface, you would need a different Green's function (using method of images or some other technique, see, e.g., *Elementary Applied Partial Differential Equations* by Haberman), but since our coating is thin compared to the thermal length scale (remember we are carrying out this equation to learn about behavior at low frequencies), this approximation should suffice.

Given our Green's function

$$G(t, \vec{x}, \tau, \vec{x}') = 2G_\infty(t, \vec{x}, \tau, \vec{x}') \quad (3.28)$$

we can immediately write down the general solution to the temperature profile as

$$\begin{aligned} u(t, x, y, z) &= \int_{-\infty}^t d\tau \iiint d^3x' G(t, \vec{x}, \tau, \vec{x}') F(\vec{x}', \tau) \\ &= \frac{2P_0}{\rho C_P \pi r_0^2 z_0} \int_{-\infty}^t d\tau \iiint dz' dy' dx' \left[\frac{1}{4\pi a(t-\tau)} \right]^{3/2} e^{-\frac{|\vec{x}-\vec{x}'|^2}{4a(t-\tau)}} e^{-\frac{x'^2+y'^2}{r_0^2}} e^{i\omega\tau}. \end{aligned}$$

The z' -integral (from 0 to z_0) is simple to carry out under the assumption that z_0 is small, and gives us simply a factor of z_0 and replaces z' with 0. Again this approximation is good as long as the coating is thin compared to the thermal wavelength, which leaves us with

$$u(t, x, y, z) = \frac{2P_0}{\rho C_P \pi r_0^2} \int_{-\infty}^t d\tau \left[\frac{1}{4\pi a(t-\tau)} \right]^{3/2} e^{i\omega\tau} e^{\frac{-z^2}{4\alpha(t-\tau)}} \iint_{-\infty}^{\infty} dy' dx' e^{\frac{-[(x-x')^2+(y-y')^2]}{4a(t-\tau)}} e^{-\frac{x'^2+y'^2}{r_0^2}}.$$

Evaluating the overlap integral in x and y gives

$$\begin{aligned} u(t, x, y, z) &= \frac{2P_0}{\rho C_P \pi r_0^2} \int_{-\infty}^t d\tau \left[\frac{1}{4\pi a(t-\tau)} \right]^{3/2} e^{i\omega\tau} e^{\frac{-z^2}{4\alpha(t-\tau)}} \pi \frac{r_0^2 4\alpha(t-\tau)}{r_0^2 + 4\alpha(t-\tau)} e^{\frac{-(x^2+y^2)}{r_0^2 + 4\alpha(t-\tau)}} \\ &= \frac{2P_0}{\rho C_P \pi \sqrt{4\pi a}} \int_{-\infty}^t d\tau \frac{1}{\sqrt{t-\tau}(r_0^2 + 4a(t-\tau))} e^{i\omega\tau} e^{\frac{-z^2}{4\alpha(t-\tau)}} e^{\frac{-(x^2+y^2)}{r_0^2 + 4\alpha(t-\tau)}}. \end{aligned}$$

Now to find the TR coefficient, let's look at the temperature at the center of the heating spot, i.e., at $x = y = z = 0$:

$$u(t, 0, 0, 0) = \frac{2P_0}{\rho C_P \pi \sqrt{4\pi a}} \int_{-\infty}^t d\tau \frac{e^{i\omega\tau}}{\sqrt{t-\tau}(r_0^2 + 4a(t-\tau))}.$$

With a suitable change of variables, this integral can be re-written in a form that Mathematica can carry out (or at least rewrite in terms of well-studied integral forms):

$$\begin{aligned} u(t, 0, 0, 0) &= \frac{4P_0 e^{i\omega t}}{\rho C_P \pi \sqrt{4\pi a}} \int_0^{\infty} dq \frac{e^{-i\omega q^2}}{r_0^2 + 4a q^2} \\ &= \frac{P_0 e^{i\omega t} e^{i\frac{r_0^2 \omega}{4a}}}{2\rho C_P a r_0 \sqrt{\pi}} \left(-1 + (1+i) \text{FresnelC} \sqrt{\frac{r_0^2 \omega}{2\pi a}} + (1-i) \text{FresnelS} \sqrt{\frac{r_0^2 \omega}{2\pi a}} \right). \end{aligned}$$

While slightly unwieldy, this function behaves as one would expect – for large frequencies it levels off at -45° of phase and the amplitude asymptotically approaches $f^{-1/2}$. At low frequencies the phase is proportional to frequency (going to zero in the

limit of a DC signal), and the amplitude levels off to a constant.

Now let's push ahead and calculate the TE expansion at the center of the spot. We'll start with

$$u(t, 0, 0, z) = \frac{2P_0}{\rho C_P \pi \sqrt{4\pi a}} \int_{-\infty}^t d\tau \frac{1}{\sqrt{t - \tau} (r_0^2 + 4a(t - \tau))} e^{i\omega\tau} e^{\frac{-z^2}{4\alpha(t - \tau)}}$$

and calculate the expansion of the mirror surface in the standard way

$$\begin{aligned} \delta l_{TE} &= -\alpha \int_0^\infty u(t, 0, 0, z) dz \\ &= \frac{\alpha 2P_0}{\rho C_P \pi \sqrt{4\pi a}} \int_0^\infty dz \int_{-\infty}^t d\tau \frac{1}{\sqrt{t - \tau} (r_0^2 + 4a(t - \tau))} e^{i\omega\tau} e^{\frac{-z^2}{4\alpha(t - \tau)}} \\ &= \frac{\alpha P_0}{\rho C_P \pi} \int_{-\infty}^t d\tau \frac{e^{i\omega\tau}}{r_0^2 + 4a(t - \tau)} \end{aligned}$$

which again, switching variables of integration, gives us an integral to be evaluated of

$$\begin{aligned} \delta l_{TE} &= \frac{\alpha P_0}{\rho C_P \pi} e^{i\omega t} \int_0^\infty \frac{e^{-i\omega\tau}}{r_0^2 + 4a\tau} \\ &= \frac{\alpha P_0}{\rho C_P \pi} e^{i\omega t} \left[-\frac{i e^{ir_0^2\omega/4a}}{4a} (\pi - i \cdot \text{ExpIntegralEi}[-ir_0^2\omega/4a]) \right] \end{aligned}$$

Again, this function behaves as we'd expect, and in the limit of large frequencies asymptotically approaches

$$\delta l_{TE} \longrightarrow -i \frac{\alpha P_0}{\rho C_P \pi r_0^2 \omega} e^{i\omega t}$$

which is the same result as equation 3.9 if you note that $a/\kappa = 1/\rho C_P$ and $A = \pi r_0^2$ for the purposes of calculating power intensity at the center of the Gaussian beam.

3.2.5 Combined Signal

For our main substrates (fused silica and a similar glass called Infrasil), we will not need to use the low-frequency complex integrals, and the previous formulae will suffice. The full response of one of our mirrors will then be

$$\delta l(\omega) = \frac{\alpha P_0}{\rho C_P \pi r_0^2 \omega} e^{i90^\circ} + \frac{\lambda}{4\pi} \beta_{eff} \frac{P_0 \sqrt{a}}{\kappa A} \frac{1}{\sqrt{\omega}} e^{-i45^\circ}. \quad (3.29)$$

We emphasize here again that since we are working at frequencies low enough to put us in the long wavelength limit, the relevant material parameters listed here are all substrate parameters except for those hidden inside β_{eff} . The first term will dominate at low frequencies and will only contain substrate information. The second term will eventually take over at a frequency where

$$\frac{\alpha P_0}{\rho C_P \pi r_0^2 \omega} = \frac{\lambda}{4\pi} \beta_{eff} \frac{P_0 \sqrt{a}}{\kappa A} \frac{1}{\sqrt{\omega}}$$

or

$$\omega = \frac{1}{a} \frac{(4\pi)^2}{\lambda^2} \frac{\alpha^2}{\beta_{eff}^2}. \quad (3.30)$$

The form of this crossover frequency is quite interesting – it is proportional essentially to the ratio of the substrate expansion coefficient and the effective phase change of the coating (both reflection coefficient and volume expansion). This also hints at our overall strategy for measuring β_{eff} , which will involve fitting to both the low-frequency part of the response and the high-frequency part — effectively comparing the two and eliminating the need for specific knowledge of the incident power, beam area, etc.

3.3 Simulation

We performed simulations of the experiment in COMSOL (version 3.5a) to understand two effects that are harder to deal with theoretically.

The first of these is a nonzero Poisson ratio. The traditional assumption is that disturbances in the substrate are small and limited to the extent of the spot width. This means that expansion is not allowed in directions parallel to the surface, suggesting that perhaps there will be more expansion in the direction of the surface normal. In the future it would be useful to go through the theory of the temperature and elastic constitutive equations, but for now we note the form that is given in the literature for these effects, and, using the finite element model, confirm the stated dependence on the Poisson ratio.

The second effect we look to COMSOL to help us understand is boundary conditions. Although it is true that thermal disturbances will be confined to very small regions (the thermal distance scale being on the order of 10s of microns), the elastic (“sound”) wavelength is large, 10s to 100s of meters, and one might thus expect the boundary conditions (finite mirror size, clamping technique, etc.) to affect the resulting physical expansion of the mirror. We use COMSOL to help us get a handle on boundary condition effects.

Shown in Figure 3.5 is the block diagram used for our model. The thermal and mechanical properties of the top and bottom blocks were those of aluminum, and in most of our simulations the back surface of the mount (the surface farthest from the mirror) was taken to be fixed and held at room temperature. These conditions were intended to ensure we suppressed “center of mass” modes in position and temperature. Both were justified by testing with varied constraints (which made no difference) and ensuring that the assumption that the temperature and motion that far from the heating spot was negligible. The model validated all these assumptions.

Next, in figures 3.6 and 3.7 we show the grid of elements used for the simulations. We took care to have very small elements (smaller than the thermal wavelengths we used) in the region where the thermal waves would penetrate so that the effects we were interested in didn’t get washed out by averaging over too large an area.

For reference, we also include in figures 3.8 and 3.9 some graphical plotting of temperatures in the model at the end of 6 cycles at 100 Hz. The thermal wave is clearly visible, as is the fact that temperatures are stable any real distance from the

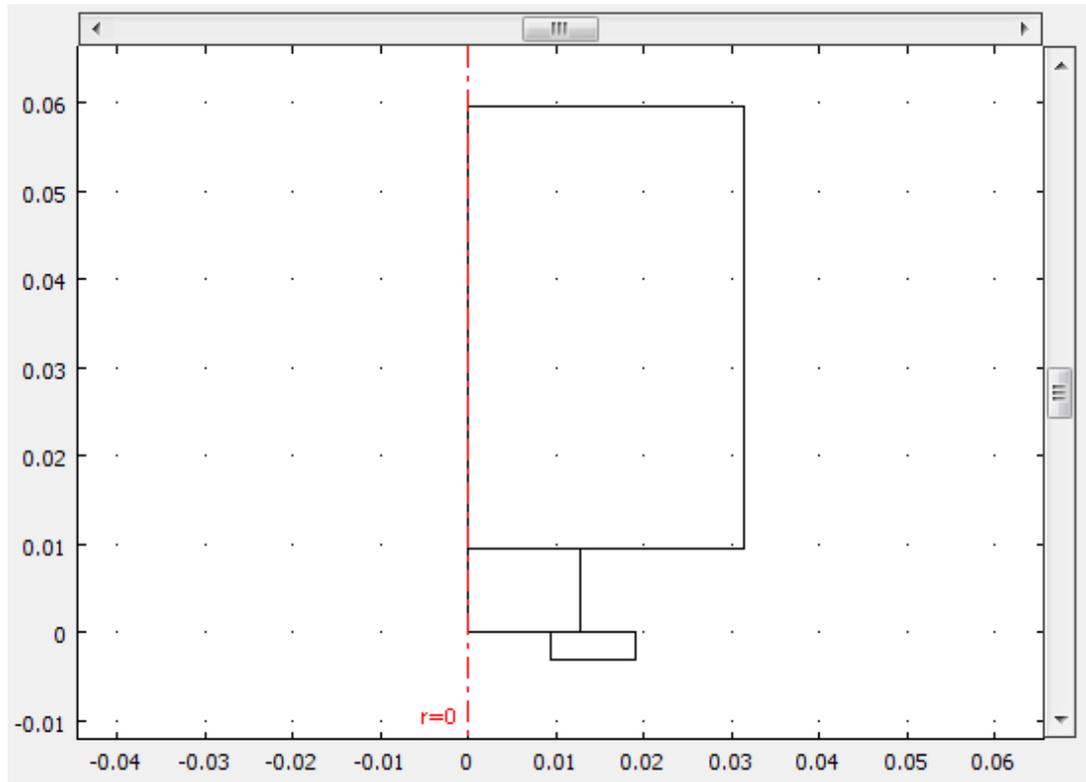


Figure 3.5: Block view of our COMSOL model. From top to bottom, the elements are the mirror mount (aluminum), the mirror itself (various materials), and the metal washer that was screwed to the mount to hold the mirror in place.

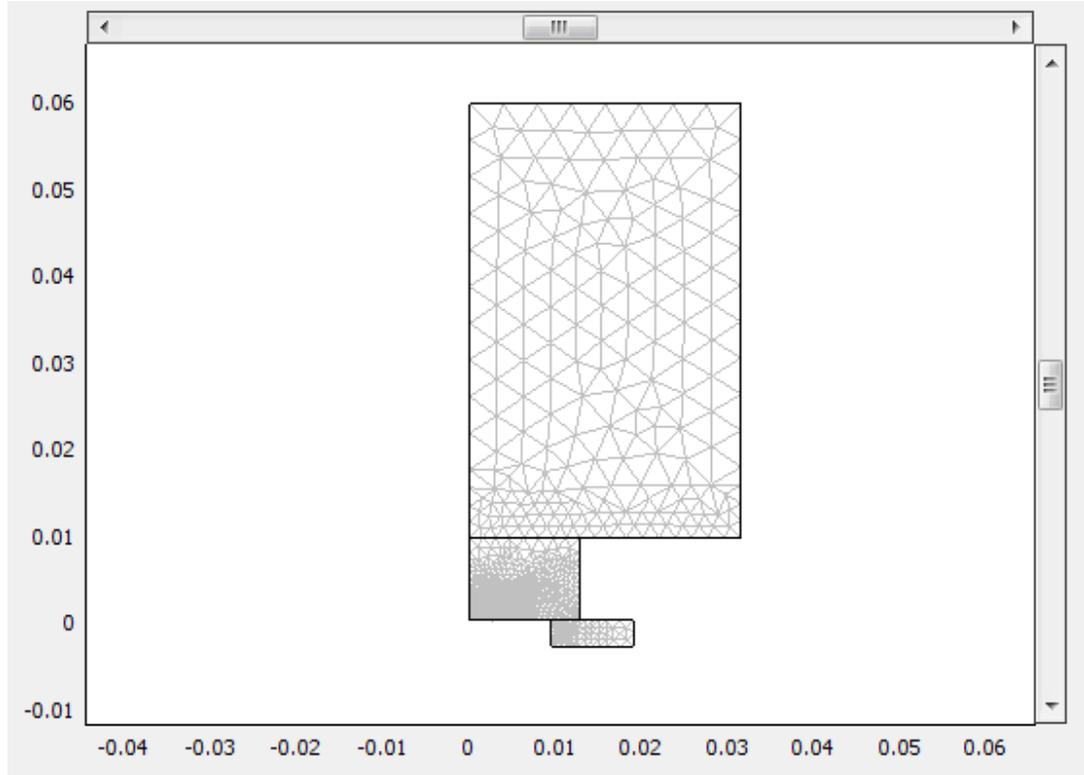


Figure 3.6: Full grid for COMSOL model

heated spot (and importantly, at the back of the mirror mount).

3.3.1 Poisson Ratio Correction

In literature where such things are accounted for carefully (e.g., [14], [10]), the factor α is frequently replaced with a corrected factor $\alpha(1 + \sigma)$. This factor comes from the fact that a small volume element, when heated, tries to expand by a factor of $\alpha\delta l$ in all three directions, but in our application is constrained in such a way that it does not move in the plane of the mirror (at infinity, there should be no motion, and this is a good approximation when the thermal wavelength is small). This suppressed expansion in plane is coupled through the Poisson ratio to extra expansion perpendicular to the plane. We did not do a proper derivation of this factor from the underlying thermoelastic equations, rather we performed a simulation of the experiment and altered the values of the Poisson ratio to test for the solution's dependence on this parameter.

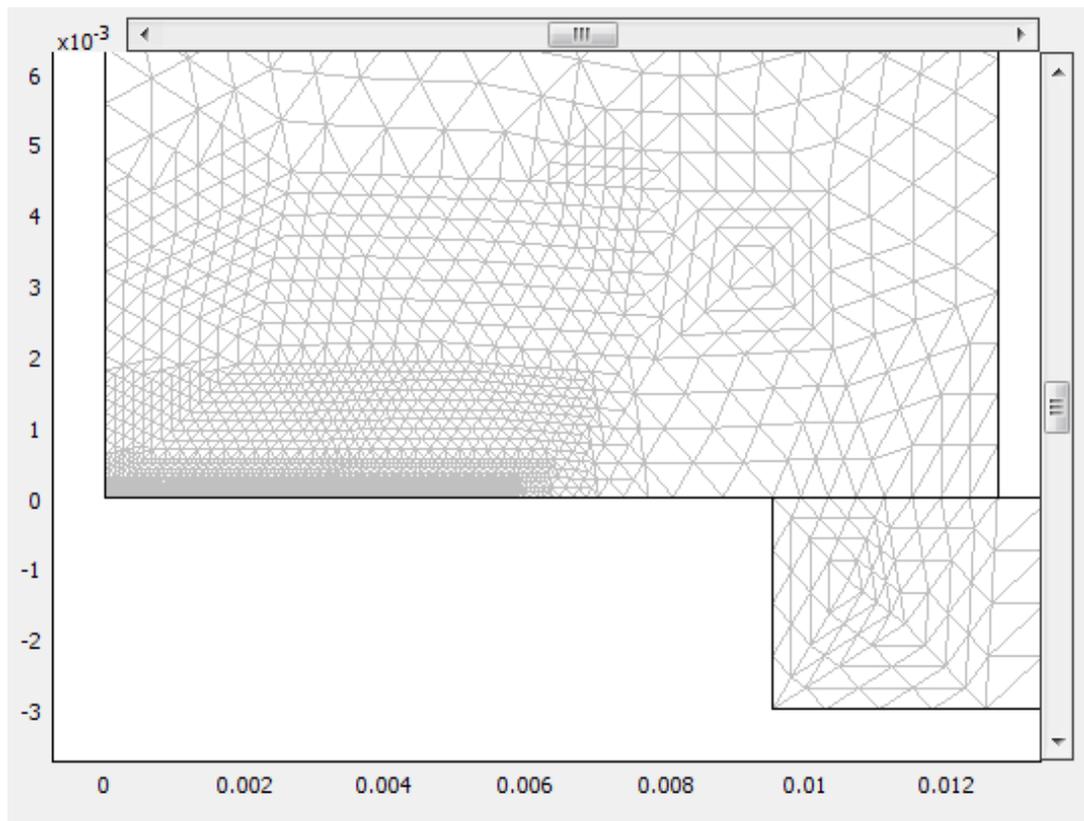


Figure 3.7: Grid for COMSOL model zoomed in on thermal spot region

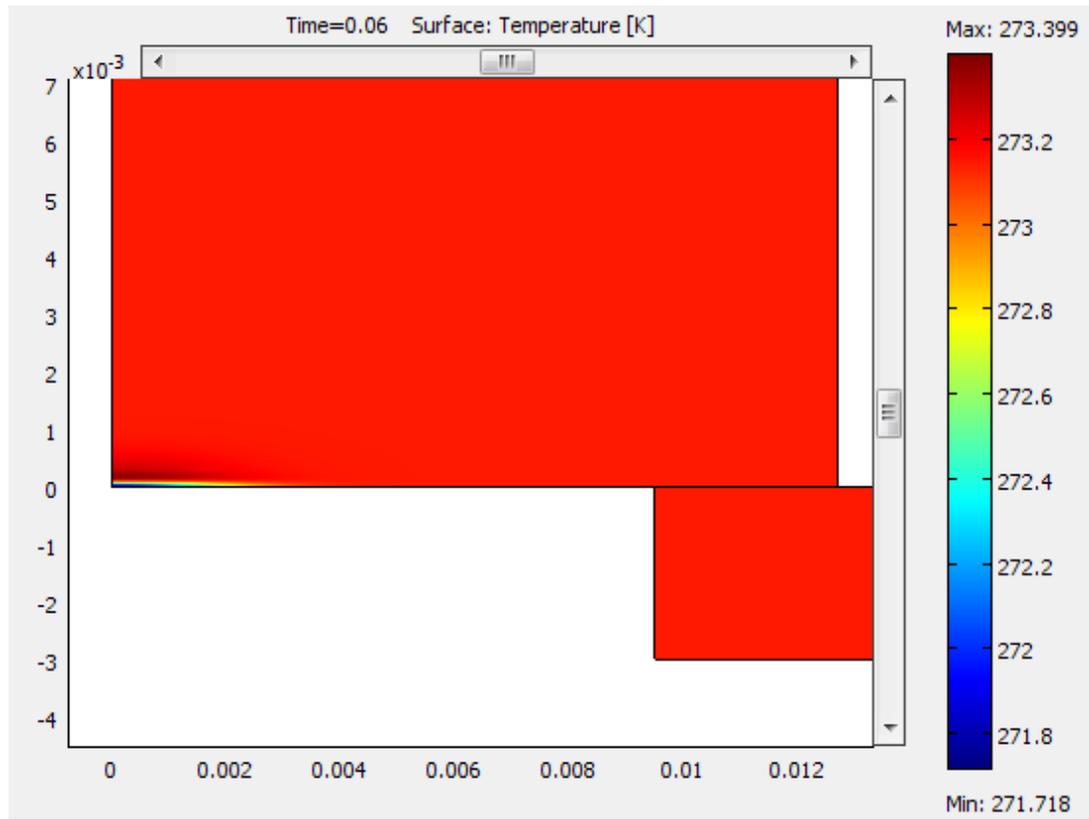


Figure 3.8: Snapshot temperature profile at the end of a simulation (after 6 cycles at 100 Hz)

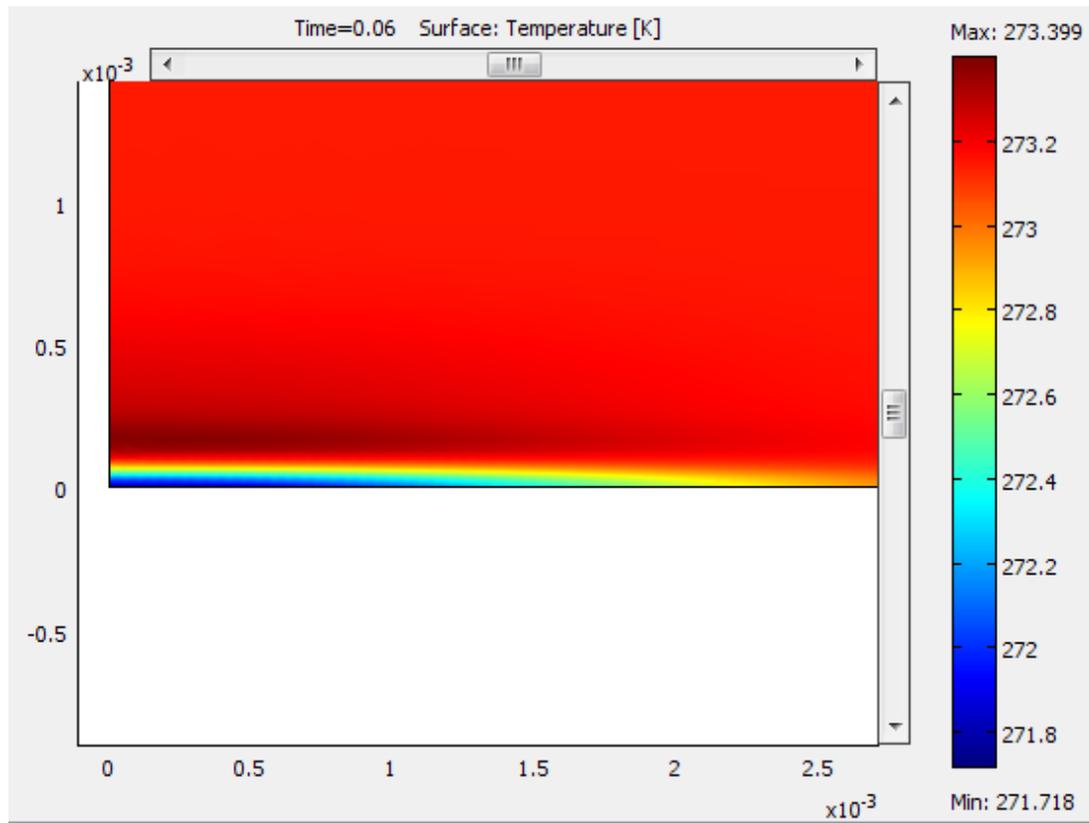


Figure 3.9: Snapshot temperature profile at the end of a simulation (after 6 cycles at 100 Hz), zoomed in on the region exhibiting thermal waves

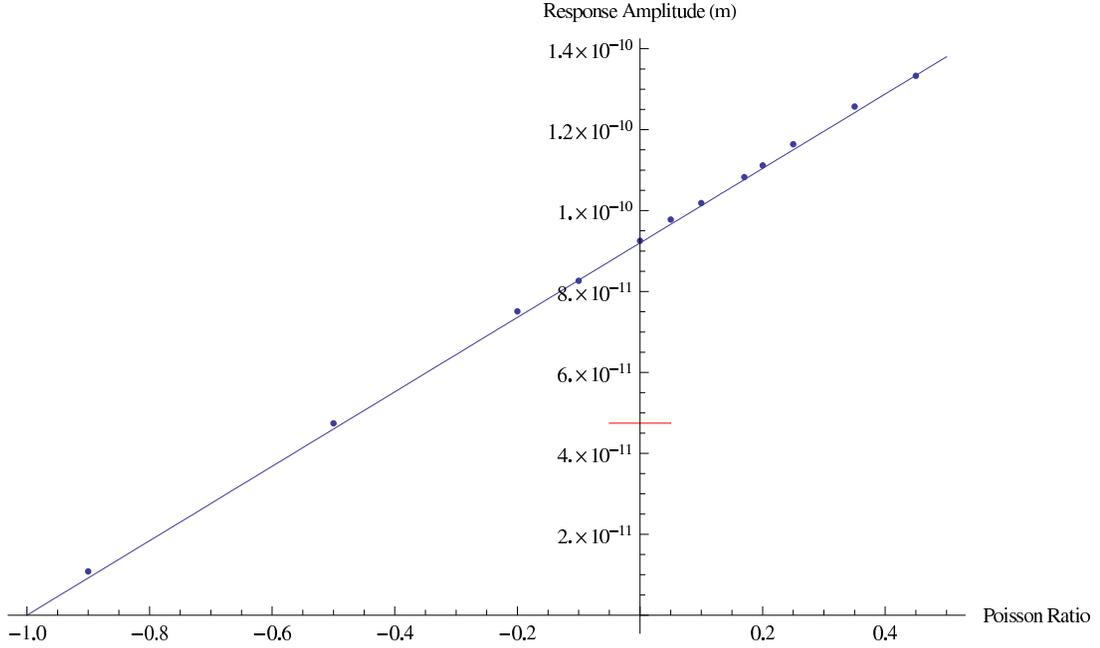


Figure 3.10: COMSOL simulated response amplitude at 100 Hz as a function of Poisson ratio. Points are simulation results, the blue line is a fit to an $\alpha(1 + \sigma)$ model.

Figure 3.10 shows results of our testing the model above with varying values of the Poisson ratio in the fused silica mirror body (all tests run at 100 Hz thermal driving frequency). The first thing we point out is that the simulations agree very well with dependence proportional to $(1 + \sigma)$. We will take this as justification to use the factor $\alpha(1 + \sigma)$ as our expansion coefficient for the substrate, and also as the parameter we measure when we look at substrate expansion.

The second thing we point out in Figure 3.10 is the red line shown — this is the level of mirror motion predicted by our theory (with zero Poisson ratio). It is, to good precision, exactly half the motion reported by the model. We currently do not understand why this is the case, but we note that asking COMSOL to report the temperature integral along the axis of symmetry and calculating

$$\delta l = \alpha \int (T - T_0) dz$$

gives the same result as our theoretical calculations (and again, half the result we get when we ask COMSOL to report the motion of the central surface point of the mirror).

We concede that this is a problem that needs to be understood, but for the purposes of this document we will assume that it is an error inherent in the understanding of COMSOL’s position reporting but that COMSOL is still otherwise a good source for understanding the relative effects of things like Poisson ratio effects and boundary condition effects. This seems to be reasonable given that the temperature integral matches our theory so well, and that the theory matches the experiment to within 20% — there is no factor of two between the theory, the temperature integral, and the experiment.

3.3.2 Boundary Conditions

3.3.2.1 Finite Boundaries

The speed of sound in glass is on the order of 4000 m/s; this means that at any of the frequencies we are interested in, the elastic response of our materials (which are of order 0.1 m) will be instantaneous throughout the entire body. This in turn means the elastic solution will “feel” its boundaries, and we should take care to check that they don’t change our solution too much.

To determine the effects of boundary conditions on our solution, we re-ran the simulation above with varied boundary conditions on the mirror, the washer, and the mount. The front surface was always free (except for the portion that was in contact with the washer, which we let COMSOL treat as an internal boundary). The other two sides of the mirror (the lateral side and the base in contact with the mount) were the ones we were interested in. In *all simulations* the back of the mount was fixed in space to suppress the center of mass motion of the system (descriptions below where the block is described as “free” mean except for this back boundary). We ran tests and looked at the magnitude of the 100 Hz response for the following sets of boundary conditions:

- Baseline; lateral mirror boundary free, washer free, aluminum block free
- Mirror lateral wall fixed, washer and mount free

- Mirror lateral wall fixed, washer body area completely fixed, mount free
- Mirror lateral wall fixed, washer and mount body area fixed
- Mirror lateral wall fixed, mount body fixed, no washer
- Mirror lateral wall free, mount body fixed, no washer
- Mirror lateral wall free, mount body free, no washer.

The first three tests, with the mount free to expand and compress, gave results identical to each other. The next three tests, with the mount body entirely fixed, gave results identical to each other and 1.4% smaller than the previous three. The final test with the mirror as free as we could imagine it without losing semblance to our experiment, returned to the value of the first three.

It is worth restating that these results also were a factor of 2 higher than we expected (to better than 2%).

It seems that our approximation that the mirror material on heating is not allowed to expand laterally is a good one, based on the results with the clamped lateral wall. The largest effect seems to be from whether or not we allow the mount to expand with the mirror or not, but even this is on the order of a few percent. Since this is a much smaller error than other sources we will come across, we will neglect boundary condition corrections in our analysis.

3.3.2.2 Clamping Pressure

In order to determine the effects of our particular mounting on the outcome of the experiment, we ran the simulation of the geometry in Figure 3.5 and applied various pressure boundary conditions to the front surface of the metal washer. First, we left the front surface free, then we tried applying a constant inward pressure of 10,000 N/m^2 (or 10 N over the surface of the washer, the total force of a 1 kg weight), then we tried 1,000,000 N/m^2 (1,000 N or 100 kg weight).

To within less than 1%, these gave no change in the Fourier component of the surface motion. We therefore conclude that the setup is quite insensitive to our clamping method.

3.4 Materials and Methods

3.4.1 Mirrors and Coatings

Mirror substrates were ordered from CVI-Melles Griot. We used various materials to allow for calibrations and checks of the coating signals as they balanced against substrates with differing thermal conductivities, densities, specific heats, and coefficients of thermal expansion; basically, we adjusted the thermal wavelength and the TE response of the substrate by changing materials. At one end, we had sapphire with a huge thermal conductivity and very large coefficient of thermal expansion, and at the other end we had Zerodur with a thermal conductivity comparable to other glasses, but a very small coefficient of thermal expansion.

The substrates were all from CVI/Melles Griot's PM (Plane Mirror) and PW (Plane Window) lines (the difference being whether the back surface was optically flat and/or slightly wedged, neither of which matter in our application). All were 1" diameter and most were .375" thick, but due to availability issues some of the samples were .25" thick and one sapphire mirror was .125" thick. All flat surfaces that we coated were specified to have transmitted wavefront error (at 633 nm) of $\lambda/10$ or less.

Coatings were applied by Thin Film Lab in Pennsylvania using electron beam deposition with materials of 99.9% purity. We realize there will likely be significant differences between the numbers we measure for these coatings and the numbers that will be measured for Advanced LIGO suppliers' coatings, but we started with a relatively inexpensive supplier for proof-of-principle.

The coating specifications were supplied by me and were designed to balance TE and TR effects in the coating. Since the experiment involves measuring the TE effect and a sum of TE and TR effects, it is desirable to have them on the same order

of magnitude. If the TE effect is too small (i.e. for very few coating layers), then this effect will be difficult to observe by itself, but if it is too large (i.e. too many coating layers), then it will be difficult to accurately subtract out of the the combined measurement. We settled on 11 doublets with a capping half-wavelength of low index material.

Both measurements give a volume-weighted average for the parameters measured. For example, in a traditional quarter-wavelength stack (where each layer is $1/4$ of an optical wavelength for light in the medium) of alternating silica and tantala, the thickness of the silica layers is 183nm, while the thickness of the tantala layers is 129nm. This means that in the cavity photothermal experiment the coefficient of thermal expansion of tantala will contribute to the overall measured coating coefficient of expansion with a weight of approximately $1/3$ ($129/312$), while the silica will contribute with a weight of approximately $2/3$ ($183/312$). A similar effect happens in the case of the combined TE/TR measurement, as outlined in section 2.3.

In order to extract coating parameters for the individual materials, we use two different coatings, both optimized to be highly reflective at $1064nm$. The first is the quarter-wave coating described above, with 11 doublets and a half-wave silica cap. The second is a so-called "1/8 - 3/8" coating, with doublets that combined still add up to $1/4$ wavelength, but with the tantala layers being $1/8$ wavelength (64.6 nm) and the silica layers being $3/8$ wavelength (275 nm), also comprised of 11 doublets and a half-wave silica cap.

3.4.2 Experimental Layout

3.4.2.1 Optical Probe Layout

The experimental layout is shown in Figure 3.11. The Lightwave/JDSU M126N-1064-200 NPRO laser is run at full output power of ~ 200 mW, and the power used in the Michelson interferometer is adjusted using a half-wave plate and polarizing beamsplitter combination. A 254-mm-focal length lens is used to roughly collimate the beam coming out of the NPRO. After the power control stage, a 356-mm-focal

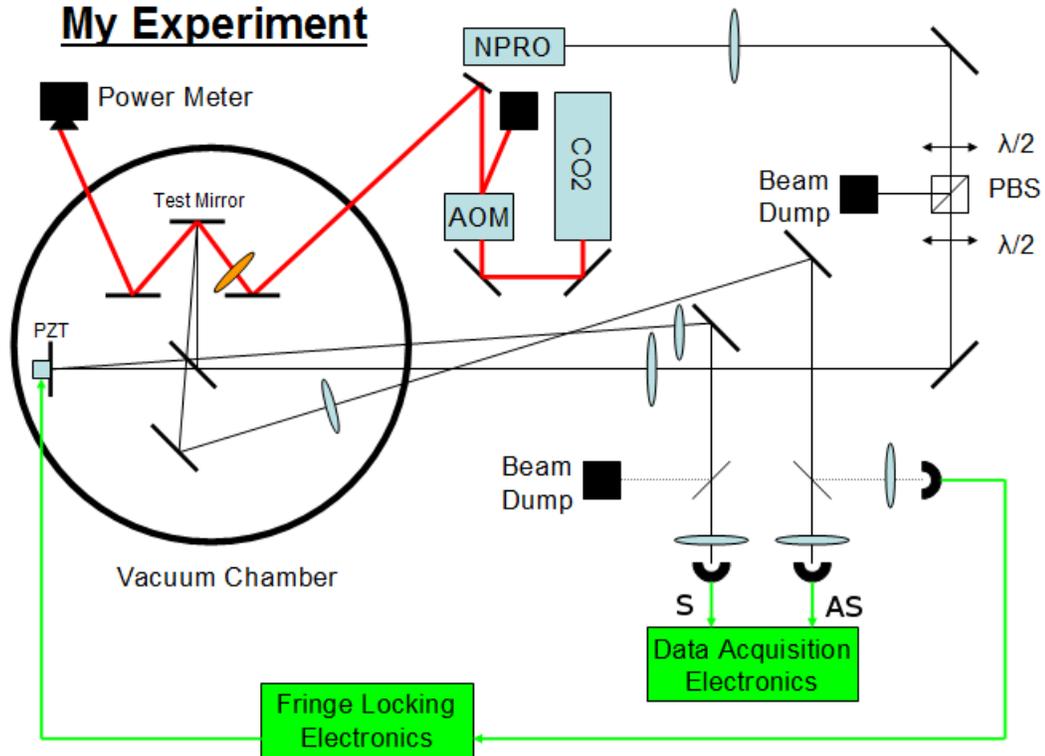


Figure 3.11: Experimental layout. Black thin lines: 1064 nm radiation, red thick lines: 10.6 μm radiation

length lens is used to create a waist 37 cm away on the beamsplitter, with radius $\cong .270$ mm. Our test mirror is then within the 21 cm Rayleigh range of this waist, so the spot on the test mirror is expected to be less than a factor of $\sqrt{2}$ larger than this, or .370 mm — thus satisfying our condition that $r_{probe} \ll r_{pump}$ (the pump radius will be on the order of 2 mm).

The optics shown in the vacuum chamber are indeed in vacuum (generally pumped down to ~ 70 mtorr), and on an optical table which sits on a 3-stage seismic isolation stack with resonant frequency on the order of 3 Hz. Electrical connecting wires are attached to the side of each level of the stack on their way to the table to reduce their contribution to mechanical coupling with the external environment.

The Michelson interferometer is very slightly misaligned so that the output symmetric port beam is not retro-reflected onto its incoming path, making it easier to measure (there is no need for a circulator, e.g., to separate the incoming and outgo-

ing beams). As was pointed out to us, this technically makes ours a Mach-Zender interferometer since the beams are not recombined at the point where they are split, but a Mach-Zender interferometer where the splitting beamsplitter and the recombining beamsplitter are just different points on the same physical object. We believe conceptually it is simpler and has just as much explanatory power to refer to our configuration as a Michelson and will continue to do so.

Once the beams are recombined they are sent to the symmetric (S) and antisymmetric (AS) ports, refocused along the way as needed, and onto the photodiodes. We used Thorlabs Det110 photodiodes, powered by 12 volts supplied by a Tektroinx CPS250 DC power supply.

In order to isolate the data signal as much as possible from extraneous coupling (e.g., through ground loops), we have a completely separate electronic path to lock the interferometer to mid-fringe. A small amount ($\sim 10\%$) of the AS signal is picked off and sent to an auxiliary locking photodiode (NewFocus 2001 photodiode box). A similar amount of light is picked off and dumped in the S optical path to preserve as much as possible the matching gains in the two paths.

3.4.2.2 Optical Pump Layout

The pump laser is a Synrad 48-1 10 W CO₂ laser with UC-2000 controller. The laser's power control is achieved by means of pulse-width modulation of the RF discharge current in the lasing gas mixture. For tests of mirrors without gold coatings, the PWM duty cycle was set at 20%, giving just over 2 W DC power in the beam exiting the laser and an AC absorbed power of approximately 0.6 W in amplitude absorbed by the test mirror. For tests of mirrors with gold coatings, the PWM duty cycle was set much higher, at 90%, giving almost 10 W DC power in the beam exiting the laser, approximately 8 W on the test mirror, of which approximately .05% is absorbed. The amount absorbed by the gold-coated mirrors was too small to measure directly with the power meter, so all we had to go on was the substrate calibration technique.

Amplitude modulation of the beam is accomplished using an AGM-406B1 infrared AOM from Intraaction Corp., driven by their GE-4030 driver. We use the

0th-order transmitted beam as our modulated signal, with the 1st-order deflected beam dumped. With a maximum efficiency of 85%, this means a decently large DC component even at the maximum amount of light deflected away, but the alternative (using the deflected beam which does go all the way to zero) is unacceptable due to the relatively large angular motion of the deflected beam during operation (presumably due to thermal effects in the AOM crystal — at the full deflection setting, there is 30 W of power being dissipated in the AOM).

The beam then travels into the vacuum chamber (through custom zinc-selenide windows from II-VI Infrared) and is focused onto the test mirror by a 7.5-inch-focal-length zinc-selenide lens. This creates a Gaussian spot on the mirror with $1/e$ power radius of ~ 2 mm. We verified the diameter of the heating spot by replacing the test mirror mount with knife edge on a micrometer in front of a power meter. The knife edge was scanned across the beam and the power transmitted recorded, then fitted to an appropriate *erfc* function, shown in Figure 3.24 and described in Section 3.4.3.3.

We could conceivably do a better job characterizing the size of the beam at various points along the path (since the beam is converging over a roughly 7.5 inch distance in the vicinity of the test mirror, displacements of the sample can make a difference in the power density deposited), but we believe that the best characterization of the spot size and absorbed power will be achieved by looking at the low-frequency response of the sample in our readout. The low-frequency regime will be dominated by the bulk properties of the substrate, which are well known, and the only unknown is the power density. Thus we take this measurement for ballpark calibration, but we get precision by fitting the low-frequency substrate response to find the power level.

After hitting the sample mirror, the reflected power ($\sim 20\%$ for stacks without gold overlayers, $\sim 99\%$ for those with) is guided out of the chamber and dumped in a power meter. This power meter is monitored during the taking of data to ensure we have a good understanding of the power on the mirror. Occasionally the CO₂ laser will mode-hop, especially if it has not been on long enough to achieve a steady thermal state, and these mode hops are accompanied by significant jumps in power output. If the subsequent jump in power is not large enough to significantly alter

the lock point of the interferometer, then the new power level is adjusted for in the data. If the jump in power is large enough to significantly move the lock point of the interferometer, we discard the data set.

A point worth mentioning is that we found it important to dump the unused pump beam off the Michelson table (somewhere on the other side of the seismic isolation). While the radiation pressure of the light is negligible, in dumping this much power there is bound to be thermal expansion which, by changing the center of mass of objects such as a beam dump, can cause vibrations. These induced vibrations can ring up mechanical modes of the power dump mount, the optical table, and the mounts for the interferometer optics, which then contaminate the output signal of the interferometer. This is a particularly devious effect because the vibrations are at the AOM drive frequency and thus can not be separated using lock-in techniques from the signal we're interested.

A similar thing can be said for any scattered CO₂ light which happens to hit interferometer optics. We found it to be important to avoid scattering pump light as much as possible (presumably this is also aided by dumping the beam off-table), and also to shield the interferometer optics as much as possible from the CO₂ beam path using foil barriers wherever possible.

3.4.2.3 Data Readout Electronics

The data readout path is as simple as we could make it. The two voltages from the S and AS photodiodes are fed directly into the A and B inputs (respectively) of an SR830 lock-in amplifier, set to measure the difference ($A - B$). The frequency to be used as the lock-in reference is taken from the reference output of the function generator that controls the CO₂ AOM driver.

The frequency of the measurement is set by a DS345 function generator which provides the 5 V synchronization signal to the lock-in amplifier, and the 0.6V peak-to-peak signal (with 0.5 V DC offset) to the AOM driver. The AOM driver has an input voltage range of 0–1 V, but we found the response of the AOM starts to become nonlinear in turn-on/saturation fashion below 0.15 V and above 0.85 V respectively,

as described in Section 3.4.3.2.

3.4.2.4 Locking Electronics

In order to linearize the readout of the Michelson interferometer, we lock the interferometer to mid-fringe using proportional gain with a low-pass filter stage. This is accomplished by feeding the difference between the auxiliary output photodiode and a reference voltage back to a PZT actuator on the end mirror in the “non-test” arm of the Michelson interferometer.

A portion ($\sim 10\%$) of the AS light output is diverted to an auxiliary locking photodiode (NewFocus model 2001, with a gain of 10 and a low-pass filter at 10 kHz). This signal is then differenced with the desired reference voltage (set on an Agilent E3644A DC power supply) that corresponds to the Michelson mid-fringe point in an SR560 preamplifier. This stage is set to have a gain of 20 and a low-pass filter that was set to -12 dB/oct at 0.1 Hz. This cutoff frequency is low enough that the effects of the feedback loop can be safely neglected in the frequency range where we are taking data (10 Hz – 5000 Hz) as shown in the next section. We chose to do this because it turned out to be nontrivial to measure (and thus divide out) the effects of the feedback loop in lock.

The next stage is another SR560 preamplifier with no filtering and unity gain, which is used to inject a swept sine wave for the purpose of taking in-loop transfer functions from time to time. When not in use, we ground the second (injection) input with a 50 ohm terminator.

The final electronic stage is a PZT driver with a gain of ~ 15 V/V. The output of the driver is connected to the PZT on the control arm of the Michelson interferometer which can change the length of that arm, which changes the power at the AS port, thus closing the loop.

3.4.2.5 The Servo Loop

In order to simplify calibration, we designed our servo loop to operate at low frequencies outside our measurement band. It is essentially a proportional feedback loop with

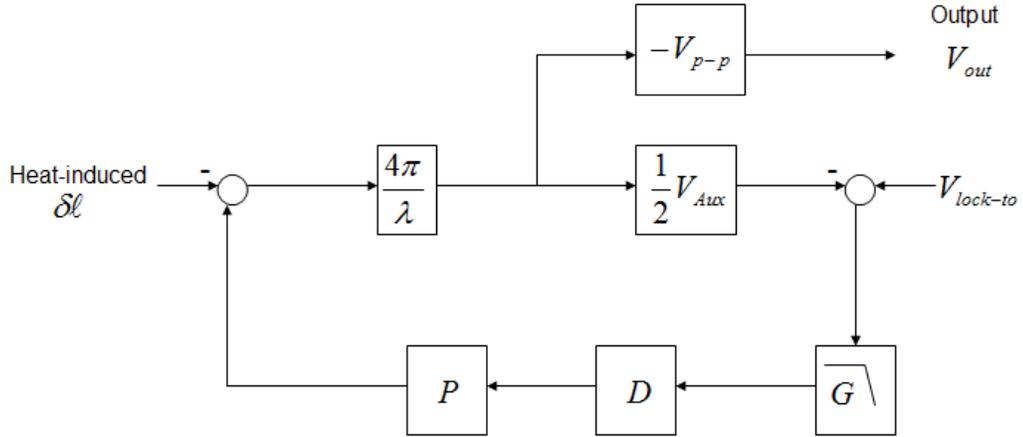


Figure 3.12: Full control loop diagram

a DC gain of order 50, with two poles at 0.1 Hz (giving us a unity gain frequency of roughly 1 Hz).

Figure 3.12 shows the schematic of our control loop. The factor $4\pi/\lambda$ is the interferometer's conversion from length change into optical phase, the factors of $-V_{p-p}$ and $(1/2)V_{Aux}$ are the readout voltages as described in Section 3.4.3.1. The “lock-to” voltage is set by hand to mid-range of the Aux photodiode readout, and is supplied by an Agilent E3644A DC power supply. Typical values for this voltage are in the range of 1 V - 1.3 V. The differencing of the photodiode voltage and the reference signal, as well as the gain G (typically 20 V/V) and a low pass filter (two poles at 0.1 Hz as mentioned) are performed in a Stanford SR560 preamplifier. The box D is a piezo driver, measured to have a flat gain of 15.1 V/V across our entire measurement band. The box P is a piezoelectric element which moves the reference mirror of the Michelson interferometer. This has a typical gain of 15 nm/V, but displays hysteresis and fairly radical variation from lock to lock (up to 50%). The difficulty this hysteresis presented to the process of exact calibration of the piezo element was our motivation for pushing the response of the servo down to low enough frequencies that it does not affect our measurement band.

Figure 3.13 shows the control loop re-written in such a way that it is clear how the

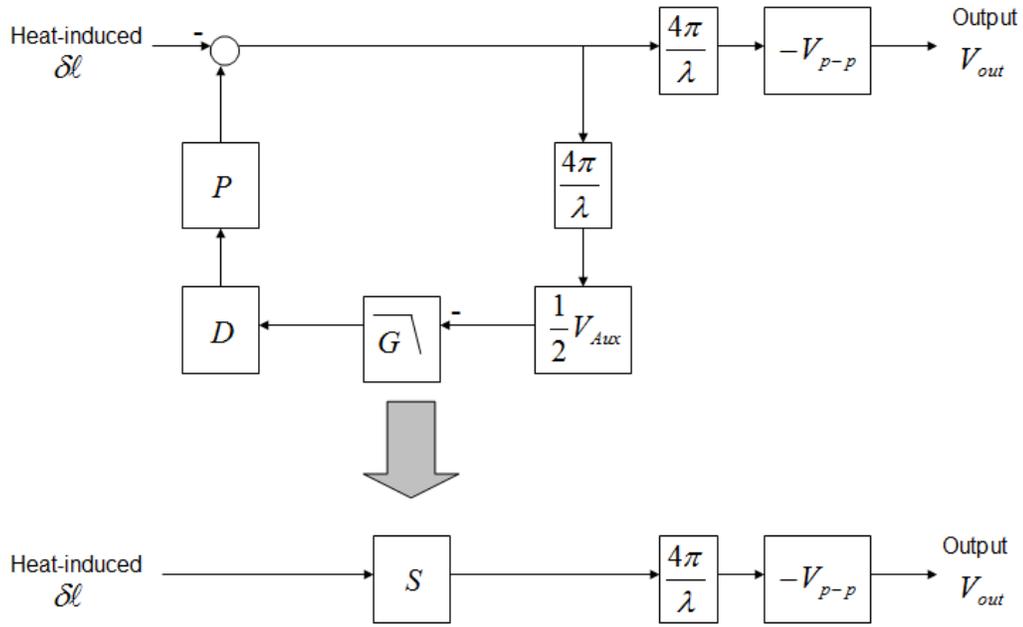


Figure 3.13: Simplified control loop schematic

servo will affect our measurement. The box S is algebraically related to the product of the loop transfer functions by “unity-forward-gain” loop equivalent factor,

$$S = \frac{1}{1 + \frac{4\pi}{\lambda} \frac{1}{2} V_{Aux} G D P}. \quad (3.31)$$

The factor G contains an aggressive low-pass filter, so the transfer function S will be close to unity by the 8 Hz start of our measurement band.

In Figures 3.14 and 3.15 we show the loop gain transfer function — i.e., the term added to 1 in the denominator of equation 3.31. The data points were measured with the two-pole low-pass filter at 10 Hz, rather than the 0.1 Hz we usually use, to facilitate the measurement and to demonstrate some of the features of this transfer function (the DC gain is visible, the corner and a table resonance at 300 Hz are visible above the noise — this is not the case when the filter is set to 0.1 Hz). The dark- and light-blue points are data taken with the loop open and closed, respectively, (the latter is an in-loop measurement of the one-way transfer function around the loop). The blue line is a model of the transfer function, and the purple line is the same

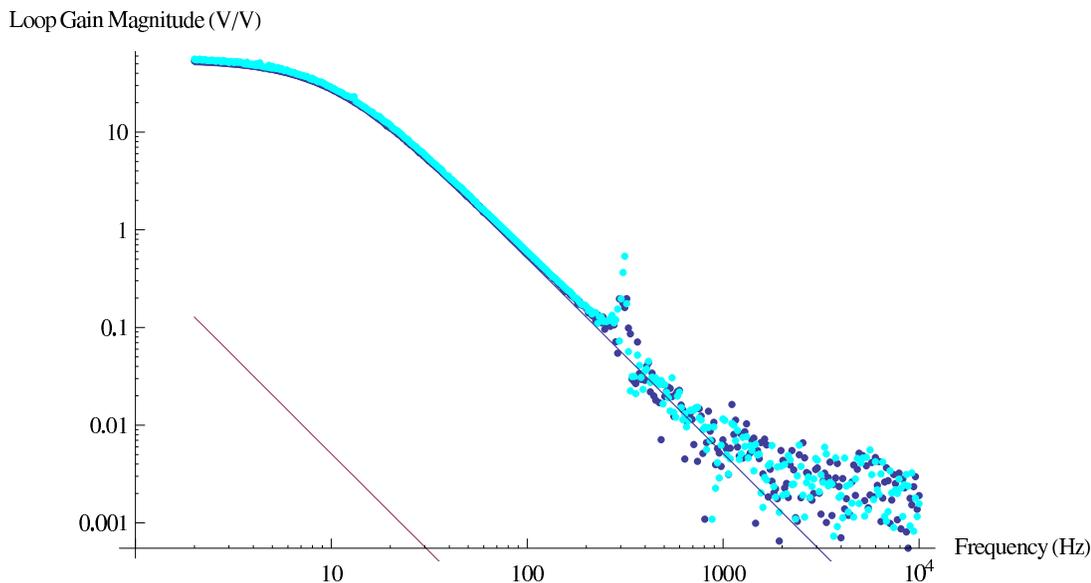


Figure 3.14: Servo loop gain magnitude

model with the cutoff frequency set to 0.1 Hz rather than 10 Hz.

The first comments we will make here regard stability. Our model gives a unity gain frequency of ~ 0.7 Hz, and an associated phase margin of less than a degree. This is close to unstable — indeed it is generally known that two-pole low-pass filters can be tough, since they achieve a nominal 180 degrees of phase asymptotically, so any extra phase (say due to a delay somewhere in the loop) can make the system unstable. However such systems are theoretically fine, and we experimentally observed that our system is quiet enough around this frequency to maintain lock for indefinitely long periods of time.

Second, we note that at 8 Hz, our model gives a loop gain of .008, and a phase of 0.1 degree. This is sufficient to state that to better than 1% we can treat the control loop transfer function S as unity over our measurement band. This low-pass filtering is sufficient to ignore the action of the servo on the measurement.

3.4.3 Rough Calibration

In this section we will describe base calibration of our signal. Aside from frequency, which we intentionally vary, there will be three experimental parameters of critical

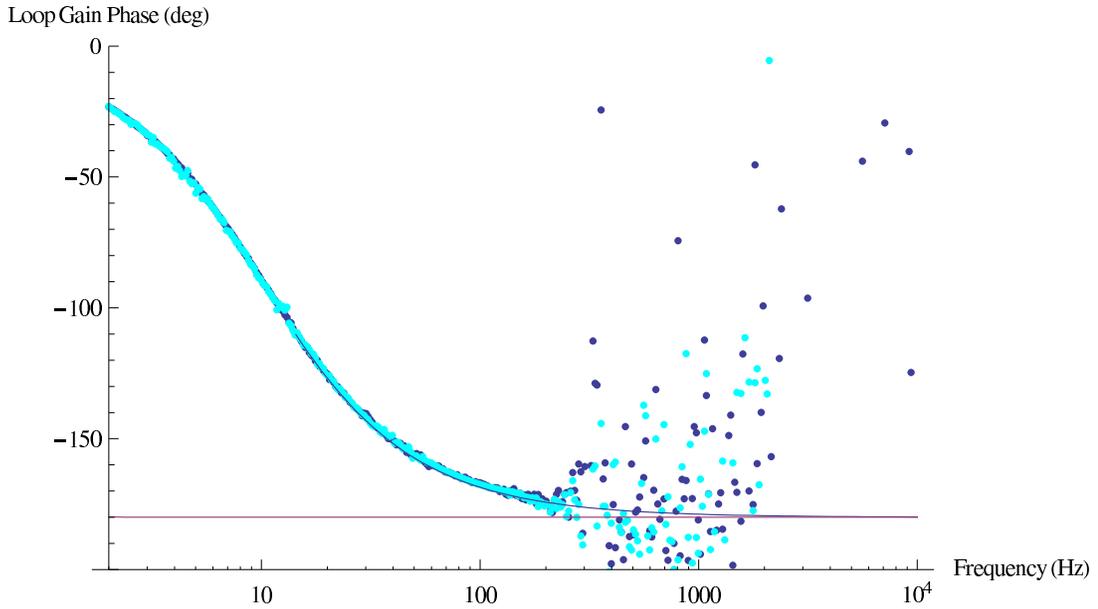


Figure 3.15: Servo loop gain phase

importance. The first two can be seen in equations 3.8 and 3.9, and are P_0 , the overall amplitude of the Fourier component of the power variations incident on the mirror, and A , the area of the heating beam. Indeed, as long as we are in the regime where λ_{therm} is much smaller than r_0 , the heating spot size, the important parameter that dictates the size of the thermal response is P_0/A , the absorbed intensity.

Unfortunately, absolute power calibrations are notoriously tricky to perform, especially in our setup where we are modulating a powerful beam with an AOM. Traditional power measurements are made with thermopile-type power meters that depend on changing the temperature of a relatively large thermal body, limiting them to very slow “DC” measurements. To get true power readings of Fourier components at high frequency it is necessary to switch to a photodiode measurement. Unfortunately, CO_2 photodiodes are extremely inefficient, so it takes a lot of power in a small area to get any signal at all, and being small devices with large power fluxes they are subject to thermal effects, especially at low frequencies. This makes it very difficult to calibrate their response at-frequency with DC power levels.

Our solution to the above problem will be to do our best with the AOM/PD combination for the power level and measure the spot size at the mirror as best

as possible with a scanning knife-edge technique, but to use these measurements as only a baseline sanity check. If we are within a few tens of percent (easily the range of combined error on power meter/PD calibration, the spot size error due to placement of the sample, and inhomogeneities in the heating beam profile) we will use this as confirmation that we understand the system, and then use the low-frequency substrate response as our power calibration. Since the signal at low frequencies will be dominated by substrate expansion, and all the quantities in equation 3.9 are well-known, except for the combination P_0/A , we will claim that the level of the low frequency response is the most accurate measurement of the heating intensity.

The third and final number we need to measure is the conversion factor of the Michelson interferometer readout. This is a very simple number to measure, and we will get that out of the way first. Also we note that although it is customary to include the fringe locking servo in the calibration section, we will not do so, since we have shown in a previous section that it does not significantly affect the output of the interferometer.

3.4.3.1 Michelson Readout

Here we describe the readout of a Michelson interferometer, showing how a change in the phase of our test arm (either through a length change or through an apparent length change due to thermo-refractive effects) will map into a change in the measured output voltage. This will also give us a formula to recreate the motion (or apparent motion) from the voltage we read out.

We can write the output of a Michelson interferometer at the symmetric (S) and antisymmetric (AS) ports as

$$\begin{aligned} P_S &= \frac{P_0}{2}(1 + \cos(\delta)) \\ P_{AS} &= \frac{P_0}{2}(1 - \cos(\delta)) \end{aligned} \tag{3.32}$$

where δ is the optical path length difference between the arms, measured in ra-

dians, and without loss of generality is taken to be zero when there is no extension/displacement of the test mirror. For pure surface motion, this phase will be given by

$$\delta = -\frac{4\pi}{\lambda}dz$$

where dz is the small mirror motion we wish to measure.

When we collect these two optical power levels with photodetectors, we will get the two voltages

$$\begin{aligned} V_S &= G_1 \frac{P_0}{2} (1 + \cos(\delta)) \\ V_{AS} &= G_2 \frac{P_0}{2} (1 - \cos(\delta)) \end{aligned} \quad (3.33)$$

and the signal voltage we measure is the difference between the two:

$$V_{Sig} = V_S - V_{AS} = (G_1 - G_2) \frac{P_0}{2} + (G_1 + G_2) \frac{P_0}{2} \cos(\delta). \quad (3.34)$$

If we equalize the gains as best we can in the two paths, $G_1 = G_2 = G$, then the first term drops out and we are left with

$$V_{Sig} = GP_0 \cos(\delta) = V_{p-p} \cos(\delta). \quad (3.35)$$

I have replaced GP_0 with V_{p-p} in the second equality, where V_{p-p} is the difference between the maximum voltage seen on a single channel and the minimum voltage seen on that channel as the interferometer is scanned through many fringes. This equivalence can be seen by looking at equation 3.33 and noticing that the minimum value of these expressions is 0 and the maximum is GP_0 .

When reading out the interferometer, we lock to the point where the cosine function goes through zero (giving us maximum change in output for a given phase change). For small variations around this operating point we have

$$\frac{\partial V_{Sig}}{\partial \delta} \Big|_{\delta=\pi/2} = -V_{p-p} \sin(\pi/2) = -V_{p-p}$$

so that for small signals we have

$$V_{Sig} = -V_{p-p} \delta = V_{p-p} \frac{4\pi}{\lambda} dz$$

and to recreate the motion (or apparent motion) of the mirror given a measured signal, we can use

$$dz = \frac{\lambda}{4\pi} \frac{1}{V_{p-p}} V_{Sig}. \quad (3.36)$$

Thus the calibration of the readout of the interferometer is simple and elegant - assuming a correctly functioning interferometer, the only number that needs to be measured is the peak-to-peak voltage on the readout channels. It is also easy to show (by going back to equation 3.34) that, should the gains be slightly unequal, the correct number to use for V_{p-p} is the average of the two values for the individual channels. Measuring this peak-to-peak voltage of the aligned (but not locked) interferometer for both channels separately and then averaging is the technique we use to calibrate the transfer function between the output of the interferometer and the actual (and/or apparent) motion of the mirror.

3.4.3.2 Pump Power Calibration

Our overall strategy for this section will be to find the conversion factor between measured DC power coming out of the chamber and P_0 , the amplitude of the Fourier component of the power absorbed by the mirror.

In operation, we will have the CO₂ laser set at a constant power output. For the full photothermal measurement, the setting was 20% full power, for the photothermal expansion (with gold topped coatings), this setting was 90% full power (in an attempt to increase absorbed power, despite the high reflectivity of the gold layer). The CO₂ achieves its power control by pulse-width-modulation (PWM) with a base

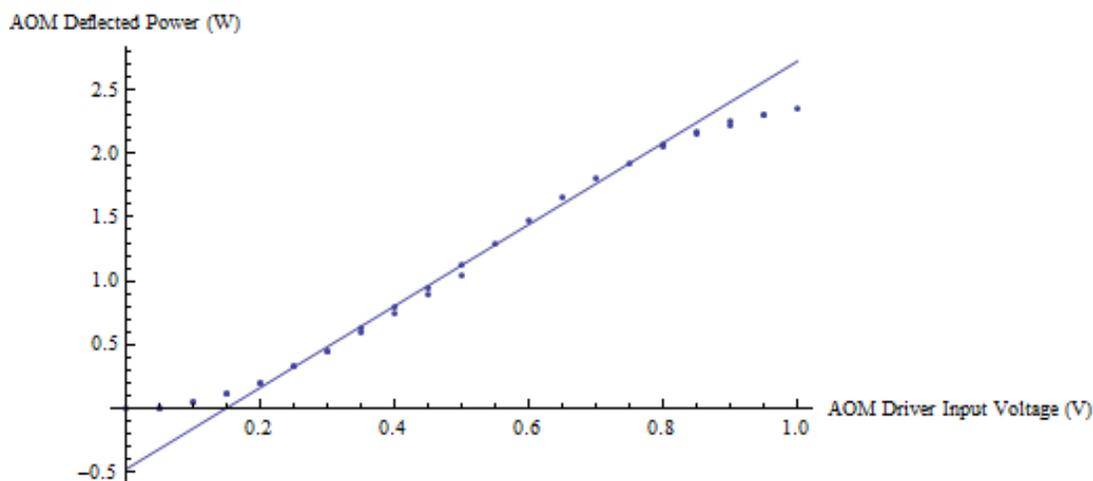


Figure 3.16: AOM deflected power vs. AOM driver input voltage, showing linearity range. Blue line is arbitrary straight line to guide the eye.

frequency of 5 kHz. This can lead to laser mode hopping accompanied by power fluctuations on the order of $1/5$ the output power of the laser — more frequent when the laser is not warmed up, far less frequent once it is. We deal with these mode hops by monitoring the power reflected off of our test mirror; if a significant jump in power occurs during data collection, that data is discarded. If not, we found the laser’s DC power output to be constant to within a few percent.

We also needed to check that the AOM’s transfer function was linear in the regime we used. Figure 3.16 shows data demonstrating good linearity of the AOM over the input range of 0.2 V - 0.8 V. Thus, we will drive the AOM with a signal from a function generator of 0.6 V peak-to-peak with a DC offset of 0.5 V, giving

$$V(t) = .5V + .3V \sin(2\pi ft).$$

This will give us a power on the mirror that has the form

$$P(t) = P_{DC} + P_0 \sin(2\pi ft).$$

Since we are going to be using a powermeter to monitor the power on the mirror, we would like to determine the ratio P_0/P_{DC} , which should not change downstream of the AOM anywhere, since the absorption coefficients of mirrors and other optics

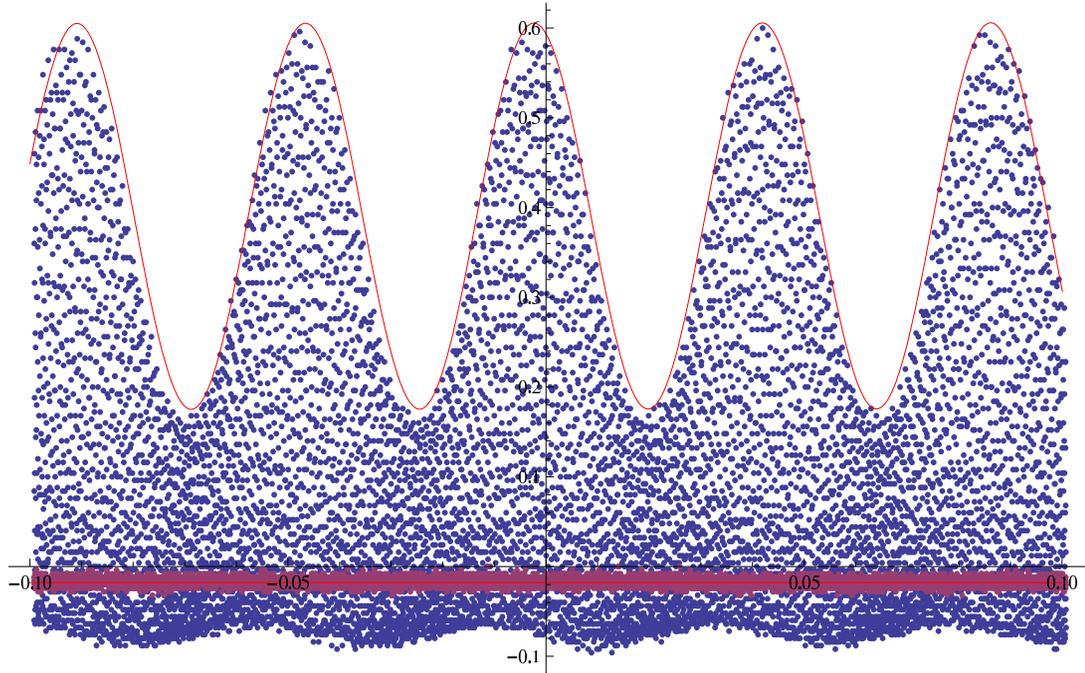


Figure 3.17: Oscilloscope trace of power as a function of time, measured with CO₂ photodiode (blue dots), with approximate sine wave fit (red line), and dark voltage (purple dots and associated flat red line)

are not frequency dependent at our measurement frequencies (which are extremely low compared to optical frequencies).

Figure 3.17 shows just such a signal, measured immediately downstream of the sample mirror with a CO₂ photodiode. The signal shown here is modulated at 22.6 Hz by the AOM, and looks “filled-in” below the AOM modulation curve because the photodetector is fast enough to also see the 5 kHz PWM of the laser.

To determine the ratio of AC power to DC power, we shift the entire signal so that the “0 V” line (measured at -0.018 V) is at zero, and write the equation that fits the red line shown in Figure 3.17. The result is

$$V(t) = .41 + .215 \cos(142t + .34)$$

which gives the ratio that we are interested in as

$$P_0/P_{DC} = .215/.408 = .53. \quad (3.37)$$

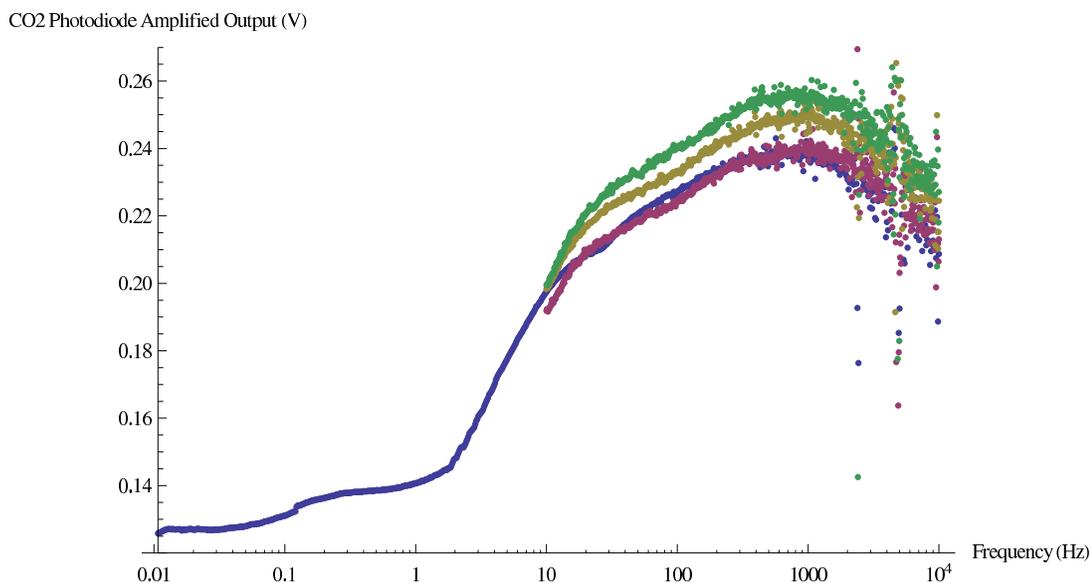


Figure 3.18: AOM+PD response magnitude, swept 5 times with exact same setup

Again, since there’s nothing downstream of the AOM to change this ratio (i.e., no absorption that acts differently on a DC power level than it would on an AC power level at these “low” frequencies), we now have a handle on how much oscillating power we have, given a measurement of the DC power present.

Next we want to refine this to account for the fact that the AOM might have a nontrivial transfer function, i.e., it might put out slightly different amplitudes or phases at different frequencies. The AOM and driver we’re using specify a modulation bandwidth of 750 kHz, and an optical rise time which gives an even faster equivalent response, but do not list specifics for things like pass-band ripple and phase shift.

First we show the transfer function from the AOM driver input to the amplified output of the photodiode, magnitude in Figure 3.18, phase in Figure 3.19. The magnitude has been adjusted to reflect the size of the voltage output we’d see from a chopped square wave with the same max power level, in anticipation of a later measurement with a chopper to determine the effects of the photodiode by itself.

A few points to make about these plots. First, there are definitely dynamics; on the order of 20% variation in magnitude across the frequency band we will be using for measurement (10 Hz – 5 kHz), and +10/-20 degrees of phase. Second, there is variation in the response, even though the measurements were taken consecutively

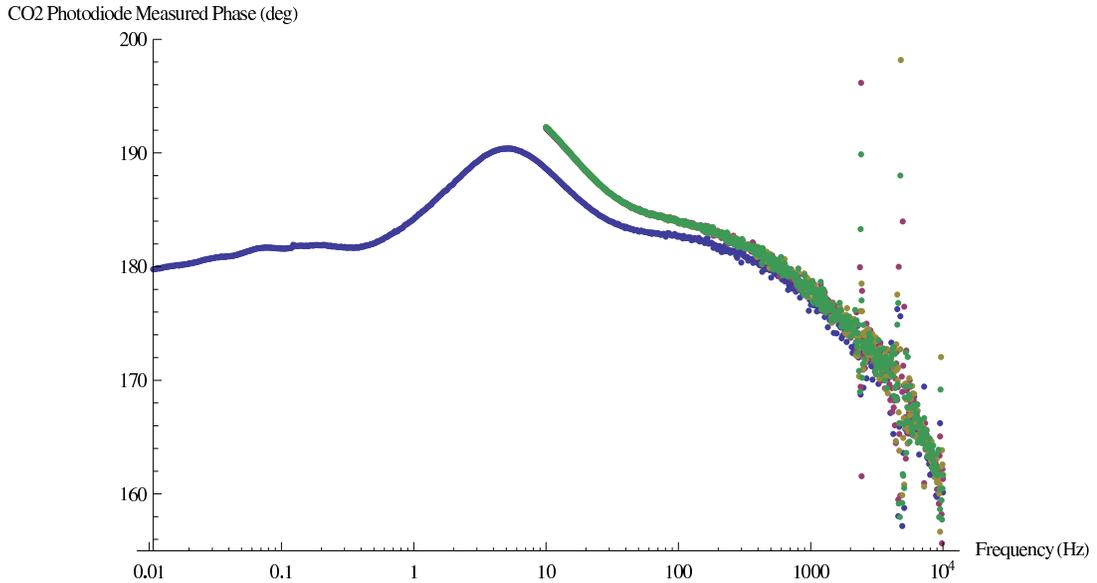


Figure 3.19: AOM+PD response phase, swept 5 times with exact same setup

with the exact same setup. There were no observed CO₂ mode hops between these data sets; we attribute the differences to thermal effects in the photodiode (which was, after all, absorbing 200 mW in a very small area). Finally, this is the combined transfer function of the AOM and the photodiode; to get the behavior of the AOM by itself, we will need to take out the effects of the photodiode's dynamics.

In order to measure the response of the photodiode, we set the CO₂ laser/AOM combination to provide a constant power (approximately equivalent to the maximum power seen by the photodiode in the previous tests), and placed a chopper in front of the AOM. We then used a lock-in amplifier to measure the amplitude of the amplified photodiode output at various chopping frequencies. Because of the limited frequency range of the chopper, we took the data in three sections. From roughly 20 Hz to just under 4 kHz we used a chopping wheel with 40 slots per revolution. From roughly 1 Hz to 100 Hz we used a hand-made chopping wheel with two slots per revolution. Below roughly 2 Hz, we performed the chopping operation by hand, letting the spectrum analyzer "control signal" tell us when to insert and remove an anodized metal plate in the CO₂ beam path. Needless to say this final method is far from precise, introducing an especially large error in the phase near and above 1 Hz (where it becomes difficult to synchronize movement of large things like arms and anodized metal plates).

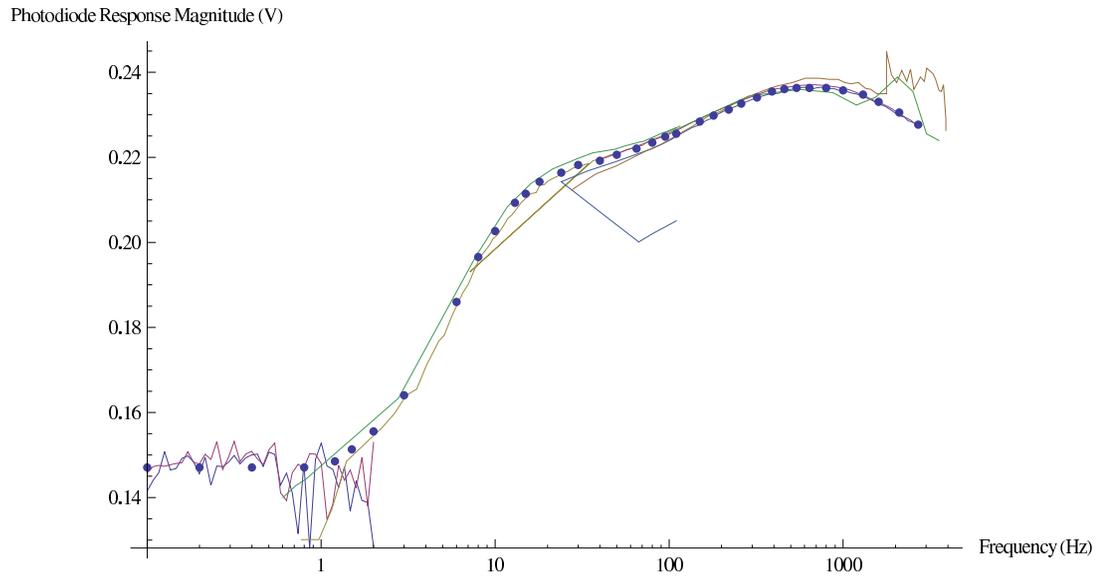


Figure 3.20: CO₂ photodiode response magnitude, measured with various chopping techniques

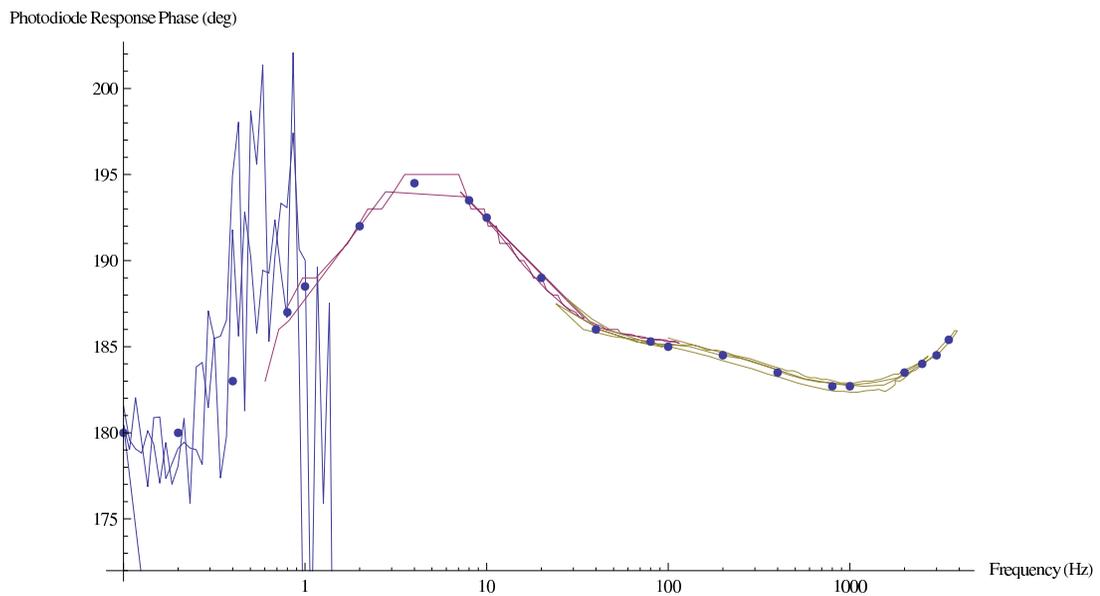


Figure 3.21: CO₂ photodiode response phase, measured with various chopping techniques

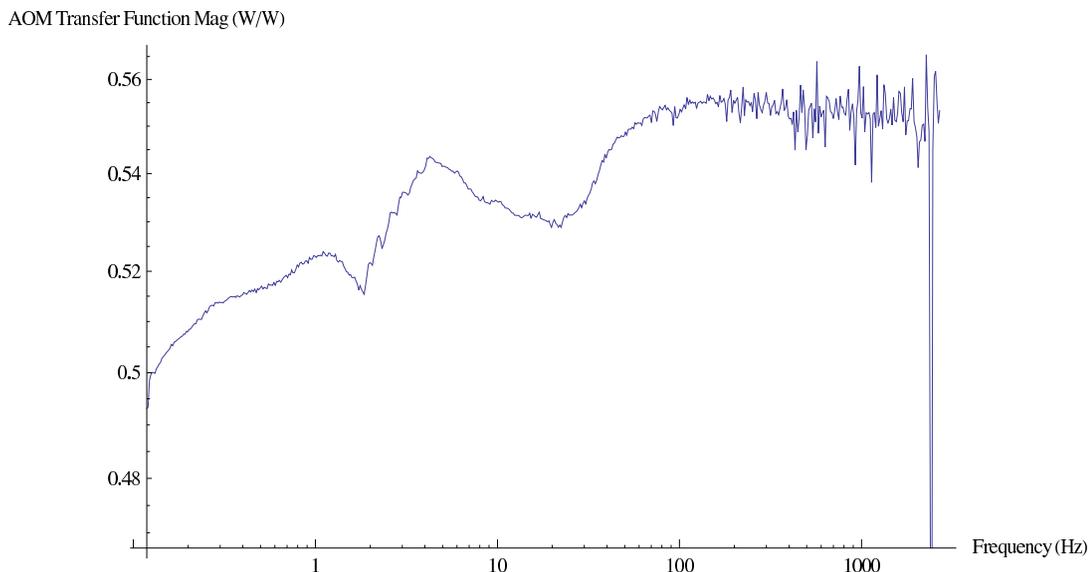


Figure 3.22: AOM transfer function magnitude, including "sine component factor"

The colored lines in Figure 3.20 show the magnitude response of the photodiode with the three chopping techniques in different colors. The magnitudes are as-measured with a lock-in amplifier for the chopping wheel data sets (though converted from RMS to amplitude measurements), and converted from a transfer function to a measured voltage with a spectrum analyzer for the lowest frequency portion.

The colored lines in Figure 3.21 show the phase response of the photodiode, also with the three chopping techniques in different colors. The noisy low-frequency regime is as-measured. The two higher-frequency regimes have been adjusted upwards or downwards to match up with the overlap in the region(s) to their left — this is simply to remove the unknown phase between the actual chopping wheel and the electronic synchronization signal sent to the lock-in amplifier.

In both figures, we also show a set of points we will use for an interpolation function to represent the photodiode response.

Finally, we show the inferred transfer function of the AOM.

Figure 3.22 shows the magnitude of the AOM transfer function, normalized to the measured AC/DC ratio from equation 3.37 (0.53 W/W at 22.6 Hz). We note two things when looking at this graph. First, fluctuations are very small — less than 10% over the entire frequency range, a few percent over most of the measurement band

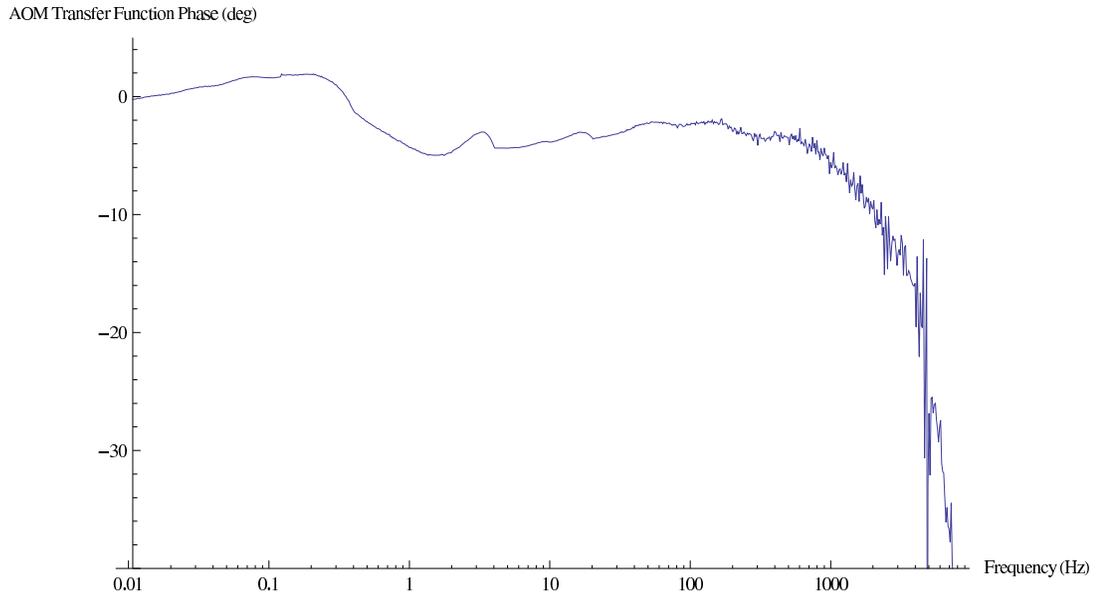


Figure 3.23: AOM transfer function phase

— and second, that suspicious behavior at lower frequencies tends to be close to the overlap regions between chopping techniques.

Figure 3.23 shows the phase of the AOM transfer function. Again there are some suspicious bumps and troughs at low frequency, especially around transition frequencies. We suspect these are artificial, and that the AOM transfer function is pretty flat with zero phase at low frequencies. There is, however, significant phase added around and above 1 kHz. This will show up in our data and need to be corrected for. As will be explained below, we hope to use these transfer functions as a sanity check and use Sapphire substrate data as a more precise calibration over the entire frequency range.

3.4.3.3 Pump Beam Spot Size

To measure the size of the pump beam on the sample, we used a scanning knife-edge technique. We replaced the mirror mount with a metal “knife edge” on a micrometer translation stage, and measured the power transmitted past the knife edge using a power meter, as a function of the position of the edge.

To determine the amount of power we should see as we scan the knife edge, we

start with a Gaussian intensity profile normalized to have a total power P_0

$$I(x, y) = \frac{P_0}{\pi r_0^2} e^{-\frac{x^2+y^2}{r_0^2}}.$$

With a knife edge at position x , the power that is not blocked will be

$$\begin{aligned} P(x) &= \int_x^\infty dx' \int_{-\infty}^\infty dy \frac{P_0}{\pi r_0^2} e^{-\frac{x'^2+y^2}{r_0^2}} \\ &= \frac{P_0}{\sqrt{\pi} r_0} \int_x^\infty e^{-\frac{x'^2}{r_0^2}} dx' \\ &= \frac{P_0}{\sqrt{\pi}} \int_{x/r_0}^\infty e^{-t^2} dt. \end{aligned}$$

This expression can be rewritten in terms of the well-known error function

$$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$$

or more succinctly in terms of the complementary error function $\operatorname{erfc}(z) = 1 - \operatorname{erf}(z)$:

$$P(x) = \frac{P_{max}}{2} \operatorname{erfc}(x/r_0). \quad (3.38)$$

Data obtained for this technique in the horizontal direction is shown in figure 3.24, along with the best fit line that gives an r_0 of 2.0 mm.

This measurement was taken in the plane of the mirror, with the knife edge scanned in the horizontal direction. In this direction the beam will be elongated due to the fact that it is coming in at a 45° angle. In reality our circular beam will have an elliptical projection on our mirror, with the profile

$$I(x, y) = \frac{P_0}{\pi r_x r_y} e^{-(x/r_x)^2 - (y/r_y)^2} \quad (3.39)$$

Assuming the beam is a symmetric Gaussian (the specification by the manufacturer has the laser's M^2 parameter less than 1.2), our x-radius and y-radius will be related by

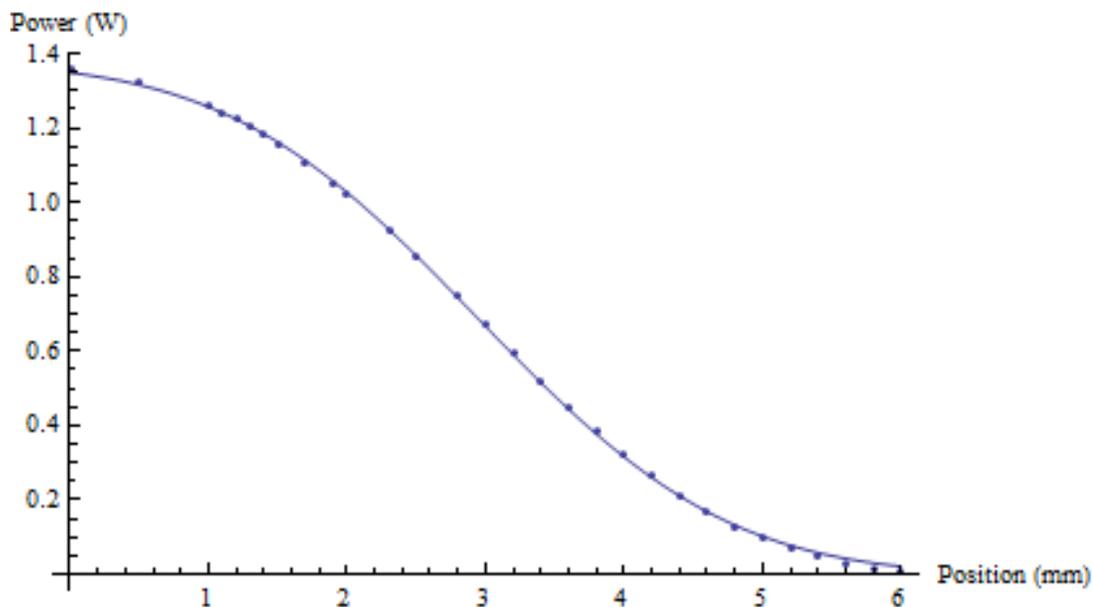


Figure 3.24: Knife edge data with fit line for $r_0 = 2.0\text{mm}$

$$r_x = r_y \sqrt{2}.$$

Or, for the purposes of making estimates of how big our signal ought to be, we will use an equivalent r_0^{eff} that is the geometric mean of these two and will give the same intensity at our probe spot. That radius will be, in terms of the x-radius we measure

$$r_0^{eff} = \sqrt{r_x r_y} = \sqrt{r_x \frac{r_x}{\sqrt{2}}}. \quad (3.40)$$

To estimate the error in our measurement of the spot size, we note that the beam is converging from approximately 4mm at the focusing lens to a small waist approximately 7.5 inches (190 mm, the focal length of the lens) from the lens. This means that the radius is changing at a rate of $4\text{mm}/190\text{mm} = .02\text{mm}/\text{mm}$. We estimate that the error in our knife-edge placement (as well as the error inherent in swapping out mirrors) is on the order of a few mm, which implies a spot size error of a few hundred microns, or a few percent, which means up to $\sim 10\%$ error in our estimate of the spot intensity.

3.4.4 Sapphire Substrate Response Calibration

Due to the difficulties listed above in accurately measuring the absolute power per unit area on our substrate, and its phase and relative magnitude as a function of frequency, we will use a technique we believe to be much more accurate and reliable in processing our data; we will use substrate responses as our overall power calibration, and to check the effects of the AOM and driver chain.

3.4.4.1 Intensity via LF Substrate Response

As was hinted at previously, we are going to do an overall power calibration by using substrates such as fused silica, in which we will see the effects of both terms in equation 3.29. We will fit to two parameters, P_0 and β_{eff} . The first parameter is essentially equivalent to fitting to an overall signal magnitude (which will give us an accurate calibration of power per unit area at the probe spot if the CO₂ spot is nonuniform, high or low absorption if the mirror absorption is slightly different at running temperature, and even overall factors like error in the Michelson readout calibration). The second parameter is equivalent to fitting the crossover frequency in 3.30, and will give us information about the coating response.

This technique does not work if the drive power changes with frequency, thus we use another technique to give us the frequency dependence of the heating power, which we divide out before fitting to 3.29.

3.4.4.2 Sapphire Response Calibration Function

Given the difficulties in measuring the AOM transfer function (combined with its driver), we decided to use the response of a sample mirror on a sapphire substrate as our “standard candle” for calibration purposes. Since sapphire does have a long thermal wavelength, we will need to include the Green’s function correction for the low-frequency region.

Figure 3.25 shows the raw sapphire data (uncorrected for AOM transfer function, but otherwise processed as normal data) along with a predicted value which is scaled

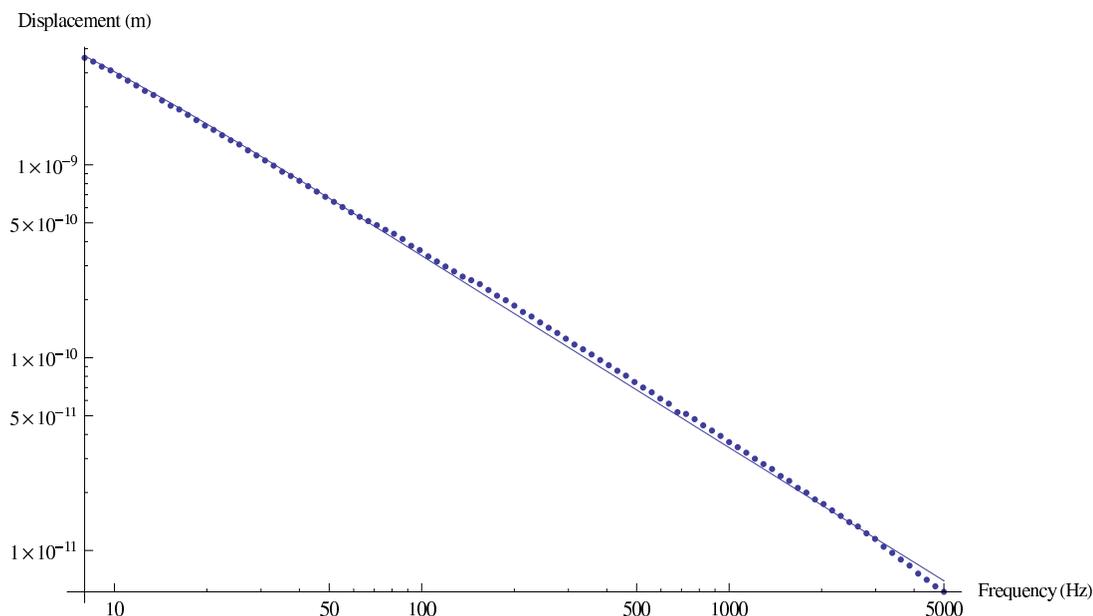


Figure 3.25: Sapphire response magnitude, raw data (uncorrected for AOM transmission) - dots, prediction - solid line.

to agree at 22.6 Hz, where we have our measured DC-Fourier component calibration from Section 3.4.3.2. From here we will infer our relative correction factor induced by the AOM (recall that we expect our overall magnitude factor to be calibrated out in the final fit, and in the end it won't play a role in determining coating response).

Figure 3.26 shows the two AOM correction functions. They show very similar behaviors, with an increase in the transmitted magnitude somewhere between 30 and 100 Hz, to 5% or so for the PD method, 10% or so with the sapphire calibration method.

Figure 3.27 shows the raw sapphire phase data (uncorrected for AOM transfer function, but otherwise processed as normal data) along with a predicted value. From here we will infer our relative phase added by the AOM. It is worth noting that the phase added at low frequencies due to the Green's function "low-frequency correction" shows remarkably good agreement.

Figure 3.28 shows the two AOM phase correction functions. Again, they show similar behaviors, with little to no phase added at low frequencies, and a sharp drop in phase above about 1 kHz. Again there is slight disagreement, but we suspect

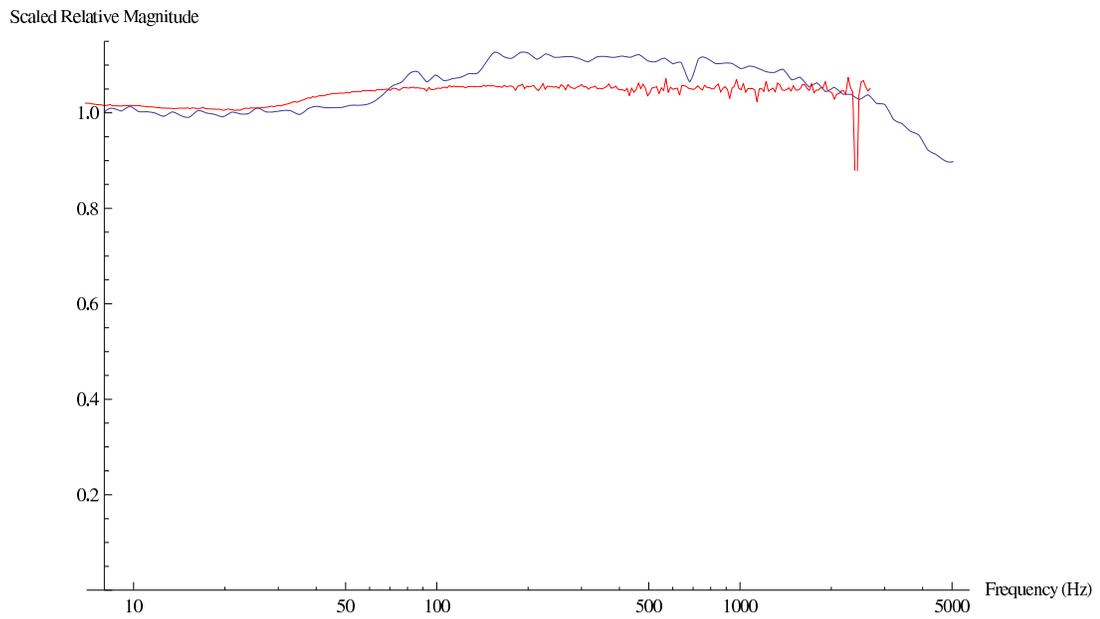


Figure 3.26: Sapphire magnitude correction: sapphire reference method (blue), AOM/photodiode calibration (red)

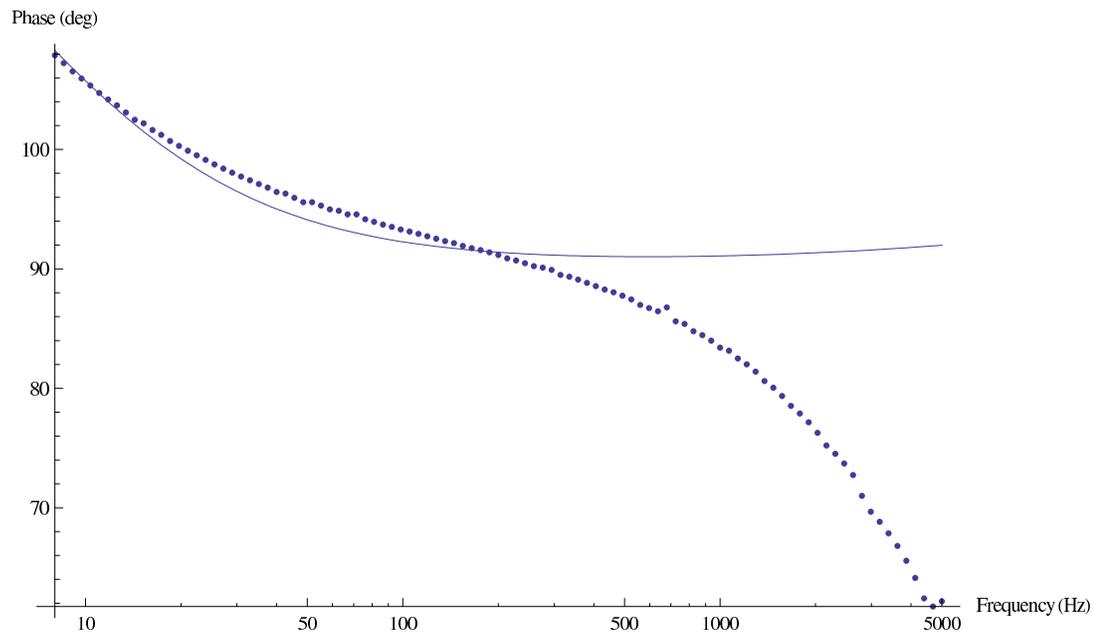


Figure 3.27: Sapphire response phase: raw data (uncorrected for AOM transmission) (dots), prediction (solid line)

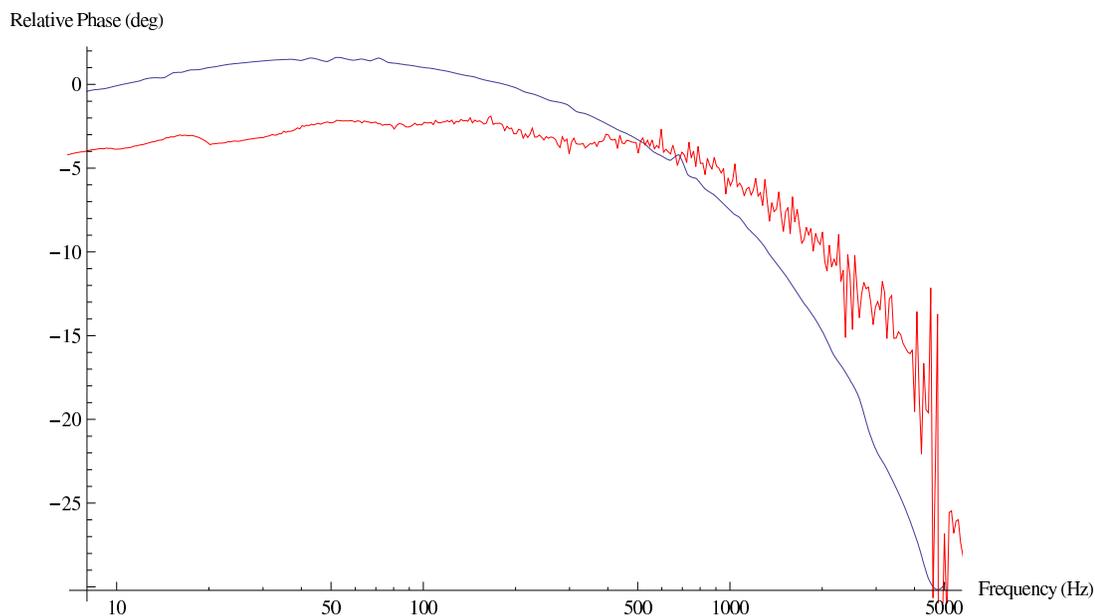


Figure 3.28: Sapphire phase correction: sapphire reference method (blue), AOM/photodiode calibration (red)

this is due to thermal difficulties in the photodiode and the inherent difficulties in determining the phase of a chopping wheel (which depends on spot size relative to chopping wheel slot width), and believe the sapphire will be the better guide.

For the rest of the data processing in this paper we will use the correction functions in blue to remove the frequency-dependent magnitude and phase effects of the AOM from the data we take.

3.4.5 Data Acquisition

3.4.5.1 Laser Alignment

The first step to taking data is to insert the mirror to be measured and align the Michelson interferometer (ensuring both beams are centered on the photodiode and then tweaking alignment for maximum visibility of fringes). We easily achieved greater than 90% visibility.

Next, we test the overlap of the heating beam with the probe beam. For visualizing the position of the CO₂ beam we use fluorescent viewing cards made by Macken Instruments. These cards fluoresce when exposed to the UV lamp that comes with

the kit, but the CO₂ laser disrupts the fluorescence. Thus, when the CO₂ laser hits the card a dark spot appears. Fortunately this same effect is observed with the probe laser, so both are visible on the card at the same time. We ensure that the spot centers match vertically, and are appropriately offset horizontally. Due to geometry, we can not place the fluorescing surface on the surface of the mirror, and because the CO₂ laser is approaching at a 45° angle, its center on the card will not coincide with the center of the heating beam. We carefully measured that the displacement of the fluorescent surface is 8 mm in front of the mirror surface, so at that point the CO₂ beam properly aligned needs to be 8 mm offset horizontally.

In the early days of the experiment we routinely checked (with a thermal imaging camera) to ensure that the CO₂ light was not scattering off the edge of the mirror mount. This was never seen to be a problem, so we do not routinely check this in the data sets reported here.

3.4.5.2 Warming Up and Pumping Down

The next step was to pump the air out of the chamber and allow the sample to reach a steady thermal state. We routinely operated around 70 mtorr with a turbopump backed by a rotary vane roughing pump.

While pumping down, we would set the CO₂ laser and AOM to their operational settings, as if we were taking data, to let the equipment and mirror warm up. By pumping down without the heating laser on, locking the servo, and then turning on the heating beam and watching the lock-to point change over time, we were able to determine that the setup (with a fused silica mirror) has a thermal time constant of approximately 10 minutes (see Figure 3.29). Other glassy mirrors would have similar time constants, and sapphire would equilibrate much faster.

Given the thermal time constant of 10 min, we would turn on the vacuum pump and lasers, pump down, and then allow another 30–40 minutes of warm up time before taking data.

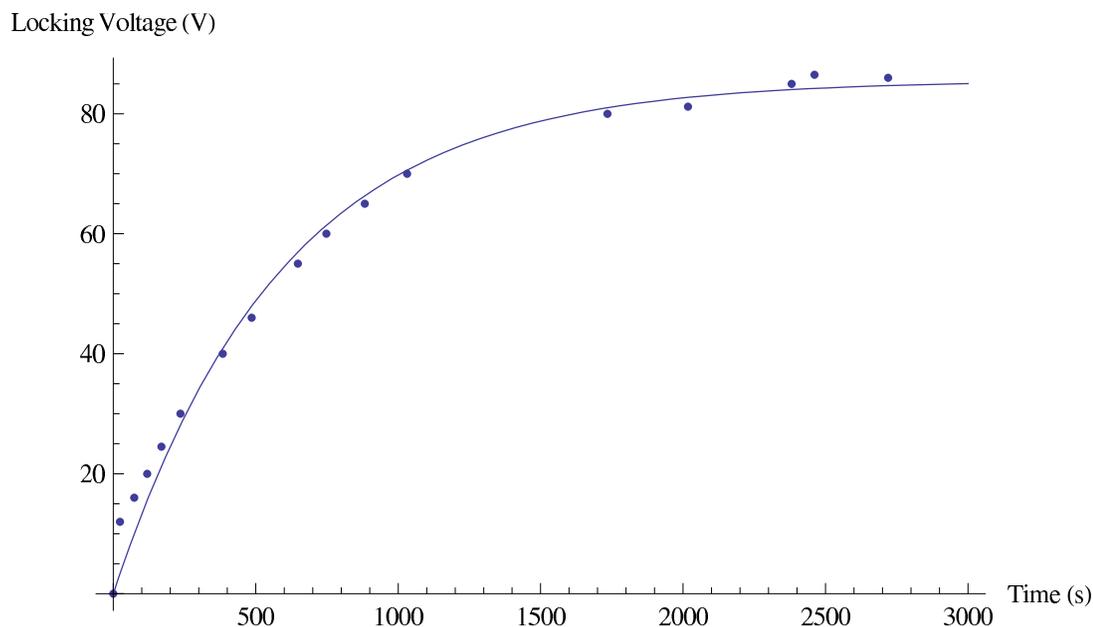


Figure 3.29: Servo locking voltage over time as sample heats up. Also shown is an exponential fit with a time constant of 590 seconds.

3.4.5.3 Pre-Run Data

Once the apparatus and mirror were warm, we would fill out a data sheet in our “T.O. Experiment Auxiliary Data” binder. This binder is filled with forms with questions designed to ensure that all relevant information about particular runs are recorded. The important details about the sample are there (what substrate, what number, etc.), as well as a record of CO₂ power levels, servo settings, lock-in settings, and, critically, the visibility on each channel (for calibrating the output signal in terms of effective displacement - see Section 3.4.3.1).

3.4.5.4 Data Acquisition Routine

Data was taken using a laptop running Windows and LabView to coordinate a lock-in amplifier (which took the data), a frequency generator (which provided the modulation signal to the AOM driver and a synchronization signal to the lock-in), and an ADC which measured the output of a power meter to monitor the power level of the CO₂ laser (looking at the power in the reflected beam after it left the sample and exited the chamber).

Frequencies for data points are selected by giving the program a low and high frequency, and a number of data points desired — the program then distributes the measurement frequencies exponentially so that they are evenly spread on a log-log plot. The user can also specify a number of independent measurements per frequency if he/she wants more than one reading per frequency. We generally used a range of 8 Hz – 5000 Hz with 100 frequency points and 1 or 2 measurements per frequency.

After switching from one frequency to the next, the program waits a specific time before taking its first data point, and waits that amount of time again before taking multiple subsequent data points. We set these wait times to be the required settling time as detailed in the lock-in amplifier’s manual. For the 12 dB/oct low-pass filtering we used, the settling time was 7 times the lock-in time constant, or 21 seconds for our most common setting of a 3 second time constant.

With these values, a run took a little longer than an hour. We also wrote a looping program that would run the above multiple times in succession. For the combined α/β measurements which had relatively large signal we would use 3-4 runs for our data, mostly to ensure consistency, since the data was very clear. For the α alone measurements we frequently wound up using 10–15 runs, which were then processed separately to give us an estimate of our statistical error.

3.5 Results

In this section we will list the results of the individual measurements. First we will describe how data recorded directly from the experiment was transformed into apparent length changes in the interferometer arm, then we will show results first from the TE/TR combined measurements and then from the TE alone (gold-coated sample) measurements, along with sample data sets and model fits.

3.5.1 Data Processing

The data recorded from this experiment was in the form of a list of data points with three parameters: a voltage magnitude (the RMS value of the Fourier component of

interest), phase, and a reading of the DC powermeter.

In converting the voltage data to length data, we take the following steps:

1. Convert the RMS value from the lock-in amplifier into an amplitude by multiplying by $\sqrt{2}$.
2. Divide out the AOM transfer function calculated as above, normalized to unity magnitude at 22.6 Hz (divide the voltage magnitude by the value of the AOM magnitude transfer function, and subtract the phase of the transfer function from the phase of the measurement).
3. Divide by the amplitude of the Fourier component of the power delivered to the mirror, using factors measured as in the calibration section above as well as the power value of the data point. This does two things. First it removes the effect of power fluctuations (which we allow as long as they aren't large). Second, it normalizes the response to what we would see if we applied 1 W of AC power, so that we can compare the results of multiple runs even if they had somewhat different power levels.
4. Multiply by a conversion factor from voltage to interferometer length, using measured voltages as in Section 3.4.3.1. Again, the conversion factor is $\delta l = (\lambda/4\pi V_{p-p}^{ave})\delta V$.
5. Subtract 19° of phase that is artificially inserted between the function generator's output and the measured phase of the lock-in amplifier with respect to its reference input.
6. Cut out data with known issues (high-frequency spurious couplings, known noisy areas around 60 Hz and just below 20 Hz for the low-amplitude gold-coated measurements).

Once we have our data cleaned up in this way, we fit the model to extract the two parameters that define it, P_0 and β_{eff} . To simplify the fitting, we use only

the magnitude data for fitting, and then compare the phase prediction using the fit parameters with the phase data as a consistency check.

Our data spans multiple orders of magnitude, so early attempts using a traditional fitting function tended to ignore the data at high frequencies (small magnitudes); unfortunately these are exactly the data points that carry information about the coating response.

Our solution to this problem was to take the logarithm (base 10, though any base would do) of the data and do a fit to a new function that is the logarithm of the old function. By defining our errors logarithmically, we effectively ask the fitting function to minimize the sum (or more accurately sum of squares) of the percent differences between the data and model.

To perform the fits, we used Mathematica’s nonlinear “FindFit” function with the “Levenberg-Marquardt” method specified.

3.5.2 Combined TE/TR

Here we will show sample data for the combined measurement, where we are seeing the effects of both the coating expansion and the index of refraction changes.

Figure 3.30 shows sample data from a silica substrate with a quarter wavelength coating. The fit for this particular data set finds $P_0 = 1.1W$ and $\beta = .0001188$. Our average for four such data sets was .000118. Figure 3.31 shows a sample plotting of the phase recorded in our data, along with a prediction using the fit parameters from the magnitude data. Since we are within a few degrees over most of the measurement range, we believe our measurements are seeing the coating response we are interested in.

Figure 3.32 shows sample data from a silica substrate with a Bragg coating. The fit for this particular data set finds $P_0 = 0.77W$ and $\beta = .0001079$. Our average for four such data sets was .000108.

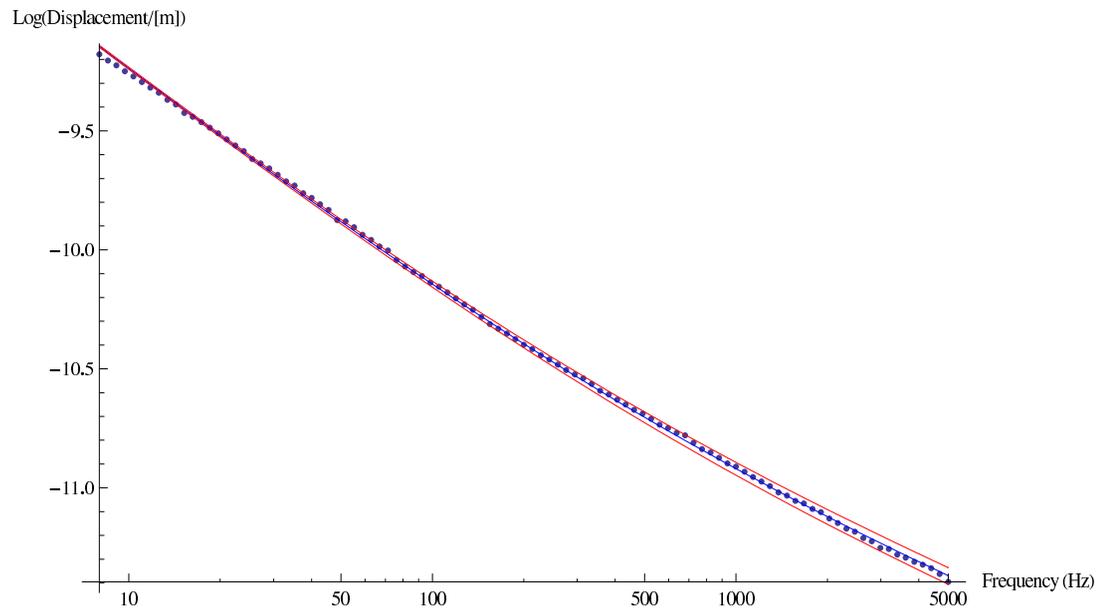


Figure 3.30: Sample silica substrate data with quarter-wave coating. Blue line is fit, red lines are “chi-by-eye” error bars (roughly double statistical errors).

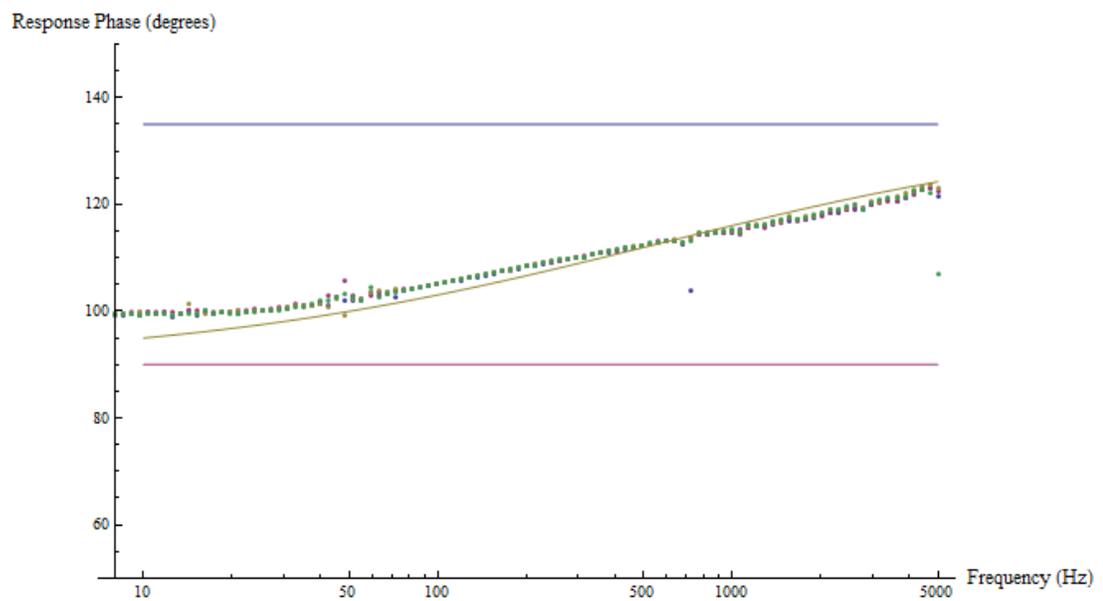


Figure 3.31: Silica QWL phase using fit parameters from magnitude data

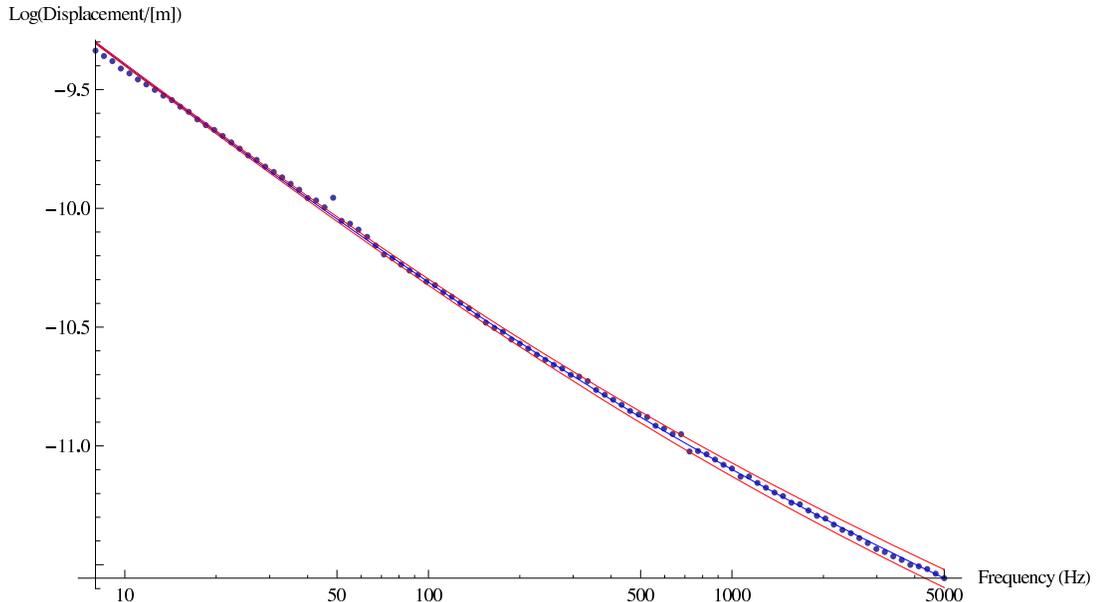


Figure 3.32: Sample silica substrate data with Bragg coating. Blue line is fit, red lines are “chi-by-eye” error bars (roughly double statistical errors).

3.5.3 TE Alone

Figure 3.33 shows sample data from an Infrasil substrate with a QWL coating and full $0.9 \mu\text{m}$ chromium layer. The fit for this particular data set finds $P_0 = .030$ and $\beta = .000337$. Figures 3.34 and 3.35 show histograms of the fitted β_{eff} for the same Infrasil sample — between the two sets of runs we removed the sample took measurements of other samples, and reinstalled it. The inferred values and errors for the two figures are $(3.2 \pm 0.11) \times 10^{-4}$ and $(3.5 \pm .15) \times 10^{-4}$, respectively. As discussed in the following section on error, we have what we believe to be a random systematic uncertainty displayed here which needs to be addressed in the future. For the purposes of this thesis we will claim the combination of these numbers to give $(3.3 \pm 0.2) \times 10^{-4}$, and also attribute this larger uncertainty in future gold-coated data sets (rather than the smaller uncertainty we measured within a set of runs).

Figure 3.36 shows sample data from an infrasil substrate with a Bragg coating and full $0.9 \mu\text{m}$ chromium layer. The fit for this particular data set finds $P_0 = .0169$ and $\beta = .0003118$. Figure 3.37 shows a histogram of our β_{eff} values for an Infrasil substrate with a Bragg coating on it. The inferred value and error for β_{eff}

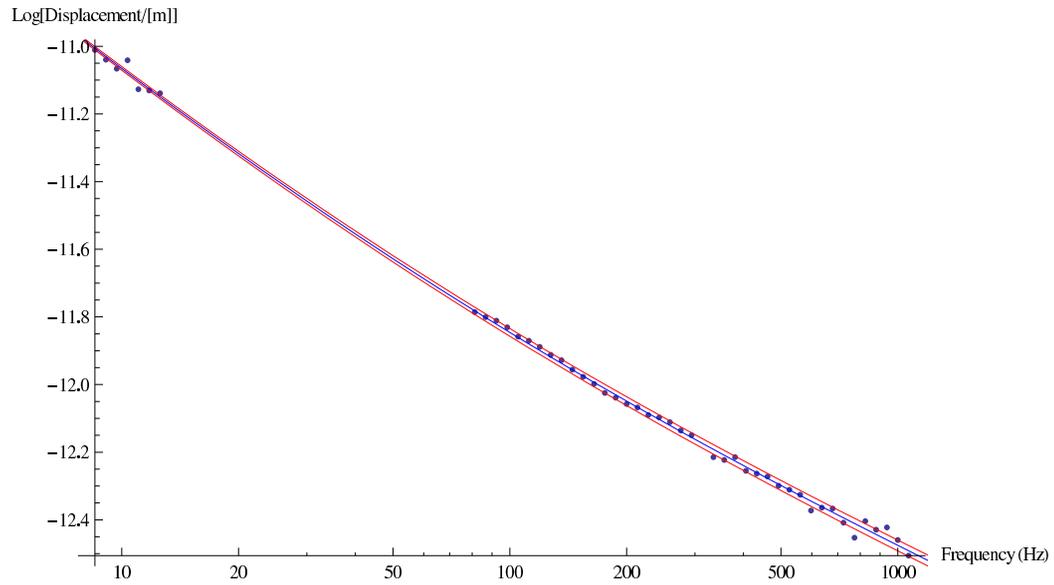


Figure 3.33: Sample Infrasil substrate with QWL + gold coating. Blue line is fit, red lines are statistical error bars derived from multiple runs

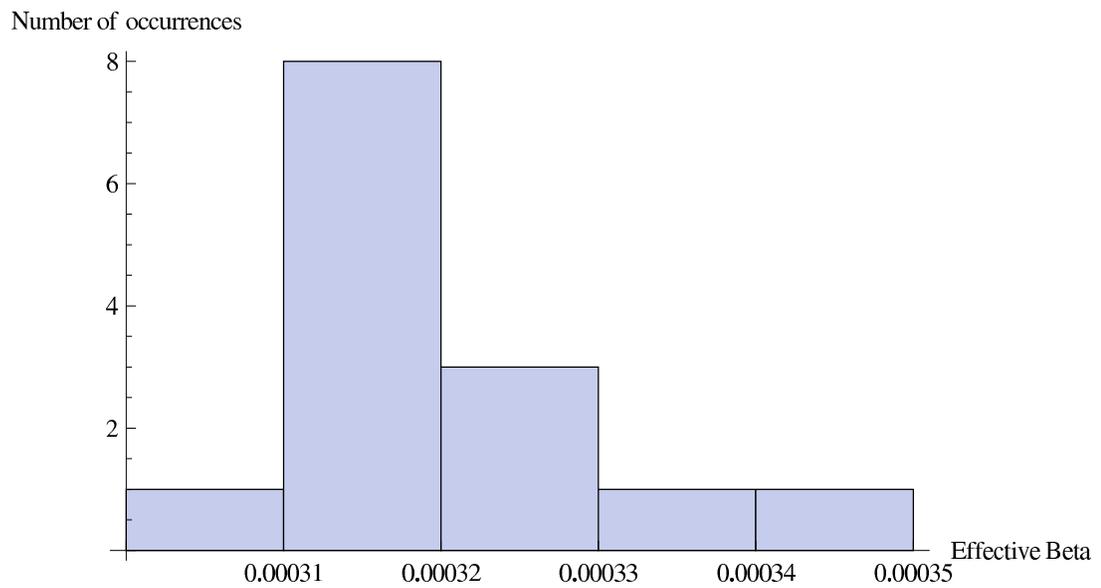


Figure 3.34: Histogram of β_{eff} values for Infrasil substrate QWL coating

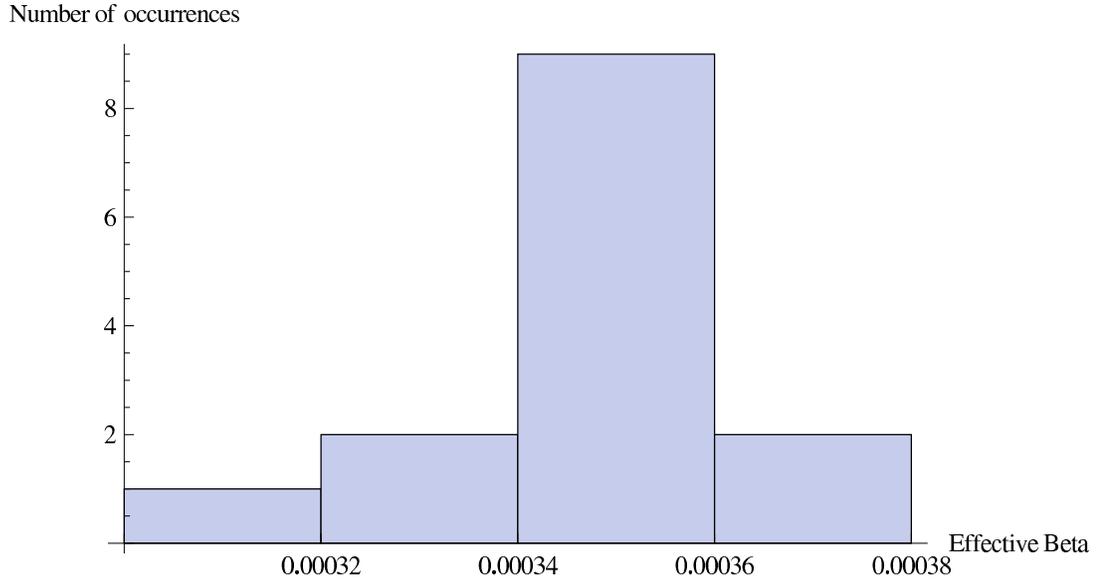


Figure 3.35: Histogram of β_{eff} values for Infrasil substrate QWL coating after re-installation

is $(3.1 \pm 0.12) \times 10^{-4}$, although, as mentioned, we will assign a larger 0.2×10^{-4} uncertainty to this number due to our belief that these numbers suffer from the same systematic uncertainty as the previous gold-coated measurements.

Figure 3.38 shows sample data from an infrasil substrate with a QWL coating with a thin $0.128 \mu\text{m}$ chromium layer. The fit for this particular data set finds $P_0 = .0156$ and $\beta = .0003103$. Figure 3.39 shows a histogram of our β_{eff} values for an Infrasil substrate with a QWL coating that has the extra-thin Cr layer. The inferred value and error for β_{eff} is $(3.2 \pm 0.14) \times 10^{-4}$, although as mentioned we will assign a larger 0.2×10^{-4} uncertainty to this number due to our belief that these numbers suffer from the same systematic uncertainty as the previous gold coated measurements.

3.5.4 Error Analysis

To get a handle on our error, we first tried the time-tested “chi-by-eye” method of varying the parameter of interest (β_{eff}) upward and downward until the fit looked clearly wrong. Later, we switched to a more rigorous method; we took multiple data sets (4 each for the combined numbers, which presented very clean signals, and 14 each for the gold-coated TE numbers) and performed fits to each, treating each run

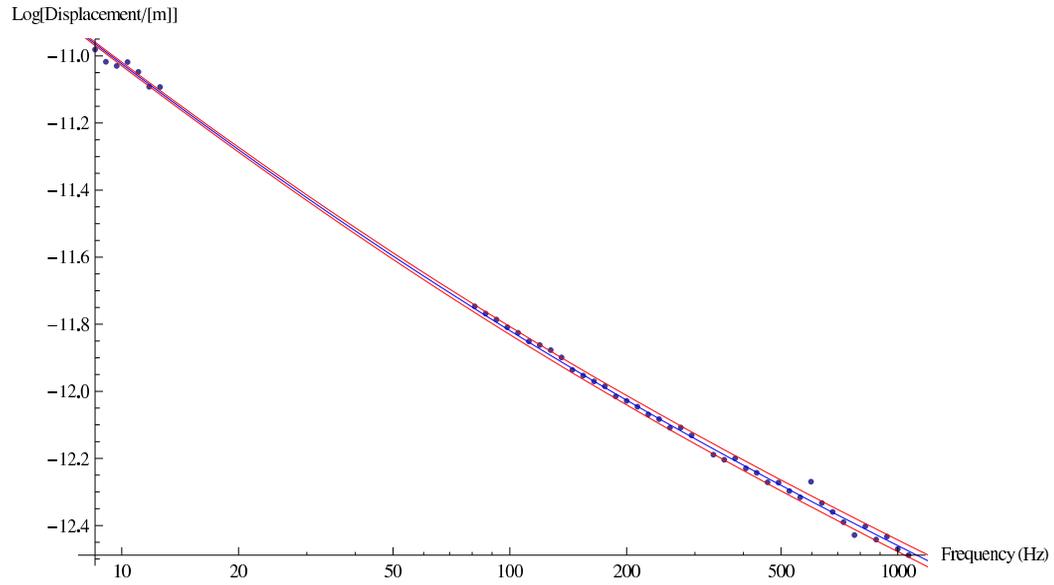


Figure 3.36: Sample Infrasil substrate with Bragg + gold coating. Blue line is fit, red lines are statistical error bars derived from multiple runs.

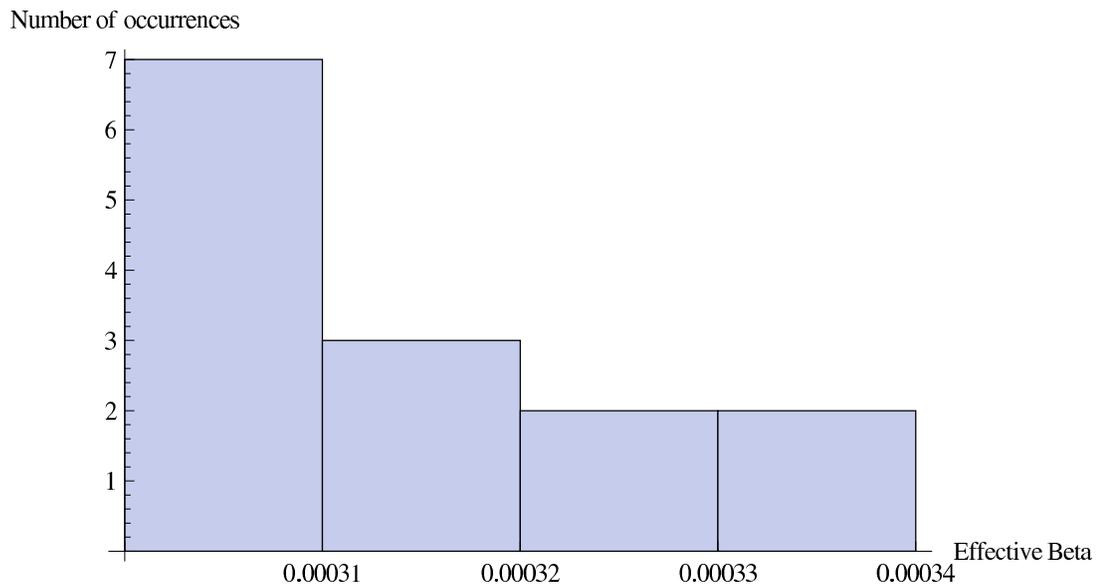


Figure 3.37: Histogram of β_{eff} values for Infrasil substrate Bragg coating

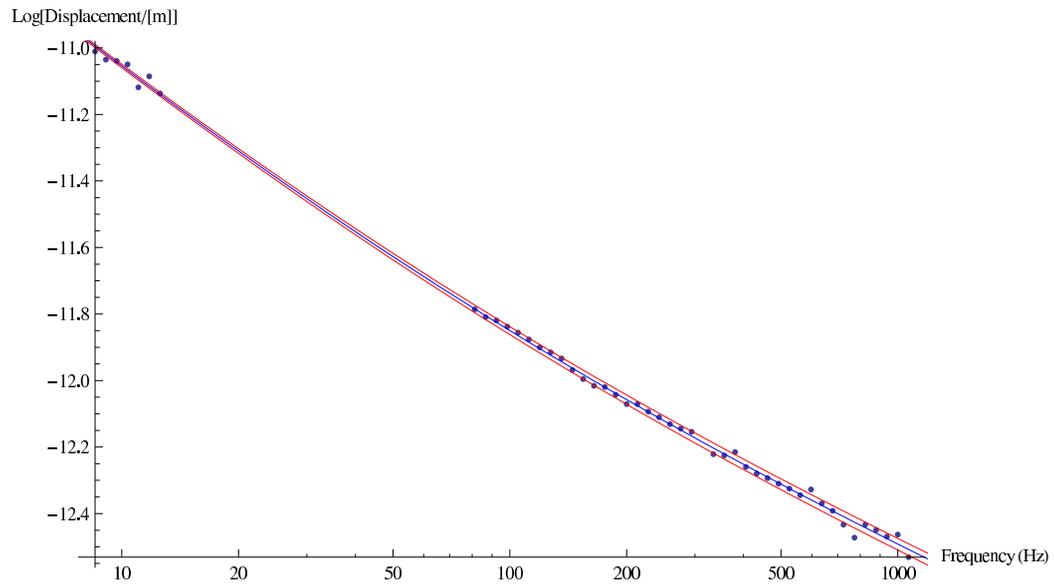


Figure 3.38: Sample Infrasil substrate with QWL + gold with thin Cr layer coating. Blue line is fit, red lines are statistical error bars derived from multiple runs.

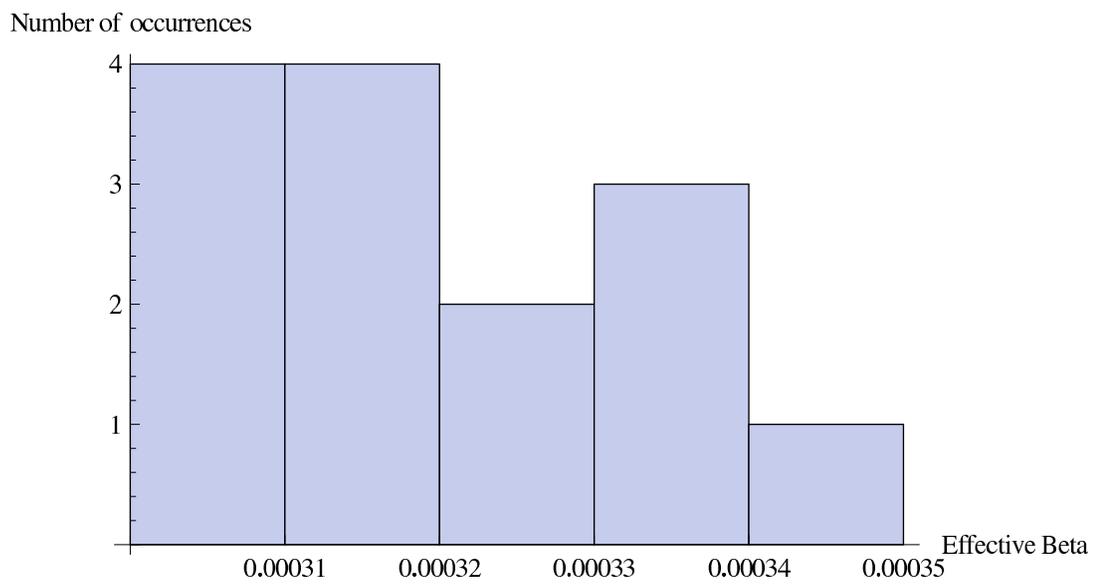


Figure 3.39: Histogram of β_{eff} values for Infrasil substrate with QWL coating with thin underlying Cr layer

as an independent test. Then we took all the recorded parameters from the batch of runs and calculated the mean and standard deviation of the β_{eff} values. This gave us error bars a factor of 2 or so lower than “chi-by-eye”, as one might expect. The mean and standard deviation from these sets of tests and fits are the values and errors that we report for the combined TE+TR measurements. We report something slightly larger for the gold-coated TE-alone results, for reasons described below.

As far as consistency, the combined TE+TR measurements with their large SNR were extremely consistent. The four traces we used for each result (QWL and Bragg) were from four different mirrors. For the QWL coating, two of the substrates were fused silica and two were Infrasil. Obviously, the fits took into account the different substrate parameters. The numbers thus derived for β_{eff} were all within 4% of the mean. The four substrates used for the Bragg coatings were all fused silica, and yielded similar consistency.

We had planned on assuming this sort of consistency for the TE measurements (indeed we had only one fused silica and one infrasil sample each of QWL, Bragg, and QWL-thin-Cr gold-coated), but discovered at a first pass that our errors between immediate subsequent measurements of the same sample were slightly larger, on the order of 10–15%, and on reinstalling one mirror to investigate this, we observed on the order of 25% change when removing, reinstalling, and re-measuring *one sample*.

Clearly in the gold-coated samples we had a) larger error in repeated runs, and b) an even larger random systematic error. We hypothesized that this had to do with the larger noise levels at low frequency around 20–60 Hz effectively giving an inaccurate power calibration to the fit, and spurious coupling starting to become visible above approximately 1 kHz. These spurious couplings can be seen starting to sneak in at the high-frequency ends of Figures 3.33, 3.36, and 3.38. If these couplings were, for example, mechanical coupling (where the expansion of the mirror caused the mount or table to physically shake), we might expect them to have different amplitudes when we removed the mirror/mount and reinstalled it. And their magnitudes, being far out at the end of a linear (in log space) fit, will tend to contribute heavily to the inferred value of β_{eff} .

In order to decrease this error, we did two things. First, we removed low-frequency data from 20–70 Hz from the fits. Second, we cut out as much high-frequency data as we thought we could get away with while still seeing coating effects. The ranges and data sets we came up with can be seen in the figures above. Once this was done, the spread within consecutive runs decreased dramatically, to approximately 4%, and the difference between the same Infrasil sample removed and reinstalled for that second batch of runs was 8%.

We have not had time to take more repeat data to better characterize this systematic error; in our later data analysis we will assume all gold measurements have the larger 8% systematic error bars. These errors need to be studied. We suspect that these systematic errors can be brought down to the level of the consecutive-run error with some work. This will be discussed in the future work section.

3.6 Poisson Factor Correction

All of the above analysis was completed without taking into account the Poisson factor correction. We will work out here the effect this correction has on our measurement.

The data in the points we fit to will not change; the motion we observed is the motion we observed. What will change is the equation we fit to. Originally we fit to equation 3.29, which we will reproduce here:

$$\delta l(\omega) = \frac{\alpha P_0}{\rho C_P \pi r_0^2 \omega} e^{i90^\circ} + \frac{\lambda}{4\pi} \beta_{eff} \frac{P_0 \sqrt{a}}{\kappa A} \frac{1}{\sqrt{\omega}} e^{-i45^\circ}$$

Remembering where this equation comes from, most of it is a determination of the temperature profile; none of that derivation will change since the Poisson ratio is a structural-elastic correction. What will change is the projected motion due to that expansion. The places where we predicted a structural expansion were the places where we inserted α coefficients. The first term on the right-hand side is the obvious one. We will now be fitting to a new equation

$$\delta l(\omega) = \frac{\alpha(1 + \sigma)P_0}{\rho C_P \pi r_0^2 \omega} e^{i90^\circ} + \frac{\lambda}{4\pi} \beta_{eff} \frac{P_0 \sqrt{a}}{\kappa A} \frac{1}{\sqrt{\omega}} e^{-i45^\circ}$$

where σ is the Poisson ratio of the substrate. How will this change the parameters we fit? Since the δl we are dealing with is measured experimentally, our estimated length change will not be different. What will be different is that in order for this new calculation to match the old length measurements, we will estimate a new power

$$P'_0 = \frac{P_0}{1 + \sigma}.$$

An increase in the effective coefficient of expansion by a factor of $(1 + \sigma)$ will lead us to infer that the power we observed was smaller by that same factor.

Given that we infer a smaller power by a factor of $(1 + \sigma)$, how does that change our estimation of β_{eff} ? We will again make the same argument that the second term on the right-hand side must remain the same since our data has not changed. Now that our estimated power is smaller though, we need to estimate a larger β_{eff} by the same factor, $(1 + \sigma)$, to ensure our predicted/estimated line matches the data points, which have not moved.

Therefore we conclude (and verified on a few data sets though not all) that the values estimated above for the measured β_{eff} are too small by a factor of $(1 + \sigma)$. Below are the measured values after we apply the correction factor, with the value $\sigma = 0.17$ for both fused silica and Infrasil.

$$\beta_{eff}^{QWL} = (1.39 \pm 0.03) \times 10^{-4} K^{-1} \quad (3.41)$$

$$\beta_{eff}^{Bragg} = (1.26 \pm 0.04) \times 10^{-4} K^{-1} \quad (3.42)$$

$$\beta_{eff}^{QWL-TE} = (3.9 \pm 0.2) \times 10^{-4} K^{-1} \quad (3.43)$$

$$\beta_{eff}^{Bragg-TE} = (3.6 \pm 0.2) \times 10^{-4} K^{-1} \quad (3.44)$$

$$\beta_{eff}^{QWL-TE-ThinCr} = (3.7 \pm 0.2) \times 10^{-4} K^{-1} \quad (3.45)$$

The other Poisson ratio correction we will need to make is in our interpretation of the coefficients of thermal expansion that went into these effective β s. We assumed there was a certain value α for each of the layers that related the change in length to temperature. Really, since the coating layers are subject to the same constraint as the substrate, i.e., no lateral expansion, the number there should be $\alpha(1 + \sigma_C)$, where the Poisson ratios are those for the coating layers. Since these are also not particularly well known, we will at the end of this thesis quote values for $\alpha_L(1 + \sigma_L)$ and $\alpha_H(1 + \sigma_H)$, since these are the parameters that will be of interest in calculating the thermal noise effects in interferometers anyway.

Chapter 4

Parameter Extraction

Now that we have four measurements (TE/TR combined for QW and Bragg coatings, as well as TE alone for QW and Bragg coatings), we can extract the four material parameters we are interested in.

4.1 The Measurement Matrix

In Section 2.3 we calculated the thermo-refractive response coefficients of the coatings to be

$$\overline{\beta}_{QWL}^{TR} = 6.4\alpha_L + 2.9\alpha_H + 4.5\beta_L + 1.4\beta_H \quad (4.1)$$

$$\overline{\beta}_{Bragg}^{TR} = 9.6\alpha_L + 2.18\alpha_H + 6.9\beta_L + 1.18\beta_H. \quad (4.2)$$

This is not the full story, however. In addition to the altered reflection coefficient, the surface of the mirror will be moving itself due to the expansion of the layers (including the expansion of the chromium underlayer). Fortunately, we need to calculate this effect anyway since it is exactly the effect we measure in the gold-coated samples.

For a multilayer mirror, we can define an effective coefficient of thermal expansion for the entire mirror that is equivalent to the appropriate thickness-weighted sum of coefficients of the individual layers

$$\bar{\alpha} = \frac{1}{L} \sum_{layers} \alpha_i l_i \quad (4.3)$$

where L is the total thickness of the coating, and l_i is the thickness of the i^{th} layer. Then, when we have an expansion of the coating, it will appear to contribute an added phase to the laser beam of

$$\Delta\phi^{TE} = -\frac{4\pi}{\lambda} \bar{\alpha} L \Delta T.$$

Thus, we define an effective β due to the expansion of the coating by

$$\bar{\beta}^{TE} = -\frac{4\pi L}{\lambda} \bar{\alpha}.$$

Since we know the layer thicknesses, we can calculate numerically the coefficients of α_L and α_H for this expression. Unfortunately, we must also include a term for the expansion of the chromium layer underneath the coating, an effect which is too large to neglect.

$$\bar{\beta}_{QWL}^{TE} = -28\alpha_L - 17\alpha_H - 10.6\alpha_{Cr} \quad (4.4)$$

$$\bar{\beta}_{Bragg}^{TE} = -40\alpha_L - 8.4\alpha_H - 10.6\alpha_{Cr} \quad (4.5)$$

With this added variable, the coefficient of thermal expansion for thin-film chromium, we need one more measurement to pin down all the unknowns. Fortunately, we had in the lab a set of samples from an earlier run where the chromium layer was designed incorrectly and was too thin to stop the penetration of the CO₂ radiation. We sent these to be gold coated with the rest of the ‘‘TE only’’ samples, and with the thinner layer of chromium they provide us one more data point:

$$\bar{\beta}_{ThinCr}^{TE} = -28\alpha_L - 17\alpha_H - 1.5\alpha_{Cr}. \quad (4.6)$$

We are making the assumption here that the coefficient of thermal expansion is

the same for 128 nm of chromium as it is for 900 nm of chromium, but we believe deviations from that assumption should not significantly affect our outcome.

Finally, noting that the response of the bare coatings (without gold) is simply the sum of the TR effects from equations 4.1 and 4.2, and the TE effects from equations 4.4 and 4.5, we can write out equations in matrix form relating the measurements we make to the parameters we are interested in.

$$\begin{bmatrix} \overline{\beta}_{QWL}^{Combined} \\ \overline{\beta}_{Bragg}^{Combined} \\ \overline{\beta}_{QWL}^{TE} \\ \overline{\beta}_{Bragg}^{TE} \\ \overline{\beta}_{ThinCr}^{TE} \end{bmatrix} = \begin{bmatrix} -21.6 & -14.1 & 4.5 & 1.4 & -10.6 \\ -30.4 & -6.22 & 6.9 & 1.18 & -10.6 \\ -28 & -17 & 0 & 0 & -10.6 \\ -40 & -8.4 & 0 & 0 & -10.6 \\ -28 & -17 & 0 & 0 & -1.5 \end{bmatrix} \begin{bmatrix} \alpha_L \\ \alpha_H \\ \beta_L \\ \beta_H \\ \alpha_{Cr} \end{bmatrix}. \quad (4.7)$$

The vectors in the matrix are linearly independent, and thus it is nonsingular and invertible. Using Mathematica to invert this matrix, we find our parameter extraction matrix

$$\begin{bmatrix} \alpha_L \\ \alpha_H \\ \beta_L \\ \beta_H \\ \alpha_{Cr} \end{bmatrix} = \begin{bmatrix} 0 & 0 & .041 & -.038 & -.023 \\ 0 & 0 & -.059 & .063 & -.031 \\ -.27 & .32 & .21 & -.26 & .028 \\ 1.6 & -1.0 & -1.5 & .90 & .08 \\ 0 & 0 & -.11 & 0 & .11 \end{bmatrix} \begin{bmatrix} \overline{\beta}_{QWL}^{Combined} \\ \overline{\beta}_{Bragg}^{Combined} \\ \overline{\beta}_{QWL}^{TE} \\ \overline{\beta}_{Bragg}^{TE} \\ \overline{\beta}_{ThinCr}^{TE} \end{bmatrix}. \quad (4.8)$$

4.2 Error Analysis

4.2.1 Statistical Error

We will assume there is no error in the matrices in equations 4.7 and 4.8. This is not strictly right — if we recall where these coefficients came from, they will depend on two quantities. First, they depend explicitly on the indices of refraction of the materials in the layers. This is well known, however, and will be a much smaller

source of error than our experimental uncertainty. The second factor that goes into these coefficients is the coating layer thicknesses — the thicknesses factor implicitly in our assumption that we were expanding around a quarter-wave thickness, or $1/8$, or $3/8$ wave thicknesses. On private communication from our coating manufacturer, I have been informed that I can expect $\sim 2\%$ variation in layer optical path lengths, from which I expect the same order variation or less in matrix coefficients (being at a reflectance plateau, leading order corrections would be expected to go as the square of the errors in optical path lengths). Again, this will be smaller than the experimental error we observe, and so will be neglected here.

Given this assumption, we can translate measurement error directly into parameter error by the formula

$$\sigma_{parameter-i} = \sqrt{m_{ij}^2 \sigma_{measurement-j}^2} \quad (4.9)$$

where the m_{ij} are the matrix elements in the parameter extraction matrix from equation 4.8, and there is an implied sum over j under the square root.

4.2.2 Known Unknowns

There are two effects that we have mentioned before that are worth bringing up again in this error analysis section, before we examine the results of our parameter estimation.

The first is the factor of 2 in our COMSOL simulation. As mentioned, we believe this to be an artifact that is safe to neglect. If it turns out to be real, that means that elastic effects are causing the surface of the mirror to actually move twice as much as we expect it to move for a given input power and temperature profile. We would then expect this effect to be confirmed by a complete elastic theory analysis. This would clearly affect the way we fit to our data sets; our given substrate motion would be fit with approximately half the power, and with a halved power estimate, we would estimate our β_{eff} to be twice as large to fit the data points taken.

The second effect that is worth mentioning again is the random systematic un-

certainty in the low-signal gold-coated measurements. We have not yet characterized this uncertainty well, and our estimate of their size is based on a very small set of measurements. With the data we have, we believe the error bars we present are reasonable, but they could be larger than we estimate.

4.3 Results

Using the techniques described in the previous sections, along with the results of our measurements, we get the following measured values for the parameters of interest:

$$\alpha_L(1 + \sigma_L) = (6.4 \pm 1.4) \times 10^{-6} K^{-1} \quad (4.10)$$

$$\alpha_H(1 + \sigma_H) = (1.1 \pm 0.2) \times 10^{-5} K^{-1} \quad (4.11)$$

$$\beta_L = (1.9 \pm 8.0) \times 10^{-6} K^{-1} \quad (4.12)$$

$$\beta_H = (1.2 \pm 0.4) \times 10^{-4} K^{-1} \quad (4.13)$$

which use the derived value for our thin film of chromium

$$\alpha_{Cr}(1 + \sigma_{Cr}) = (1.3 \pm 3.6) \times 10^{-6} K^{-1}. \quad (4.14)$$

These error bars are rather large (compared to our $\sim 5\%$ errors on individual β_{eff} results) due to the fact that we are subtracting numbers close to the same magnitude. Since the errors add in quadrature, they stay at the same absolute order of magnitude while the differencing process can mean our results for actual parameter values are down an order of magnitude or so from the measured numbers. This means our fractional errors go up by an order of magnitude. In the next chapter, on future directions, we will discuss possible ways to decrease these errors. They will include decreasing the error on the original measurements as well as increasing the size of the

difference of measurements by making coating variations more extreme.

Chapter 5

Conclusions and Future Directions

5.1 Conclusions

We report here, to the best of our knowledge, the first measurement of the thermo-optic parameters α and β (dn/dT) for individual layers of thin-film materials in a dielectric stack. Our error bars are larger than we would like and we have systematic effects to study, but we have shown the technique to be feasible and believe this to be a solid starting point for further research.

5.1.1 Comparison with Published Values

We compare our measured values with published values from tables 2.1 to 2.4. For the α values, we will divide by $(1 + \sigma_c)$, using values for Poisson ratios listed in GWINC. We have not studied the literature for these Poisson ratio values, so there is potentially some error induced in the use of these values, but we believe there would be more without using them.

For the coefficient of expansion of thin-film fused silica, our value of $(5.5 \pm 1.2) \times 10^{-6}$ compares reasonably with Centigoru's value. Ours is higher than theirs, but within a few standard deviations and, as noted, they used different deposition technology and a much thicker layer. We are a factor of 10 larger than Braginsky's number used in GWINC, but again we do not have a reference for this and it appears to be a bulk value.

For the coefficient of thermal expansion of thin-film tantala, our value of $(8.9 \pm 1.8) \times 10^{-6}$ is approximately a factor of 2 larger than the measurements listed, but does seem to contradict the negative sign observed by Inci.

Our dn/dT measurement for silica has unfortunately large error bars, which do include the value suggested in the GWINC documentation attributed to Braginsky.

Our dn/dT measurement for tantala agrees with the value quoted by Inci (surprisingly, since the situations are so different and we get radically different α values). We are larger than Andri Gretarsson's measurement by a factor of 2, but his value is within two of our standard deviations, and we unfortunately don't know at the time of writing what values he assumed for the other material parameters.

Finally, we compare our results with the measurement published by Farsi [10]. This was an experiment where the power in a Fabry-Perot cavity was modulated, and the resulting change in power absorbed in the mirrors was used to measure a version of β_{eff} and $\bar{\alpha}$ (presumably weighted by volume) for their coatings. When we compare our $\bar{\alpha}$ to theirs, we get $(6.9 \pm 2) \times 10^{-6}$ versus their measured $(2.3 \pm 0.3) \times 10^{-6}$. Our measurement here is significantly higher. We can also use our numbers to calculate their version of β_{eff} , which is slightly different from ours and includes only TR effects. The result is our $(1.7 \pm 0.6) \times 10^{-5}$ versus their $(4.4 \pm 0.6) \times 10^{-6}$. Our values are again larger, and more than two standard deviations from theirs.

We observe a trend where our numbers are consistently large compared to the rest of the literature. We can not say at this point whether or not this is due to systematic errors in our experiment, or the noted differences in our coatings (likely some combination of both).

5.1.2 Possible Implications for Future Gravity Wave Detectors

The material parameters we measure are substantially different from values commonly used in Advanced LIGO noise predictions. A summary of the changes between our numbers and the numbers currently used in LIGO's Gravitational Wave Interferom-

Table 5.1: Comparison of GWINC v3 parameters and our measured parameters.

Parameter	GWINC value	Our Value	New/Old ratio
α_L	5.1×10^{-7}	5.5×10^{-6}	10.8
α_H	3.6×10^{-6}	8.9×10^{-6}	2.5
β_L	8×10^{-6}	1.9×10^{-6}	0.24
β_H	1.4×10^{-5}	1.2×10^{-4}	8.6

eter Noise Calculator (GWINC) version 3 (the current version) is listed in Table 5.1. Again, we have divided our numbers by $(1 + \sigma_c)$, using the values for coating Poisson ratios listed in the program, in order to make an “apples to apples” comparison.

The major changes between our measurements and the widely used values are factors of ~ 10 in α_L and β_H , though it is important to note that α_H also went up significantly. These changes, if also measured in the coatings LIGO will use for its advanced interferometers, could have significant implications for the noise performance in the interferometers’ most sensitive frequency range.

We repeat our standard AdLIGO curve from Chapter 1 in Figure 5.1. In Figure 5.2, we show a GWINC calculation of the AdLIGO noise curve based on our parameters. Figures 5.3 and 5.4 show our estimates of thermo-optic noise with the best-case parameters and worst-case parameters within our one-standard-deviation errors.

The main cause of the raised thermo-optic noise levels is the increased α s with a relatively small increase in the β s. The earlier cancellation predictions had parameter values that fortuitously put TE and TR effects on the same footing, so coherent addition had them largely cancelling. With our numbers, the TE effect is much larger due to the increased α values, the effect of which is leveraged by the fact that AdLIGO coatings will have approximately twice as many layers as our samples, and all contribute to the expansion. We do have larger TR effects predicted, but the coating reflection coefficient effects decrease with variations deeper in the coating, so the higher cancelling effects scale with the parameters, but are not leveraged by the number of layers.

It is not clear that the numbers we have measured here imply such large thermo-optic noise levels for Advanced LIGO and future detectors. First, we used a different

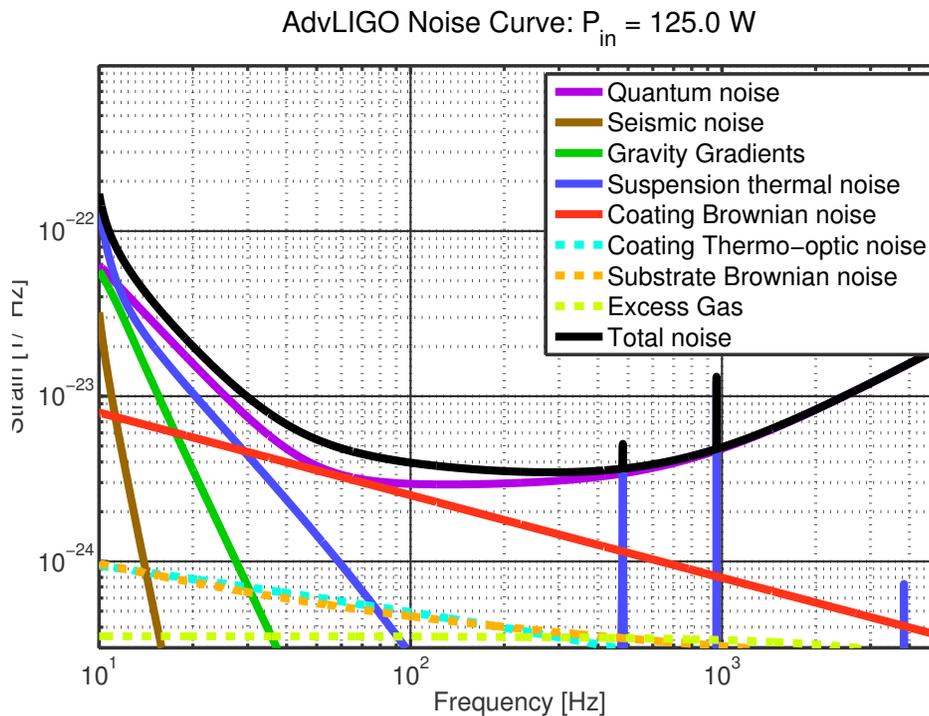


Figure 5.1: Baseline Advanced LIGO noise curve

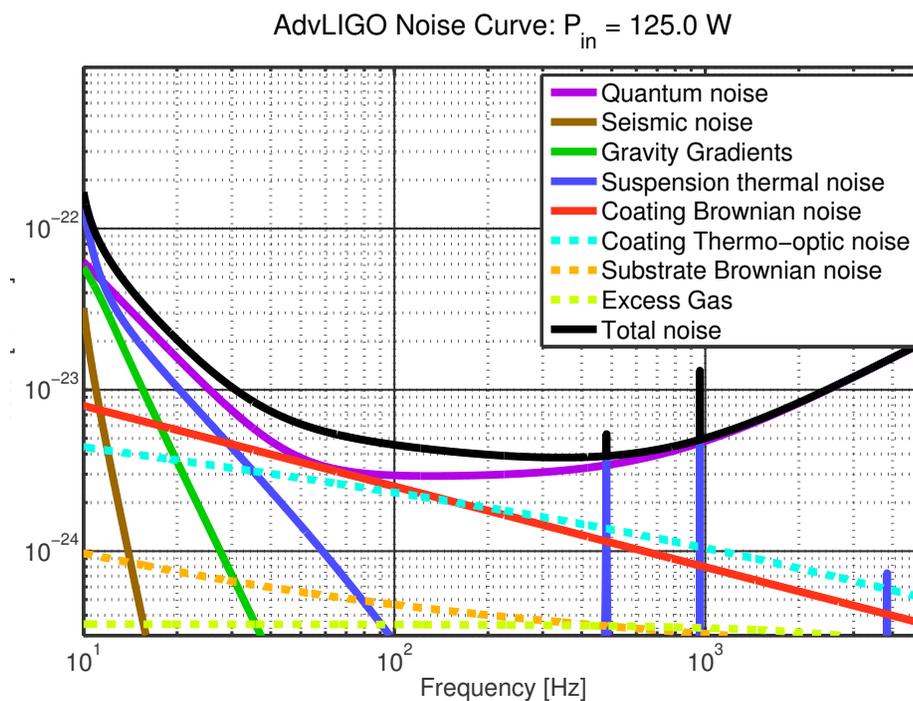


Figure 5.2: Predicted AdLIGO noise using our measured material parameters. Thermo-optic noise may be substantially different with different coating processes, manufacturers

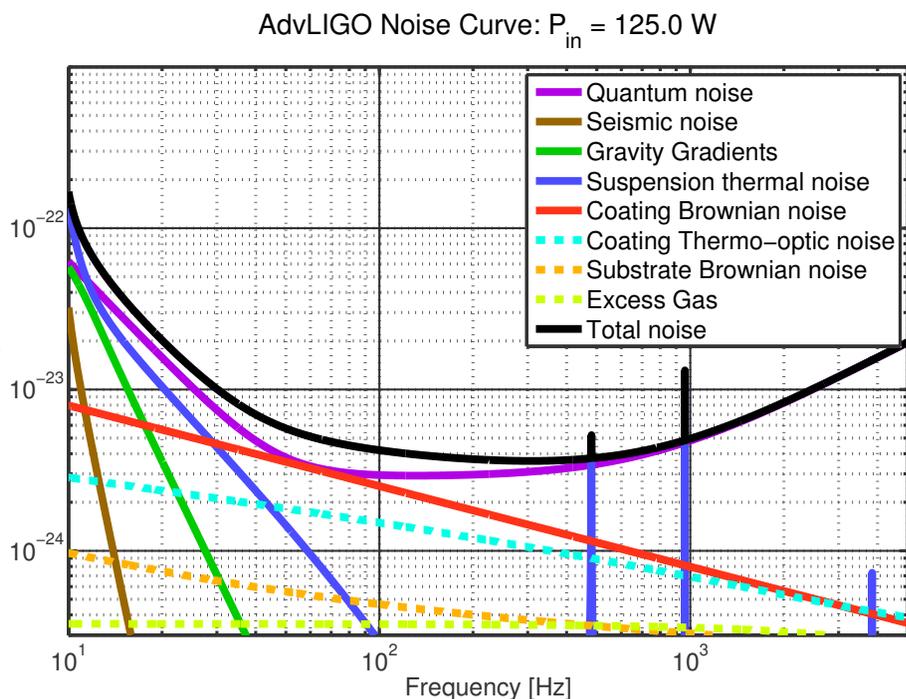


Figure 5.3: One-sigma best-case AdLIGO noise curve using our measured material parameters. Thermo-optic noise may be substantially different with different coating processes, manufacturers

manufacturer and a different deposition technology (electron beam sputtering rather than ion beam sputtering), so our coatings were almost certainly structurally different from those that will be used in AdLIGO. Second, as mentioned in our error analysis section, there are systematic effects in our small-signal TE measurements that we have not studied well; since the TE effects are the ones in question, and also since all of our measurements depend on them (we can not extract the β numbers without first getting the α s from the TE experiment), there could be change in the future as this research continues.

What we do hope to assert in this thesis is that this issue is important. It is quite conceivable for coating parameters to differ from those that cancel nicely in Advanced LIGO or future detectors' coatings, so the parameters need to be studied and measured closely. The technique described here is starting to yield good results, and in the near future can and should be pushed to better precision and understanding.

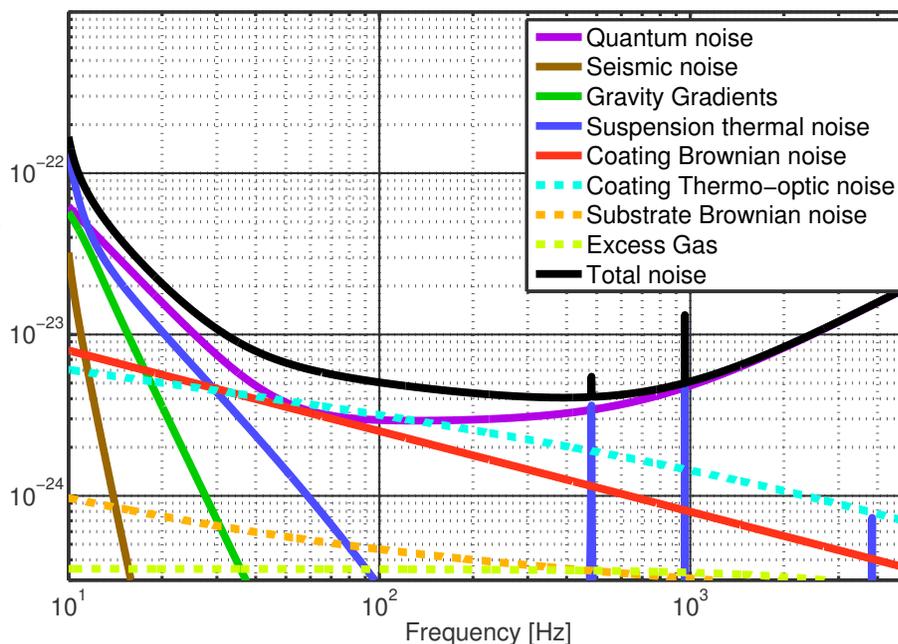
AdvLIGO Noise Curve: $P_{in} = 125.0$ W

Figure 5.4: One-sigma worst-case AdvLIGO noise curve using our measured material parameters. Thermo-optic noise may be substantially different with different coating processes, manufacturers

5.2 Future Work

5.2.1 Decrease Noise in Gold-Coated Measurements

As we mentioned in the section on gold-coated error analysis, we believe we have systematic errors that are not well characterized. These errors need to be better characterized and reduced.

As mentioned, we suspect these errors are induced by noise at low frequencies and spurious couplings (mechanical and/or electrical) at high frequencies. The noise at low frequencies can be averaged down by increasing the time constant on the lock-in amplifier, which we effectively tried to do originally by averaging many data sets. However, even with the averaging we observed what looked to be increased power in that region at low frequencies. A better strategy would be to increase the signal, either by increasing the heating laser power or by taking data at lower frequencies where the signal should continue to rise as $1/f$ until we hit low-frequency “Cerdonio-

type” effects at tens to hundreds of mHz (assuming of course we don’t run into more low-frequency seismic or flicker noise).

The high-frequency coupling is a more challenging problem. If they are truly vibrational couplings, increasing pump power won’t help because that will increase the mechanical shaking along with our signal. Possible solutions in this case would include determining exactly what’s moving (e.g., the sample mount vs. table resonance vs. other IFO mounts responding to scattered CO₂ light) and damping it or otherwise redesigning mechanical portions of the system. If the couplings are electrical (ground loops or radiative coupling of the 40 MHz AOM driver signal) then they can be tracked down and eliminated with more careful grounding and/or filtering. It might also be possible, if the couplings aren’t too pervasive, to simply characterize them well enough to remove data only in the vicinity of known couplings, rather than cutting out large swathes of high-frequency data.

5.2.2 Coating Design for Increased Signal

As we discussed in our error analysis, the errors on our coating parameters were larger than the errors on our measurements due to the fact that we were looking for the difference in approximately equal numbers. In hindsight, anything we could do to increase the difference between these numbers would decrease the relative error in our material parameter measurements. An easy example of this would be to repeat the experiment with 1/8–3/8 and 3/8–1/8 mirrors rather than the QWL design (though using all three would certainly provide extra information and consistency checks).

One could also envision relatively larger differences in ratios — the Bragg condition will give high reflectivity for any configuration in which the layer doublets add up to $\lambda/2$. However, one would want to take care to consider effects that kick in on the length scales over which layers might diffuse into each other (see Section 5.2.5).

5.2.3 Boundary Condition and Poisson Ratio Effects

It would be useful to work with theorists to write up comprehensible descriptions of why the correct Poisson ratio factor for expansion constrained to zero in-plane expansion is $\alpha(1 + \sigma)$.

Another next step is to work on understanding what our COMSOL simulations are telling us. First and foremost it is critical to sort out the embarrassing factor of 2 between the reported z-displacement of the surface and α times the integrated temperature profile. We do not believe this doubling of expected surface motion is a real effect, since our experimental results with the combined TE and TR effects agree with our naive prediction to within the $\sim 20\%$ error that we estimate exists on our overall power levels. We believe it to be extremely unlikely that we are missing a factor of 2. However, if we were, this would definitely change the power we get from fitting our low-frequency data, and thus also change our estimates of the coating response at high frequency.

5.2.4 Advanced LIGO Coating Measurements

As noted in our Theory and Background chapter, these parameters are quite likely dependent on deposition technology, and quite possibly dependent on coating stress environment. The numbers we report here are useful to demonstrate this as a valid technique, but the numbers we measure are for electron-beam-gun-deposited coatings of 99.9% pure silica and tantala. Advanced LIGO will use coatings of silica and titania-doped tantala deposited by ion beam deposition, with approximately twice as many layers. These numbers will need to be remeasured for the deposition technology and coating materials used by the project.

The thicker coatings for Advanced LIGO should be easier to measure. We designed our coatings so that the TE effect would be on par with the TR effect, thinking that would make the cancellation as easy to observe as possible. It turned out our main limitation was by far the noise in the measurement of the α s due to the small signal involved once the entire coating is hidden under a gold layer. Increasing the thickness

of coating that participates in our α measurements will increase that signal by a factor of 2 or so with respect to the spurious couplings we're seeing and decrease that noise source, and should not significantly decrease the observability of the TR effects without the gold coating.

5.2.5 Further Coating Studies

As mentioned earlier, there are many hypotheses as to why thin film materials properties are so different from their bulk values. Groups in LIGO are studying the actual structure of the material through computational methods, low frequency mechanical loss measurements, and, in the near future, through X-ray diffraction.

The technique outlined here is an additional way of studying these materials, and has the benefit of providing measurements that are as close to real world “in-situ” measurements as one can get. This technique samples coating parameters as they are in coatings — high stress, and in contact with other layers of similar makeup.

One possible source of the difference between coating material parameters is the high stresses these coatings experience. Under high stress, any temperature dependence of the Young's modulus could cause unanticipated motion of the coating surface.

Another possibility is diffusion from adjacent layers — it has been shown via SEM investigation of coatings [12] that material types don't change discontinuously across layer boundaries, but transition gradually over a region of tens of nanometers. Layer thicknesses for our $\lambda/4$ coatings are 129 nm of tantala and 183 nm for silica. In our Bragg coatings, $\lambda/8$ of tantala was 64.6 nm. It shouldn't take much more of a push in this direction before tens of nm of diffusion start to have a significant impact on the thermal and mechanical properties of the layers. We envision a parametric study with various Bragg coating designs in which the thin layers approach these diffusion lengths to confirm or exclude models of layer diffusion and the resulting impacts on materials parameters.

Now that this technique has been shown to be feasible, there are many avenues of study that are open for us to learn more about the physical and optical properties of

these exciting and critical materials.

Bibliography

- [1] H. C. Lin A. K. Chu and W. H. Cheng. Temperature dependence of refractive index of ta2o5 dielectric films. *Journal of Electronic Materials*, 26, 1997.
- [2] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A*, 264(1):1 – 10, 1999.
- [3] V. B. Braginsky and A. A. Samoilenko. Measurements of the optical mirror coating properties. *Physics Letters A*, 315(34):175 – 177, 2003.
- [4] V. B. Braginsky and S. P. Vyatchanin. Thermodynamical fluctuations in optical mirror coatings. *Physics Letters A*, 312:244 – 255, 2003.
- [5] K. Chuang C. Tien, C. Lee and C. Jaing. Simultaneous determination of the thermal expansion coefficient and the elastic modulus of ta2o5 thin film using phase shifting interferometry. *Journal of Modern Optics*, 47(10):1681–1691, 2000.
- [6] H. B. Callen and T. A. Welton. Irreversibility and generalized noise. *Phys. Rev.*, 83:34–40, Jul 1951.
- [7] S. M. Carroll. *An Introduction to General Relativity Spacetime and Geometry*. Addison Wesley, 2004.
- [8] M. M. Fejer A. Gretarsson G. Harry J. Hough N. Nakagawa S. Penn R. Route S. Rowan D. R. M. Crooks, G. Cagnoli and P. H. Sneddon. Experimental measurements of coating mechanical loss factors. *Classical and Quantum Gravity*, 21(5):S1059, 2004.

- [9] S. R. Rao E. D. Black, I. S. Grudinin and K. G. Libbrecht. Enhanced photothermal displacement spectroscopy for thin-film characterization using a fabry-perot resonator. *Journal of Applied Physics*, 95(12):7655–7659, 2004.
- [10] A. Farsi et al. Photothermal and thermo-refractive effects in high reflectivity mirrors at room and cryogenic temperature. *Journal of Applied Physics*, 111(12):043101, 2012.
- [11] E. Cetinorgu et al. Mechanical and thermoelastic characteristics of optical thin films deposited by dual ion beam sputtering. *Applied Optics*, 48:4536 – 4544, Aug 2009.
- [12] G. Cagnoli et al. Coating research and development at lma. *LIGO DCC, GWADW 2012*, pages G080151–00, 2008.
- [13] G. R. Fowles. *Introduction to Modern Optics, second edition*. Dover Publications, Inc., New York, 1975.
- [14] M. L. Gorodetsky. Thermal noises and noise compensation in high-reflection multilayer coating. *Physics Letters A*, 372(46):6813 – 6822, 2008.
- [15] A. Gretarsson. Thermo-optic noise from doped tantala/silica coatings. *LIGO DCC*, pages G080151–00, 2008.
- [16] A. Gurkovsky and S. Vyatchanin. The thermal noise in multilayer coating. *Physics Letters A*, 374:3267 – 3274, 2010.
- [17] M. N. Inci. Simultaneous measurements of the thermal optical and linear thermal expansion coefficients of a thin film etalon from the reflection spectra of a superluminescent diode. *Journal of Physics D*, 37:3151 – 3154, 2004.
- [18] L. D. Landau and E. M. Lifshitz. *Theory of Elasticity, third ed.* Pergamon, Oxford, 1986.
- [19] Y. Levin. Fluctuation-dissipation theorem for thermo-refractive noise. *Physics Letters A*, 372:1941 – 1944, 2008.

- [20] Y. T. Liu and K. S. Thorne. Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses. *Phys. Rev. D*, 62:122002, Nov 2000.
- [21] A. Heidmann M. Cerdonio, L. Conti and M. Pinard. Thermoelastic effects at low temperatures and quantum limits in displacement measurements. *Phys. Rev. D*, 63:082003, Mar 2001.
- [22] M. Cerdonio M. Pinard M. De Rosa, L. Conti and F. Marin. Experimental measurement of the dynamic photothermal effect in fabry-perot cavities for gravitational wave detectors. *Phys. Rev. Lett.*, 89:237402, Nov 2002.
- [23] M. Fejer P. Fritschel G. Harry M. Evans, S. Ballmer and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Phys. Rev. D*, 78:102003, Nov 2008.
- [24] G. Cagnoli D. R. M. Crooks A. Gretarsson G. M. Harry J. Hough S. D. Penn P. H. Sneddon M. M. Fejer, S. Rowan and S. P. Vyatchanin. Thermoelastic dissipation in inhomogeneous media: loss measurements and displacement noise in coated test masses for interferometric gravitational wave detectors. *Phys. Rev. D*, 70:082003, Oct 2004.
- [25] M. L. Gorodetsky V. B. Braginsky and S. P. Vyatchanin. Thermo-refractive noise in gravitational wave antennae. *Physics Letters A*, 271(56):303 – 307, 2000.
- [26] S. F. Chi W. H. Cheng and A.K. Chu. Effect of thermal stresses on temperature dependence of refractive index for ta2o5 dielectric films. *Thin Solid Films*, 347(12):233 – 237, 1999.