

**Active Wavefront Correction in Laser
Interferometric Gravitational Wave Detectors**

by

Ryan Christopher Lawrence

A.B., University of Chicago (1997)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

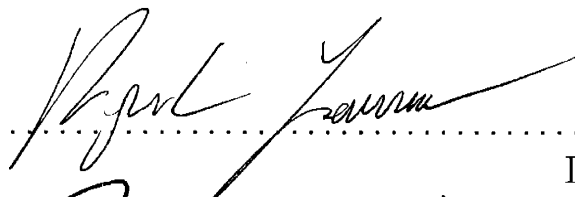
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2003

© Ryan Christopher Lawrence, MMIII. All rights reserved.

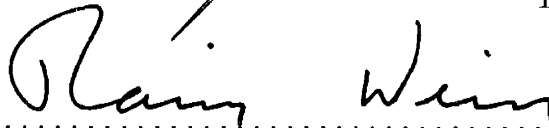
The author hereby grants to MIT permission to reproduce and
distribute publicly paper and electronic copies of this thesis document
in whole or in part.

Author



Department of Physics
January 10, 2003

Certified by




Rainer Weiss
Professor of Physics
Thesis Supervisor

Accepted by



Thomas J. Greytak
Associate Department Head for Education

ARCHIVES



Active Wavefront Correction in Laser Interferometric Gravitational Wave Detectors

by

Ryan Christopher Lawrence

Submitted to the Department of Physics
on January 10, 2003, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

As the first generation of laser interferometric gravitational wave detectors near operation, research and development has begun on increasing the instrument's sensitivity while utilizing existing infrastructure. In the Laser Interferometer Gravitational Wave Observatory (LIGO), significant improvements are being planned for installation in ~ 2007 to increase the sensitivity to test mass displacement, hence sensitivity to gravitational wave strain, by improved suspensions and test mass substrates, active seismic isolation, and higher input laser power. Even with the highest quality optics available today, however, finite absorption of laser power within transmissive optics, coupled with the tremendous amount of optical power circulating in various parts of the interferometer, result in critical wavefront deformations which will cripple the performance of the instrument. Discussed is a method of active wavefront correction via direct thermal actuation on optical elements of the interferometer; or, "thermally adaptive optics". A simple nichrome heating element suspended off the face of an affected optic will, through radiative heating, remove the gross axisymmetric part of the original thermal distortion. A scanning heating laser will then be used to remove any remaining non-axisymmetric wavefront distortion, generated by inhomogeneities in the substrate's absorption, thermal conductivity, etc. This work includes a quantitative analysis of both techniques of thermal compensation, as well as the results of a proof-of-principle experiment which verified the technical feasibility of each technique.

Thesis Supervisor: Rainer Weiss

Title: Professor of Physics

Acknowledgments

First and foremost, I wish to thank my parents Robert and Jennifer, and my brothers Matthew and Jeremy. Without their love, support, and encouragement, my accomplishments thus far would've been impossible.

In my first year in Boston I briefly ran into a friendly, lively, and beautiful young woman named Juliana Pepp. We didn't really meet again until almost two years later, but neither of us had forgotten our first encounter and, thanks to her directness overcoming my modesty, we began dating and quickly fell in love. She's quick with a word and a smile, and her big brown eyes alone are enough to lift my spirits (although I get a talking-to if I don't notice when she puts in some of those fancy colored contacts). Jules has enriched my life immeasurably. In addition, her parents Richard and Adria have shown me a great deal of warmth and kindness, making their home as open to me as my own family's.

At the beginning of my second year at MIT, just when I began looking for a new apartment, I came across a mass-mailing to the entire Physics department from a fellow grad student, Stephen Moss, who was looking for a roommate. Expecting a stereotypically small, pale, and dull MIT student with potentially poor grooming habits, I begrudgingly answered the message and went to visit this Stephen Moss at his apartment in Cambridgeport. Stephen turned out to be a normal and friendly guy, I took the apartment, and we quickly became the best of friends. In those three years at Three Rockingham Place, I'd always had a compatriot to discuss the vagaries of experimental physics, the trials and tribulations of grad school (this was especially useful to me, since Stephen was a year ahead), or how the Nebraska Cornhuskers were going to maul the other college football team that coming Saturday.

Through my time as an Officer in the United States Army Reserve I've met quite a few outstanding people, but the most significant of those to me is my good friend Timothy O'Brien. My first year in Boston, Tim took me under his wing and gave me an introduction to Boston as only a local from Dorchester could. Over the course of my stay here, he taught me things that I would've never learned or experienced

if I'd hid in the safe confines of MIT as so many students unfortunately do: how to bang nails like an experienced carpenter, how to hunt for lobster 70 feet below the surface of Boston Harbor on so many beautiful summer days, how to extract yourself from 100 angry Irishmen in a South Boston bar, and how to be an outstanding Army Officer. Most importantly, he taught me that you don't have to have a fancy degree on the wall to be considered "wicked smart".

The other graduate students in the lab provided a constant source of mutual support and companionship. We studied and/or helped each other with the dreaded General Exams, as well as keeping each other from going stark raving mad when working to get an experiment to function or to solve a problem at an ungodly hour of the morning when fatigue and the eerie quiet combine to become quite unnerving. My longtime officemate Jamie Rollins has made enjoyable the time spent in our windowless office which lacked any form of circulation, much less temperature control (the temperature controls located on the walls of both offices we shared were proven long ago to be nothing more than placeboes). He's been a constant source of conversation and good humor, and has also been an excellent friend. I'll miss our going to lunch every day and smoking out on the loading dock when we both needed a break. Rana Adhikari was the grad student who never appeared to go home, and was always willing to lend his input (often unsolicited, but also often useful) if I was doing something interesting, but was also willing to go get beer and burritos on a whim from the local Tex-Mex restaurant. Peter Csatorday always had a slightly pessimistic view of things, but was always willing to escape from the lab at the drop of a hat and continues to be my friend a few years after he graduated. Shourov Chatterji entered grad school at the same time as I, and we both helped each other soldier through classes and exams in those first few years. We also had quite a few lively snowball fights in the parking lot outside (which I, of course, won). I wish you all the best of luck, and I hope we keep in touch.

The laboratory support staff provided me with a great deal of practical help whenever I needed it. The laboratory computer-guru Keith Bayer wasn't your typical system administrator (who usually seem convinced of the utter irrelevance of

the “user”), and was always willing to help with any hardware or software-related issues. Besides going above and beyond the call of duty on many occasions in the lab, Keith introduced me to Northeastern University Hockey games which, while the team seems to lose more often than not, are nonetheless cheap and entertaining. Our laboratory technician Myron MacInnis, despite being in constant demand for high priority projects in the lab, always found time to help me attain and maintain the equipment necessary for my experiment, while also showing me how to properly use all of the neat tools in the machine shop. Left to my own devices, I don’t think I would’ve ever been able to figure out how to even identify, much less properly use, all the highly specialized, bizarre-looking, instruction-less little tools like the “edge finder”. I’d also probably be missing a couple of fingers by now.

I consider myself extremely lucky to have joined, right from the start, a laboratory which was chocked full of highly experienced and competent experimental physicists. In my first year at MIT, Mike Zucker and Peter Fritschel asked me to investigate their idea of using a heating ring to fix power absorption problems in Advanced LIGO, which eventually blossomed into the totality of this work. Mike served as my de facto thesis adviser, providing month-to-month direction for the most my graduate career, and conversations with Peter provided me with technical insight that I most certainly wouldn’t have come upon by myself. Besides the guidance of Mike and Peter, David Ottaway’s optics wisdom was crucial to getting my experiment to work properly after I wrestled with it with little or no success for a great deal of time (measured in years).

My faculty adviser, Rainer Weiss, was the type of guy who’d be the first to tell me if I was destined to fall flat on my face in some endeavor, but would also be the first to extend a hand to help me up provided I put my best foot forward in the first place. His candid and outgoing demeanor coupled with an apparent inability to mince words was quite refreshing (and quite entertaining, when directed at others). Most importantly, though, I’ve gained a knowledge and respect of experimental physics that I don’t think I would’ve gotten under any other adviser.

Finally, I want to thank everyone in the MIT LIGO Laboratory. Because of all

of you, I finish my graduate career with the feeling that I had a very pleasant and rewarding experience.

Contents

Introduction	19
1 Laser Interferometric Detection of Gravitational Radiation	21
1.1 Gravitational Radiation	22
1.2 Laser Interferometry for Gravitational Wave Detection	30
1.2.1 Basic Sensitivity Limitations	32
1.2.2 LIGO	36
2 The Problem and The Solution	47
2.1 Optical Power Handling Issues in Advanced LIGO (The Problem)	48
2.1.1 Common Mode Thermal Effects	56
2.1.2 Differential Mode Thermal Effects	67
2.1.3 Thermal Effects on Sensitivity	72
2.1.4 Summary	77
2.2 Thermally Adaptive Optics (The Solution)	77
2.3 Displacement Noise Induced by Thermal Actuation	80
3 The Numerical Tools	87
3.1 Numerical Models of a Single Optic	88
3.1.1 The Finite Element Model	89
3.1.2 The Thermo-optic Effect	90
3.1.3 Thermoelastic Deformation	102
3.1.4 The Elastooptic Effect	114

3.2	Modeling the Entire Interferometer	117
3.2.1	Initial LIGO	119
3.2.2	Sapphire Advanced LIGO	122
3.2.3	Silica Advanced LIGO	126
4	Heating Ring Thermal Compensation	131
4.1	Simple Heating Ring	132
4.1.1	Practical Limits on Heating Ring Design	134
4.1.2	Optimum Ring Compensation of Input Test Masses	136
4.1.3	Conclusions	138
4.2	Shielded Heating Ring	139
4.2.1	Optimum Shielded Ring Compensation	141
4.2.2	Optimum Shielded Ring Compensation acting on a Compensator Plate	143
4.2.3	Conclusions	147
4.3	Effects on the Entire Interferometer	148
4.3.1	Sapphire Advanced LIGO	149
4.3.2	Silica Advanced LIGO	153
4.4	Noise Considerations	157
5	Scanning Laser Thermal Compensation	161
5.1	Finding the Actuation Basis	162
5.1.1	Choosing a Scan Pattern	166
5.1.2	Using Zernike Polynomials as Basis Functions	168
5.1.3	Using Influence Functions as Basis Functions	171
5.1.4	Example: Compensating a Point-absorber	174
5.1.5	Addressing the time dependence of scanning laser actuation	178
5.2	Noise Considerations	182
6	The Proof-of-Principle Experiment	185
6.1	General Setup	186

6.2	Measuring Material Parameters	192
6.2.1	Systematic Parameters	197
6.2.2	Fused Silica Measurement	198
6.2.3	Sapphire Measurement	202
6.3	Heating Ring Thermal Compensation	209
6.3.1	Heating Ring Results	211
6.4	Scanning Laser Thermal Compensation	216
6.4.1	Scanning Laser Results	217
	Conclusion	223
	A Tables of Parameters	225
A.1	Materials Parameters	225
A.2	Instrument Parameters	227
	B Calculation of the Heating Ring Intensity Profile	233
B.1	Intensity Distribution on the Optic's Face	234
B.2	Intensity Distribution on the Optic's Edge	235
B.3	Intensity Distribution for the Shielded Ring	236

List of Figures

1-1	Laser interferometer's iteration with a passing gravitational plane wave.	29
1-2	Schematic diagram of LIGO.	36
1-3	Simple over-coupled optical cavity.	37
1-4	The spatial mode of a simple optical cavity.	43
2-1	Wavefront distortion caused by mirror heating.	50
2-2	Mirror heating in advanced LIGO.	54
2-3	Thermal effects in an arm cavity.	57
2-4	Electric fields in the Power Recycling Cavity.	59
2-5	Electric fields at the Beam Splitter.	67
2-6	Maximum operating power versus differential and common distortion parameters.	75
2-7	Thermally compensating absorption-induced thermal distortions.	78
3-1	2D Numeric vs. Analytic: Thermal Lensing, Substrate Absorption	99
3-2	2D Numeric vs. Analytic: Thermal Lensing, Coating Absorption	100
3-3	3D Numeric vs. Analytic: Thermal Lensing	101
3-4	2D Numeric vs. Analytic: Thermoelastic Deformation, Substrate Absorption	110
3-5	2D Numeric vs. Analytic: Thermoelastic Deformation, Coating Absorption	111
3-6	3D Numeric vs. Analytic: Thermoelastic Deformation, Bulk Absorption	112
3-7	3D Numeric vs. Analytic: Thermoelastic Deformation, Surface Absorption	113

3-8	The elasto-optic effect in M-axis sapphire.	116
3-9	Melody model of an idealized Initial LIGO.	120
3-10	Melody model of the Livingston interferometer.	121
3-11	Melody model of an idealized Advanced LIGO with sapphire test masses.	124
3-12	Melody model of a Advanced LIGO with sapphire test masses with differential mode absorption.	125
3-13	Melody model of an idealized Advanced LIGO with fused silica test masses.	128
3-14	Melody model of a Advanced LIGO with fused silica test masses with differential mode absorption.	129
4-1	Coordinates for the simple heating ring method of thermal compensation.	132
4-2	Parameter space for ring compensation of a Advanced LIGO Sapphire ITM	137
4-3	Advanced LIGO Sapphire ITM, 1W Probe Absorption, Ring Compensated	138
4-4	Coordinates for the shielded heating ring method of thermal compensation.	140
4-5	Advanced LIGO Sapphire ITM, 1W Probe Absorption, Shielded Ring Compensated	142
4-6	Advanced LIGO Sapphire ITM, 1W Probe Absorption, Shielded Ring on Insulated ITM	143
4-7	Shielded ring thermal compensation in Advanced LIGO.	145
4-8	Advanced LIGO Sapphire ITM, 1W Probe Absorption, Shielded Ring on Silica Compensation Plate.	146
4-9	Tolerance on shielded ring placement.	147
4-10	Melody model of a thermally compensated Advanced LIGO with sapphire test masses.	151
4-11	Melody model of a thermally compensated Advanced LIGO with sapphire test masses and differential absorption.	152

4-12	Melody model of a thermally compensated Advanced LIGO with fused silica test masses.	155
4-13	Melody model of a thermally compensated Advanced LIGO with fused silica test masses and differential absorption.	156
5-1	Different types of scan patterns.	167
5-2	A few Zernike polynomials (astigmatism and coma, as well as an unnamed one of higher order).	169
5-3	The scan pattern used in this example.	175
5-4	A few influence functions.	177
5-5	Thermally compensating a small region of large absorption.	179
6-1	General setup of the experiment.	187
6-2	Setup of the materials measurement.	193
6-3	Surface measurement fitting functions.	197
6-4	Fused Silica reflected data.	200
6-5	Fused Silica steady-state transmitted phase map.	201
6-6	Sapphire crystal structure	202
6-7	C-axis Sapphire reflected data.	204
6-8	C and M-axis steady-state reflected phase map.	206
6-9	M-axis steady-state transmitted phase map.	207
6-10	M-axis steady-state transmitted phase map (blue crosses), and the principal axes of the best fit model to the data (red lines).	207
6-11	Setup of the heating ring experiment.	208
6-12	Measured intensity profile of the pump beam.	210
6-13	Results of the shielded heating ring test.	212
6-14	Heating Ring Data vs. Model.	213
6-15	Heating Ring Coupling Efficiency.	214
6-16	Setup of the scanning laser test.	215
6-17	Results of the scanning laser test.	219
6-18	Scanning laser test initial and final powers.	220

B-1	Schematic of the coordinates used in the calculation of the intensity distribution of the heating ring on the face of the optic.	234
B-2	Schematic of the coordinates used in the calculation of the intensity distribution of the heating ring on the outer edge of the optic.	235
B-3	Schematic of the coordinates used in the calculation of the intensity distribution of the shielded heating ring.	237

List of Tables

2.1	Approximate strengths of thermoelastic deformation and the elastoopic effect relative to the thermooptic effect in each material.	50
2.2	Approximate strengths of thermoelastic deformation and the elastoopic effect relative to the thermooptic effect in sapphire.	52
2.3	Arm cavity failure points.	58
2.4	Arm cavity gain falloff points.	60
2.5	Carrier failure points.	61
2.6	Sideband failure points.	64
2.7	Power absorption parameters for Initial LIGO and Advanced LIGO.	66
2.8	Critical events with increasing power absorption.	66
2.9	Approximate distortion parameters for Initial LIGO and Advanced LIGO.	76
4.1	Optimum ring compensation of Sapphire Advanced LIGO ITM's.	139
4.2	Optimum shielded ring compensation of Sapphire Advanced LIGO ITM's.	143
4.3	Optimum shielded ring compensation of Advanced LIGO ITM's by actuating on a Fused Silica compensation plate.	146
4.4	Summary of the various methods of heating ring thermal compensation.	148
6.1	Systematic parameters and their errors.	199
6.2	Measured materials parameters for fused silica.	200
6.3	Initial measured materials parameters for sapphire.	204
6.4	Measured materials parameters for sapphire.	209

A.1	Fused silica (Amorphous SiO ₂) material parameters.	225
A.2	Sapphire (Crystalline Al ₂ O ₃) material parameters.	226
A.3	LIGO input laser and length parameters.	227
A.4	LIGO mirror parameters.	228
A.5	LIGO cavity parameters.	229
A.6	Heated arm cavity parameters.	230
A.7	Heated power recycling cavity parameters.	231

Introduction

As the first generation of laser interferometric gravitational wave detectors near operation, research and development has begun on increasing the instrument's sensitivity while utilizing existing infrastructure. In the Laser Interferometer Gravitational Wave Observatory (LIGO)¹[1, 46], significant improvements are being planned for installation in ~ 2007 to increase the sensitivity to test mass displacement (hence sensitivity to gravitational wave strain) by improved suspensions and test mass substrates, active seismic isolation, and higher input laser power [12]. Even with the highest quality optics available today, however, finite absorption of laser power within transmissive optics, coupled with the tremendous amount of optical power circulating in various parts of the interferometer, result in critical wavefront deformations which will cripple the performance of the instrument. Discussed is a method of active wavefront correction via direct thermal actuation on optical elements of the interferometer; or, "thermally adaptive optics".

Chapter 1 includes an overview of gravitational waves and their detection using laser interferometry, highlighting the necessity of large amounts of circulating optical power. The deleterious effects on the performance of the instrument as a result of optical absorption are discussed in Chapter 2, including a rough overview of the technique of correcting these absorption-induced optical distortions via thermally adaptive optics. Chapter 3 talks about the numerical models used to simulate thermal distortions (and corrections) in transmissive as well as reflective optics, and the resulting performance of LIGO-like instruments with absorption-induced optical distortions left unchecked. Heating ring thermal compensation, an optimized solution utilizing a

¹<http://www.ligo.caltech.edu/>


simple axisymmetric radiative heating element to remove the gross axisymmetric part of the anticipated thermal distortion in individual optics, is discussed in Chapter 4, while scanning laser thermal compensation, which utilizes a scanning heating beam to remove any unanticipated non-axisymmetric distortions, is the topic for Chapter 5. Finally, Chapter 6 discusses the results of a proof-of-principle experiment that was constructed to demonstrate the technical feasibility of both modes of thermal compensation.

This work, and the LIGO Laboratory, is supported by the National Science Foundation², grants PHY-9210038 and PHY-0107417.

²<http://www.nsf.gov/>

Chapter 1

Laser Interferometric Detection of Gravitational Radiation



This chapter represents the standard “nickel tour” of General Relativity, gravitational waves, and the detection of these waves via laser interferometry (by the Laser Interferometer Gravitational-wave Observatory [1, 46], in particular). The first section discusses General Relativity in brief, concentrating on its linear form in weakly curved spaces (such as that which anchors the reader to his chair) as well as the propagation and generation of gravitational waves. Section 1.2 discusses the practical detection of gravitational waves with a simple Michelson interferometer, leading up to the description of a few of the techniques used in LIGO to improve upon the sensitivity of the simple Michelson. In the end, we’ll find (with no real surprise) that large amounts of light circulating in the instrument is absolutely crucial to the instrument’s sensitivity to passing gravitational waves.

1.1 Gravitational Radiation

The Principal of Special Relativity states that the laws of physics appear identical to individual observers in all inertial frames. Any coordinate transformation between inertial frames which satisfies this principle (that is, leaving Newton’s Laws as well as Maxwell’s Equations invariant) is termed a “Lorentz transformation”. An upshot is that these transformations (and uniquely the Lorentz transformations, as it turns out [45]) also leave invariant the interval between events in space-time (interchangeably termed “proper distance” as well as “proper time”), defined as:

$$ds^2 \equiv -c^2 dt^2 + dx^2 + dy^2 + dz^2 \equiv \eta_{\mu\nu} dx^\mu dx^\nu$$

where η is the Minkowski tensor, which defines “flat” space-time, and we’ve used the Einstein summation convention for repeated indices. Note that if two events are connected by a pulse of light, then necessarily $ds = 0$.

In General Relativity the Principal of Equivalence is introduced, which states that at any space-time point in an arbitrary gravitational field it is possible to find locally (within a sufficiently small space-time region) an inertial reference frame. In other

words, space-time is locally flat but free to contain smooth curvature on a large scale, hence can be represented in mathematical terms as a four-dimensional differentiable manifold. As such, the space-time interval is written more generally:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu$$

where the 4×4 matrix g is the “metric tensor” which describes the deviation of local inertial coordinates ξ^μ from the global coordinates x^μ (which are free to be anything that span 4 dimensions):

$$g_{\mu\nu} \equiv \frac{\partial \xi^\alpha}{\partial x^\mu} \frac{\partial \xi^\beta}{\partial x^\nu} \eta_{\alpha\beta}. \quad (1.1)$$

If we assume that the metric is only slightly perturbed from flat, we may write:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad \text{with } h_{\mu\nu} \ll 1,$$

and the free space version of Einstein’s field equations¹ can be decoupled and linearized in the perturbation $h_{\mu\nu}$, yielding:

$$\left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2 \right) h_{\mu\nu} = 0. \quad (1.2)$$

This equation supports a generic plane wave solution for each element of h :

$$h_{\mu\nu} = A_{\mu\nu} e^{ik_\alpha x^\alpha}.$$

All *ten* elements of the symmetric matrix of wave amplitudes $A_{\mu\nu}$ are not independent, however, as we are free to specify both a direction of propagation as well as a coordinate system and gauge. Choosing Cartesian coordinates and utilizing all of the gauge freedom available, the matrix of amplitudes $A_{\mu\nu}$ can be constrained to be traceless and transverse [27, 38, 45] (meaning that the disturbance is a transverse wave, i.e., $A_{\mu z} = 0$ for z -axis propagation). Assuming propagation along the z -axis,

¹A rather intimidating set of 10 second order coupled partial differential equations over 4 variables; see the standard textbooks [27, 38, 45] for their full glory.

the generic plane wave solution now takes the form:

$$\{h_{\mu\nu}(z,t)\} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} e^{i\omega(\frac{z}{2\pi c}-t)}$$

which is seen to have *two* independent components (polarizations): h_+ (“h-plus”) associated with A_{xx} and h_\times (“h-cross”) associated with A_{xy} .

What does all this mean, in terms of something that can be measured? Consider a freely falling test particle at Cartesian coordinate (x, y) relative to a freely falling observer at the origin in flat space. The observer measures the position of the test particle by pulsing a strobe light and measuring the time of flight as well as the angle of the return from the test particle. Suppose a gravity wave passes along the z -axis, with wavelength much larger than the distance from the observer to the test particle. Due to the space-time curvature induced by the wave, the local inertial coordinates of the observer change relative to the global coordinates (on which the test particle is “attached”), as defined by the metric in equation 1.1. If the gravity wave is + polarized, the coordinates of the test particle measured by the observer in his coordinates then becomes:

$$(x, y)_{\text{new}} = (x, y)_{\text{old}} + \frac{1}{2}h_+ \sin(\omega t) \cdot (x, -y)_{\text{old}}$$

Similarly, if the wave is \times polarized:

$$(x, y)_{\text{new}} = (x, y)_{\text{old}} + \frac{1}{2}h_\times \sin(\omega t) \cdot (y, x)_{\text{old}}$$

Note that the effect caused by the \times wave is identical to that of the + polarization rotated 45 degrees about the z -axis. Figure 1-1 diagrams this result for the + polarization on a Cartesian grid of test particles.

It is clear that Einstein’s field equations support wave solutions that are measur-

able (in theory, at least), but what sort of phenomena generate gravitational waves and what strengths can we expect? Due to conservation of mass/energy there can be no fluctuating mass monopole moment, hence no gravitational radiation. Similarly, mass dipole radiation is forbidden by conservation of linear momentum and the mass analogue of magnetic dipole radiation is forbidden by the conservation of angular momentum. Now that we've run out of conservation laws, the leading term in the generation of gravitational radiation will be due to an oscillating mass quadrupole moment I_{ij} (equivalent to the system's inertia tensor), and the far-field gravitational wave takes the form:

$$h_{ij} = \frac{1}{R} \left(\frac{2G}{c^4} \ddot{I}_{ij} \right). \quad (1.3)$$

Perhaps the simplest example of an oscillating quadrupole moment is a pair of masses locked in a stable circular orbit about each other. Applying equation 1.3 then results in a circularly polarized gravitational wave with a frequency twice that of the orbital frequency and an amplitude which can be expressed in surprisingly simple terms:

$$f_{gw} = 2 \frac{\omega_{orb}}{2\pi} = \frac{c}{\pi R_S} \left(\frac{R_S}{r_0} \right)^{\frac{3}{2}} \lesssim \frac{c}{\pi R_S} \sim 30 \text{ kHz} \left(\frac{1 M_\odot}{\mu} \right) \quad (1.4)$$

$$|h| \approx \frac{R_S}{R} \left(\frac{R_S}{r_0} \right) \lesssim \frac{R_S}{R} \sim 1 \times 10^{-21} \left(\frac{\mu}{1 M_\odot} \right) \left(\frac{150 \text{ Mpc}}{R} \right)$$

where r_0 is the orbital separation of the masses, R is the distance from the system to our detector, and R_S is the Schwarzschild radius of the reduced mass of the system:

$$R_S \equiv \frac{2G}{c^2} \frac{M_1 M_2}{M_1 + M_2}.$$

The listed upper bounds assume $r_0 \sim R_S$, at which point the objects must necessarily merge (unless they did so at a larger separation, which happens for binaries containing anything other than black holes).

It's clear that gravitational waves can only be expected at relatively low frequencies (compared to electromagnetic waves) and exceedingly weak amplitudes, even for

black hole/black hole inspirals which represent the strongest and fastest conceivable sources of gravitational radiation. Other weaker and slower potential sources are any compact astrophysical phenomena which involves the rapid rotation of a body without perfect spherical symmetry (i.e., the axis of rotation is not aligned with the body's principal moment of inertia). Potential detectable examples for Advanced LIGO include (summarized from the review by Thorne [42] as well as the specific sources listed below):

- **Binary Inspiral of Compact Objects.** These are the most powerful sources of gravitational radiation, but are also short-lived (less than a few minutes in LIGO's band of detection) and rather infrequent per galaxy (optimistic estimates are about 10^{-5} per year per galaxy). A properly functioning Advanced LIGO interferometer, however, can expect to see 100 to 2000 such events per year, primarily due to $10 M_{\odot}$ Black Hole/Black Hole mergers in globular clusters at ranges out to $z = 0.4$ [33, 20].
- **Non-Axisymmetric Pulsars.** For known astrophysical objects rotating at well known frequencies within LIGO's detection band, one can coherently integrate the instrument's output over long periods of integration time to constrain the gravitational wave emission of the source. For pulsars in the galaxy, LIGO can thus put constraints on the ellipticity of the star perpendicular to its rotational axis. The gravitational wave amplitude at earth due to a pulsar of frequency f and ellipticity ϵ at a distance R goes as $h \propto \epsilon f^2/R$. For Advanced LIGO, the detectable ellipticity for *known* sources and 3 months of integration time is approximately:

$$\epsilon \gtrsim 2 \times 10^{-6} \left(\frac{100 \text{ Hz}}{f} \right)^2 \left(\frac{R}{10 \text{ kpc}} \right).$$

For 1 full year of integration time, there are nearly 100 known pulsars (including the Crab Pulsar) which may be detected by Advanced LIGO, provided that the measured spin-down is exclusively due to gravitational wave emission [19]. For unknown sources (spinning neutron stars that are radio quiet), the sensitivity

to ellipticity is diminished by ~ 10 .

- **Low-Mass X-ray Binaries.** These systems are neutron stars which occasionally accrete mass from a companion star and emit X-rays, although are not observed to increase in spin frequency (as one would expect, since angular momentum is transferred to the neutron star from the companion). A seemingly plausible explanation is that the accretion induces an asymmetry in the inertia tensor which is abruptly radiated away gravitationally in such a way as to exactly cancel the torque due to accretion [5]. There are 10 such known sources in the galaxy which may be detectable by Advanced LIGO (including Sco X-1), if this is indeed the correct mechanism [19].
- **R-Modes in Newborn Neutron Stars.** This mechanism is a resonant “sloshing” of the neutron fluid in a newly formed neutron star rotating very rapidly, and is a means by which great amounts of angular momentum can be radiated away. Although it now appears that r-modes are heavily damped by mechanisms other than the gravitational radiation reaction (e.g., a solid crust and a magnetic field), such modes have been observed in models to generate 100 second bursts of gravitational waves with frequencies slightly less than 1 kHz and an amplitude such that sources could be seen by Advanced LIGO out to 20 Mpc [31] (where the neutron star birth rate is a few per year; however, it must be noted the author claims that the amplitude found in the simulation is probably overly optimistic).
- **Rotational Instabilities during Core Collapse.** This exact magnitude or waveform of this source remains poorly understood due to the complicated dynamics involved in stellar collapse, whether it be supernovae or accretion-induced collapse. However, recent numerical simulations indicate that the core “bounces” one or more times near the end of the collapse on the scale of tens of milliseconds, each of which can excite numerous non-spherical modes (the core “rings”) which radiate sufficient gravitational waves to be detectable by Advanced LIGO out to 20 Mpc [8] (where there should be a few events per

year).

With these potential sources in mind, a properly functioning Advanced LIGO can *optimistically* expect somewhere between 0.3 and 6 events per day, primarily due to binary mergers of compact objects in globular clusters (and a few tens of events per year from the other sources mentioned). If Advanced LIGO fails to reach its designed sensitivity goal by a mere factor of 2, then this *optimistic* event rate is decreased by a factor of ~ 8 (if the instrument's effective range is decreased by n , then the volume of space observed decreases as $\sim n^3$) yielding a rate somewhere between an event per month and an event every other day for binary mergers, and only a hand-full of events per year for all the other aforementioned sources. Fully attaining the designed sensitivity is thus crucial to obtaining a guaranteed scientific payoff.

Later in this chapter we'll see that very large amounts of optical power are necessary to interferometrically measure displacements sufficient to detect gravitational waves, and these tremendous powers in turn cause the optics of the instrument to distort via nonzero levels of optical absorption, which degrade the high power beam throughout the instrument and decrease its effective sensitivity. The degree of degradation due to absorption-induced optical distortions is dependent on the choice of optical material (fused silica versus sapphire for Advanced LIGO) and configuration, but will be seen to vary from a factor of 2 for a very optimistic sapphire detector (thus decreasing the anticipated event rate by almost 10, as discussed in the previous paragraph) to a factor of ~ 10 for a moderately optimistic fused silica instrument (thus yielding an optimistic event rate of only a few per year, which is *almost no better than that of Initial LIGO*). The motivation of this thesis is to examine a means of actively correcting these thermal distortions via radiative thermal actuation, thus restoring proper high power operation of Advanced LIGO in a manner "gentle" enough that the optics under actuation do not get physically perturbed to the point of compromising the displacement sensitivity of the instrument (recall from above that gravity waves from solar mass Black Hole/Black Hole inspirals in the Virgo Cluster will register as position fluctuations $\leq 10^{-18}$ m/ $\sqrt{\text{Hz}}$ above 10 Hz over 4 km).

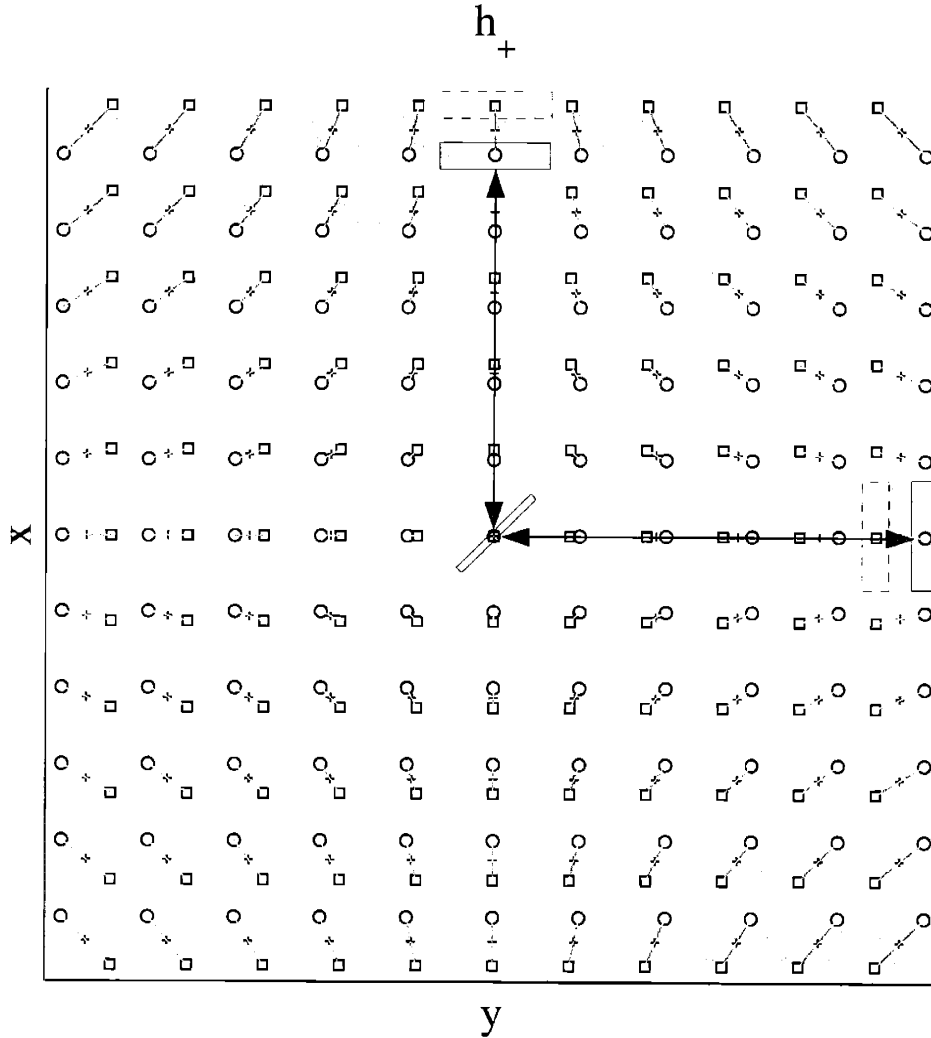


Figure 1-1: A regular grid of free test particles interacting with a passing gravitational wave of the form $h_+ \sin \omega t$, as seen by an observer at the origin. The $+$ marker indicates the position of the test particles at $t = 0$, the \square markers correspond to $\omega t = \pi/2$, and the \circ markers corresponds to $\omega t = 3\pi/2$. Interferometrically measuring the position of the free mirrors (the large boxes) by recombining coherent optical beams at a beamsplitter (the observer) allows for detection of the passing wave.

1.2 Laser Interferometry for Gravitational Wave Detection

It is now clear that a passing gravitational wave can be measured by free falling observer measuring the apparent positions of one or more freely falling test particles in the plane normal to a passing gravitational wave. A classic Michelson Interferometer with a 50/50 beamsplitter at the origin and end mirrors at Cartesian coordinates $(L_x, 0)$ and $(0, L_y)$ (recall the mirrors shown in figure 1-1) appears particularly well suited to detect the passing of an h_+ gravitational wave, in that it facilitates a precise comparison of the time light takes to travel the distance between the beamsplitter and each end mirror. More precisely, suppose that an optical field $E_{\text{BS}}e^{i\omega t}$ is input to a beamsplitter of real amplitude reflectivity r and complex transmissivity it ² and the total time of flight for light along the x and y arms are τ_x and τ_y , respectively. The fields returning to the beamsplitter will have accrued phase shifts $\phi_x = \omega\tau_x$ and $\phi_y = \omega\tau_y$:

$$E_{xr} = itE_{\text{BS}}e^{i\omega t}e^{i\phi_x} \quad E_{yr} = rE_{\text{BS}}e^{i\omega t}e^{i\phi_y}$$

which can be measured by examining the power exiting the “dark port” (the path along $-y$ from the beamsplitter, and the name will become obvious in a moment):

$$P_{\text{DP}} = |rE_{xr} + itE_{yr}|^2 = 4t^2r^2P_{\text{BS}} \cos^2\left(\frac{\Phi}{2}\right).$$

where P_{BS} is the input power and $\Phi \equiv \phi_x - \phi_y$ is the total phase difference of the light returning from each arm. We are free to choose the nominal value of Φ about which we will attempt to detect gravity waves, although the “dark fringe” ($\Phi = \frac{\pi}{2}$), is almost

²For this entire work, the amplitude reflectivity of any mirror is taken to be purely real while the amplitude transmissivity is purely complex. The *actual* phase shift between reflected and transmitted beams is dependent on the type of dielectric coating on the mirror as well as the choice of reference planes where the phase of the reflected and transmitted fields are compared. As long as energy conservation (unitarity of the total scattering matrix) and time reversal (reciprocity of the total scattering matrix) are always satisfied, we are free (although heavily constrained) to choose the phase shifts among the elements of reflection and transmission for both directions of travel through the mirror. See Chapter 11 of [39] for a detailed discussion of transmission and reflection through dielectric mirrors.

always utilized to deaden the sensitivity of the output signal P_{DP} to fluctuations in the input power P_{BS} , as well as allowing for “power recycling” of the light reflected back towards the laser (this is discussed later). Assuming a 50/50 beamsplitter, the power measured at the output is then:

$$P_{\text{DP}} \approx P_{\text{BS}} \frac{\Phi^2}{4} \quad (1.5)$$

When an arbitrary plane gravitational wave of frequency ω_g passes the interferometer in the z direction, the light transit time for each arm can be found by using the fact that the proper distance between the beamsplitter and mirrors is identically zero when connected by a beam of light, i.e.,:

$$\begin{aligned} 0 = -c^2 \frac{\tau_x^2}{4} + L_x^2 + h_+(\omega_g t) L_x^2 &\Rightarrow \tau_x \approx \frac{2L_x}{c} \left(1 + \frac{1}{2} h_+(\omega_g t) \right) \\ 0 = -c^2 \frac{\tau_y^2}{4} + L_y^2 - h_+(\omega_g t) L_y^2 &\Rightarrow \tau_y \approx \frac{2L_y}{c} \left(1 - \frac{1}{2} h_+(\omega_g t) \right) \end{aligned}$$

where we have explicitly assumed that the period of the gravitational wave is much larger than the nominal light transit time along each arm. The detected phase difference of the returning beams is now:

$$\Phi = \frac{2\pi c}{\lambda} (\tau_x - \tau_y) = 2 \frac{2\pi}{\lambda} (2\Delta L + L h_+(\omega_g t)) \quad (1.6)$$

where $L \equiv \frac{L_x + L_y}{2}$ is the nominal common length of the arms, and $\Delta L \equiv \frac{L_x - L_y}{2}$ is the differential length of the arms in the absence of the gravitational wave. Note that the passing gravitational wave has an identical effect on the measurement as differential fluctuations in the positions of the end mirrors in their own rest frames (e.g., if they’re each kicked around by some local influence):

$$2 \frac{\Delta L(t)}{L} \leftrightarrow h_+(t) \leftrightarrow \frac{1}{2L} \frac{\lambda}{2\pi} \Phi(t)$$

1.2.1 Basic Sensitivity Limitations

The measurement of a passing gravitational wave hinges on the degree to which the experimenter can resolve the phase difference Φ (termed “splitting the fringe”), as well as necessitating the reduction the local random fluctuations in mirror positions ($\Delta L(f_g)$) to a level smaller than the effect of the passing gravitational wave ($Lh_+(f_g)$) at the gravitational wave frequencies f_g of interest. While there is a veritable plethora of physical effects that can have an effect on the instruments sensitivity, we’ll briefly discuss the three major players for a ground based, room temperature laser interferometer: seismic noise, thermal noise, and optical readout noise.

Seismic noise is the coupling of ground motion into the local position of the interferometer’s mirrors, which must be held against earth’s gravity while being allowed to move freely in the plane of the earth’s surface (but kept from slowly drifting away from the nominal position which maintains a dark fringe at the instrument’s anti-symmetric port). LIGO I mirrors are suspended from a series of 4 damped oscillators (a pendulum on a heavy table on springs which rests on another table on springs, etc.) each with resonant frequency ~ 1 Hz, and each serves to reduce the coupling of ground motion into mirror motion well above resonance (1 Hz) by an additional factor of $1\text{Hz}^2/f^2$, resulting in the approximate seismic motion:

$$\delta x_{\text{seismic}}(f) = 10^{-11} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{10 \text{ Hz}}{f} \right)^{10}.$$

This passive means of isolation creates a steep “seismic wall” to sensitivity as seismic motion quickly grows larger for lower and lower frequencies. Constraints on space as well as considerations regarding creep and the large amplitude of motion near resonance complicates matters when adding more and more stages of passive isolation, however. Advanced LIGO envisions attaining seismic motions at the level of:

$$\delta x_{\text{seismic}}(f) = 10^{-19} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{10 \text{ Hz}}{f} \right)^8$$

by suspending test masses from quadruple pendula which are mounted on an actively

controlled system of seismic isolation.

Thermal noise is motion of the mirrors' surfaces caused by the fact that they are held near room temperature (determined by the temperature of the vacuum enclosure). An energy $k_B T/2$ is stored in every internal mechanical mode of the mirror; however, choosing a material with very low internal mechanical loss (i.e., a high Q) results in most of the thermal vibrational energy residing in very narrow frequency bands around the mirror's internal modes. The spectral density of thermal motion for a single mirror mode at frequency ω_0 can be approximately expressed as [37]:

$$\delta x_{\text{thermal}}(\omega) = \sqrt{\frac{8k_B T \theta}{m \omega \omega_0^2}} \sqrt{\frac{\omega_0^4}{(\omega_0^2 - \omega^2)^2 + \omega_0^4 \phi^2}}$$

where m is the mirror's mass and θ is the loss angle for that mode (equivalent to $1/Q$). Two optical materials which have been identified as suitable for gravitational wave interferometers are amorphous fused silica (which has a sound speed of about 6 km/s and loss angle of 3×10^{-8} , and is used in LIGO I) and crystalline sapphire (10 km/s sound speed and 5×10^{-9} loss angle, which is the target material for Advanced LIGO). Since the sound crossing time for both of these materials is so short, the materials' lowest frequency internal resonances will be at the very high end of the gravitational wave spectrum: $f_0 \approx \frac{10 \text{ km/s}}{0.3 \text{ m}} \approx 30 \text{ kHz}$. For frequencies well below the two lowest resonances of a sapphire test mass at room temperature (with an aspect ratio such that these two modes are degenerate), the thermal fluctuations of a sapphire mirror's surface will be approximately:

$$\begin{aligned} \delta x_{\text{thermal}}(f) &\approx 2 \sqrt{\frac{8k_B T \theta}{m \omega \omega_0^2}} \\ &\approx 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{100 \text{ Hz}}{f}\right)^{\frac{1}{2}} \left(\frac{30 \text{ kHz}}{f_0}\right) \left(\frac{\theta}{5 \times 10^{-9}}\right)^{\frac{1}{2}} \left(\frac{m}{40 \text{ kg}}\right)^{\frac{1}{2}} \end{aligned}$$

which will fully dominate over seismic noise for $f > 10 \text{ Hz}$ in Advanced LIGO.

Lastly we consider "optical readout noise", which is simply a manifestation of the fact that our gravitational wave interferometer is really just a giant Heisenberg

Microscope. To get a better fix on the position of the test masses, we'll see that we must throw more and more photons at them. Eventually, though, we'll be striking the test masses with so many photons and reaching such a fine measurement of position that the back-action of the measurement (the randomly arriving photons perturbing the mass's momentum) will begin to compromise the position measurement itself. First, let's disregard the perturbations of the test masses by our probe light (radiation pressure noise), and consider how well we can "split the fringe" with a perfect detector filled with grainy photons. If power P_{DP} is detected at the dark port, then the fluctuation in this detection is Poissonian:

$$\delta P_{\text{DP}} = \sqrt{2 \frac{hc}{\lambda} P_{\text{DP}}}.$$

Relating this to the corresponding fluctuation in measured phase gives us "shot noise":

$$\begin{aligned} \delta \Phi_{\text{shot}}(f) &= \frac{d\Phi}{dP_{\text{DP}}} \delta P_{\text{DP}}(f) = \frac{1}{2\sqrt{P_{\text{BS}}P_{\text{DP}}}} \sqrt{2 \frac{hc}{\lambda} P_{\text{DP}}} = \sqrt{\frac{2hc}{\lambda P_{\text{BS}}}} \\ &\approx 10^{-11} \frac{\text{rad}}{\sqrt{\text{Hz}}} \left(\frac{100 \text{ W}}{P_{\text{BS}}} \right)^{\frac{1}{2}}. \end{aligned} \quad (1.7)$$

Note that more power must be incident on the *beamsplitter* in order to improve the phase measurement³. This result is not so surprising when we enter it into equation 1.5, and find that this limit simply means that we can't draw any conclusions about phase below the level of one photon detected over the integration time of our detector (we haven't assumed one thus far, which is why $1/\sqrt{\text{Hz}}$'s keep popping up). In terms of the equivalent sensitivity to the differential arm length:

$$\delta x_{\text{shot}}(f) = \frac{\lambda}{2\pi} \frac{\delta \Phi}{4} = \frac{1}{8\pi} \sqrt{\frac{2hc\lambda}{P_i}} \approx 10^{-18} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{100 \text{ W}}{P_{\text{BS}}} \right)^{\frac{1}{2}}$$

which appears to heavily dominate over seismic and thermal noise (Advanced LIGO plans to utilize a 100 Watt laser); however, there are means of boosting both the power at the beamsplitter as well as the phase response to length fluctuations which

³Splitting a fringe to one part in 100 billion might seem extremely small, but experiments attaining the LIGO I level of $\delta \Phi \sim 10^{-10} \text{ rad}/\sqrt{\text{Hz}}$ were successful in 1998 [21, 11]

are discussed later in this chapter.

As hinted before, one can't simply increase the power in the interferometer *ad infinitum* because of quantum mechanics: measuring the position of an object necessarily perturbs its momentum, which in turn creates uncertainty in its position⁴. If optical power P is reflected off a mirror of mass m , the amplitude spectral density of the resulting motion can be calculated in a straightforward manner:

$$\begin{aligned}\delta x_{\text{rad pressure}}(f) &= \frac{2\frac{\delta P}{c}}{m(2\pi f)^2} = \frac{1}{mf^2} \sqrt{\frac{hP}{2\pi^4 \lambda c}} \\ &\approx 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{100 \text{ Hz}}{f}\right)^2 \left(\frac{40 \text{ kg}}{m}\right) \left(\frac{P}{800 \text{ kW}}\right)^{\frac{1}{2}}\end{aligned}\quad (1.8)$$

and is seen to dominate thermal noise at frequencies below 100 Hz for the parameters we have chosen (the choice of these parameters, the incident power in particular, will become clear in the next section).

Writing the local physical displacement noises previously discussed (seismic, radiation pressure, and thermal) in terms of equivalent gravitational wave strain for Advanced LIGO ($L = 4 \text{ km}$):

$$\delta h_{\text{motion}}(f) = \frac{2\sqrt{2}\delta x}{L} \approx \begin{cases} 10^{-22} \frac{1}{\sqrt{\text{Hz}}} \left(\frac{10 \text{ Hz}}{f}\right)^8 & f < 10 \text{ Hz} \\ 10^{-22} \frac{1}{\sqrt{\text{Hz}}} \left(\frac{10 \text{ Hz}}{f}\right)^2 & 10 \text{ Hz} < f < 100 \text{ Hz} \\ 10^{-24} \frac{1}{\sqrt{\text{Hz}}} \left(\frac{100 \text{ Hz}}{f}\right)^{\frac{1}{2}} & f > 100 \text{ Hz} \end{cases} \quad (1.9)$$

we can compare to that for shot noise to see what must be done to make it comparable:

$$\begin{aligned}\delta h_{\text{shot}} &= \frac{1}{2L} \frac{\lambda}{2\pi} \delta\Phi = \frac{1}{2L} \frac{\lambda}{2\pi} \sqrt{\frac{2hc}{\lambda P_{\text{BS}}}} \\ &\approx 10^{-21} \frac{1}{\sqrt{\text{Hz}}} \left(\frac{4 \text{ km}}{L}\right) \left(\frac{100 \text{ W}}{P_{\text{BS}}}\right)^{\frac{1}{2}}.\end{aligned}\quad (1.10)$$

⁴This is because we are measuring position, which is not an integral of the motion. Measuring an integral of the motion (i.e., momentum) while deliberately deadening the detector to the conjugate variable (i.e., position) immunizes the detector from the back-action of the measurement. This type of measurement is termed ‘‘Quantum Non-Demolition’’, and is an active area of research for ‘‘LIGO III’’.

One can reduce this by folding the optical path (effectively increasing L) as well as increasing the amount of optical power incident on the beamsplitter. In the following sections we'll see how LIGO utilizes optical cavities achieve both of these goals, thus reducing shot noise by about 3 orders of magnitude to such a level that it defines the sensitivity floor only at frequencies above 300 Hz.

1.2.2 LIGO




Figure 1-2: Schematic diagram of LIGO. The arms are replaced by resonant optical cavities to fold the optical path in the arms, thus boosting the phase response at the beamsplitter to length changes of the arms. A mirror is placed at the input of the instrument to recycle clean light which would otherwise simply exit the instrument back towards the laser. In order to actively hold all these mirrors in the proper positions to establish resonance, the laser is weakly phase modulated at two different frequencies (these sidebands are shown as dashed lines on the diagram, and are on the order of tens of MHz): one to establish control of the power recycling cavity and one to control the arm cavities.

Figure 1-2 shows a rough layout of the various optical cavities utilized in LIGO

to reduce the effective shot noise. The Power Recycling Cavity serves to recycle light reflected back towards the laser, thus increasing the amount of power incident on the beamsplitter, and the Arm Cavities serve to store light in the arms so each photon experiences many round trips before interfering at the beamsplitter. In what follows, we'll discuss optical cavities and the effect they have on the strain sensitivity of the instrument.

Optical Cavities

The simplest optical resonator consists of two parallel, partially reflective mirrors as diagrammed in figure 1-3. Any light of both proper axial mode (the proper wavelength) and spatial mode (the proper curvature and intensity pattern) that enters the cavity becomes trapped between the mirrors, slowly exiting after some number of round trips determined by the mirror reflectivities and any losses present in the cavity. The higher the mirror reflectivities, the more round trips light will undergo before exiting the cavity.




Figure 1-3: Simple over-coupled optical cavity.

Consider an overcoupled optical cavity of length L , mirror diameter d , and laser wavelength λ such that we may neglect diffraction (i.e., $\frac{L}{d} \frac{\lambda}{d} \ll 1$). Figure 1-3 diagrams the fields in such an optical cavity. Following Siegman [39], the plane-wave circulating field just inside the input surface of this cavity is related to the plane-wave incident

field by:

$$E_c(\lambda) = it_1 E_i(\lambda) + r_1 \sqrt{1 - \mathcal{S}} E_c(\lambda) e^{-2\pi i \frac{2L}{\lambda}}$$

which is simply the sum of the field immediately entering the cavity (the first term) and the circulating field which has propagated one full round trip (the second term), experiencing some small amplitude loss \mathcal{S} (e.g., due to scattering or absorption). The circulating field just inside the cavity is then:

$$E_c(\lambda) = \frac{it_1}{1 - r_1 \sqrt{1 - \mathcal{S}} e^{-2\pi i \frac{2L}{\lambda}}} E_i(\lambda)$$

which reaches a maximum (resonance) when the wavelength λ is an integer multiple of the round trip distance $2L$ (all the light traversing back and forth in the cavity then adds coherently). On resonance, the incident and circulating optical powers are related by the cavity gain \mathbb{G} :

$$\mathbb{G} \equiv \frac{P_c}{P_i} = \frac{1 - r_1^2}{1 - r_1 \sqrt{1 - \mathcal{S}}} \approx \frac{2r_1}{1 - r_1} \quad \text{For } \mathcal{S} \ll (1 - r_1) \text{ and } r_1 \sim 1 \quad (1.11)$$

which is effectively the mean number of bounces that a photon undergoes before exiting the resonator, and can be quite large for values of r_1 near 1. To see this more clearly, we can examine the phase shift of reflected field $E_r = r_1 E_i + it_1 E_c$ relative to the incident field for small deviations δx from resonance:

$$E_r = \left(r_1 - \frac{t_1^2}{1 - r_1 e^{-2\pi i \frac{2\delta x}{\lambda}}} \right) E_i \approx \left(-1 + i \frac{2r_1}{1 - r_1} \frac{2\pi}{\lambda} 2\delta x \right) E_i \quad (1.12)$$

and the phase angle of the reflected field relative to incident field is clearly:

$$\phi = \frac{2r_1}{1 - r_1} \frac{2\pi}{\lambda} 2\delta x = \mathbb{G} \frac{2\pi}{\lambda} 2\delta x.$$

Utilizing optical cavities of gain \mathbb{G}_{ARM} which stretch the full length of the instrument's arms thus serves to amplify the response to differential length, hence to gravitational wave strain. The net effect is that the shot noise limited strain sensi-

tivity is reduced:

$$\delta h_{\text{shot}} = \frac{1}{2\mathbb{G}_{\text{ARM}}L} \frac{\lambda}{2\pi} \delta\Phi = \frac{1}{2\mathbb{G}_{\text{ARM}}L} \frac{\lambda}{2\pi} \sqrt{\frac{2hc}{\lambda P_{\text{BS}}}}.$$

One cannot arbitrarily increase the gain of the arm cavities, however, since this increases the storage time of the cavities ($\tau \approx \mathbb{G}_{\text{ARM}}L/c$). For gravitational waves of periods longer than the storage time of the cavity, the result is that the response of the instrument is diminished (photons in the cavity interact with the wave for a full period or more). Advanced LIGO will utilize arm cavities of gain $\mathbb{G}_{\text{ARM}} \sim 800$ (compared to LIGO I's 120), coupled with the technique of resonant sideband extraction which bypasses the large storage time of these cavities (about 10 ms) for a relatively narrow band of signal frequencies⁵.

Finally, the amount of power incident on the beamsplitter for a fixed input laser power (7 Watts or so is available for LIGO I, and 100 Watts is the goal for Advanced LIGO) through the technique of “power recycling” [18]. When the instrument is held at a dark fringe, nearly all of the input power is reflected from the beamsplitter back into the laser, hence the instrument acts like a gigantic mirror. A partially reflective mirror placed between the laser and the beamsplitter can be used to reflect a large portion of the rejected light back into the instrument. Choosing the mirror transmissivity according to the total round trip loss light experiences between leaving and returning to the recycling mirror, the field reflected from the recycling mirror back to the laser can be made to vanish (recall equation 2.9 with losses included in the second term), thus fully utilizing all power available. LIGO I is able to attain power recycling gains of about 50, corresponding to a total round trip loss of under 2% (not too shabby considering light strikes 250 surfaces in one round trip). Advanced LIGO will only be able to realistically achieve a gain of about 20, due to the many

⁵This technique is also called “signal recycling” or “dual recycling” [18, 10] and is achieved by creating a “signal recycling cavity” though suspending a partially reflective mirror at the dark port and actually reflecting a significant part of the output signal back onto itself into the instrument. While rather counterintuitive, this technique actually serves to make the path from inside the arms to the photodetector fully transparent for a narrow band of signal frequencies determined by the length of the signal recycling cavity (the narrowness of the band is determined by the reflectivity of the signal recycling mirror).

more surfaces light strikes per round trip on account of the much larger arm cavity gain.

Through the use of arm cavities as well as power recycling, the shot noise limited strain sensitivity of Advanced LIGO is now seen to be⁶:

$$\begin{aligned} \delta h_{\text{shot}} &= \frac{1}{2\mathbb{G}_{\text{ARM}}L} \frac{\lambda}{2\pi} \sqrt{\frac{2hc}{\lambda\mathbb{G}_{\text{PRC}}P_{\text{BS}}}} \\ &\approx 10^{-24} \frac{1}{\sqrt{\text{Hz}}} \left(\frac{800}{\mathbb{G}_{\text{ARM}}} \right) \left(\frac{20}{\mathbb{G}_{\text{PRC}}} \right)^{\frac{1}{2}}. \end{aligned} \quad (1.13)$$

Readout and Control

The lengths of the various optical cavities in LIGO must be carefully controlled (at frequencies below the band where one expects to detect gravitational waves) in order to maintain resonance with the light input to the instrument. Here a very brief overview is given, mainly to emphasize the importance of phase modulated sidebands to the proper functioning of the instrument. For a detailed discussion of the readout and control scheme for LIGO, see [40].

To actively control the length of a simple optical cavity, one typically weakly phase modulates the input light and examines the reflected field to obtain an error signal. Weakly modulating the phase of an optical field at frequency ω_m is equivalent to peeling amplitude away from the original optical frequency ω (the carrier) into a sidebands of optical frequency $\omega + \omega_m$ and $\omega - \omega_m$:

$$E_i = |E_i| e^{i(\omega t + \Gamma \sin \omega_m t)} \approx |E_i| \left(e^{i\omega t} + \frac{\Gamma}{2} e^{i(\omega + \omega_m)t} - \frac{\Gamma}{2} e^{i(\omega - \omega_m)t} \right)$$

where Γ is the modulation index ($\Gamma \ll 1$), and we've dropped all terms of order Γ^2 and higher. The modulation frequency is chosen so that the sidebands are antiresonant in the optical cavity when the carrier is resonant, and the power reflected from the cavity is examined at the modulation frequency. If the cavity is slightly shorter than resonance, the reflected carrier suffers a slight phase shift and beats against the fully

⁶While waving our arms in the air and chanting “signal recycling” to make the storage time of the arms go away.

reflected sidebands such that the detected power oscillates at frequency ω_m out of phase with the signal driving the modulator; if it is slightly long, it oscillates in phase with the modulator. Demodulating this signal against the signal driving the phase modulator yields a linear error signal by which cavity length can be actively controlled. The input light in LIGO is phase modulated at two distinct frequencies, thus producing two sets of sidebands: one to control the common length of the power recycling cavity (at about 60 MHz, and never enters the power recycling cavity), and another to control the length of the arm cavities (at about 25 MHz, and is chosen to be antiresonant in the arm cavities while resonant in the power recycling cavity).

Similarly, the differential length of the instrument must be controlled to maintain a dark fringe. Simply examining the power exiting the dark port is not enough, as this yields no information regarding the direction of the differential length causing the leakage. By deliberately building a macroscopic arm length asymmetry into the power recycling cavity (the path from the beamsplitter to the X input test mass is about 30 cm longer than the path to the Y input test mass), the dark fringe condition for the sidebands resonant here (those used to control the arm cavities) becomes different than that of the carrier, and equally leak out the dark port when the carrier is held on a dark fringe. A small differential length change causes a sideband imbalance at the output which beats against the carrier leakage, so that demodulating against the modulator signal yields a linear error signal to control differential length (as well as read out the gravitational wave signal in such a way as to completely de-sensitize the signal from intensity noise on the input light [37]).

Spatial Modes

Thus far we have completely ignored diffraction, which cannot be neglected when considering finite width beams traveling finite distances. On account of diffraction, an initially well-collimated beam of finite width will accrue some amount of curvature and will begin to expand after propagating a finite distance. In an optical cavity the curvature of the end mirrors must match this accrued curvature, thus reflecting beam back onto itself into the cavity. The curvatures of the end mirrors define a single

spatial mode which “fits” into the cavity: too narrow a beam will become accrue too much curvature for the cavity mirrors to compensate, while too wide a beam will be too strongly focused.

The selectivity of spatial modes by an optical cavity is perhaps best examined in terms of Hermite-Gauss modes, which represent a complete, orthonormal set of solutions to the paraxial wave equation:

$$u_n(x, z) = \left(\frac{2}{\pi w^2(z)} \right)^{\frac{1}{2}} e^{-\frac{x^2}{w^2(z)}} e^{-i \frac{kx^2}{2R(z)}} e^{-i \frac{1}{2} \psi(z)} \times \frac{e^{-in\psi(z)}}{2^n n!} H_n \left(\frac{\sqrt{2}}{w(z)} x \right) \quad (1.14)$$

where $H_n(x)$ are Hermite polynomials, and the beam waist w (characteristic spatial width of the beam), wavefront radius of curvature R , and “Guoy phase” ψ are defined as:

$$w(z) \equiv w_0 \sqrt{1 + \frac{z^2}{z_R^2}} \quad R(z) \equiv z + \frac{z_R^2}{z} \quad \psi(z) \equiv \tan^{-1} \frac{z}{z_R}$$

all of which scale with the “Rayleigh Range”:

$$z_R \equiv \frac{\pi w_0^2}{\lambda}$$

which delineates the near-field ($z < z_R$) and the far-field ($z > z_R$) for the propagating beam. Any paraxial optical beam E can be written as a sum of these Hermite-Gauss modes (often called transverse electromagnetic modes and designated TEM_{nm}):

$$E(x, y, z, t) = \sum_n \sum_m c_{nm} u_n(x, z) u_m(y, z) e^{i(\omega t - kz)}.$$

The lowest order mode (TEM_{00}) is a Gaussian beam:

$$E_0(x, y, z) = \frac{2}{\pi w^2(z)} e^{-\frac{x^2+y^2}{w^2(z)}} e^{-i \frac{\pi x^2}{\lambda R(z)}} e^{-i \frac{1}{2} \psi(z)} \approx \frac{2}{\pi w_0^2} e^{-\frac{x^2+y^2}{w_0^2}} \quad \text{for } z \ll z_R \quad (1.15)$$

Now consider a simple optical cavity with end mirrors with radii of curvature R_1 and R_2 ⁷, as shown in figure 1-4. The curvatures of the end mirrors pick out a single

⁷ R is taken to be positive for mirrors that are concave as viewed from inside the resonator.




Figure 1-4: The spatial mode of a simple optical cavity.

Gaussian mode (i.e., uniquely defines a waist w_0 and waist position z_0 such that the wavefront curvature matches both mirror surfaces), provided that the following criterion is satisfied:

$$0 < \left(1 - \frac{L}{R_1}\right) \left(1 - \frac{L}{R_2}\right) < 1.$$

in addition to the end mirrors being large enough in diameter to fully encompass the width of this unique mode. Higher order TEM modes can also be resonant in such a cavity, since their curvatures will also match the mirror surfaces; however, their faster Guoy phase shifts and larger spatial widths⁸ typically prevents all but the fundamental Gaussian mode from resonating.

Sensitivity in the Presence of Optical Imperfections

In this modal picture, small optical imperfections can be viewed in terms of scattering out of the nominal cavity eigenmode. If an eigenmode $|E_0\rangle$ (here assumed to be approximately Gaussian) experiences a phase distortion $\phi(x, t)$, then the total optical

⁸The spatial half-width of the n th Hermite-Gauss mode w_n can be shown to be [39]:

$$w_n(z) \approx \sqrt{n}w(z)$$

power scattered out of this mode is given by:

$$\begin{aligned} \mathcal{S} &= 1 - |\langle E_0 | e^{i\phi(x,y)} | E_0 \rangle|^2 \\ &\approx 1 - \frac{4}{\pi^2 w^4} \left| \iint_{-\infty}^{\infty} e^{i\phi(x,y)} e^{-2\frac{x^2+y^2}{w^2}} dx dy \right|^2. \end{aligned} \quad (1.16)$$

Optical imperfections common to both arms will serve to scatter power (primarily sideband power, as we shall see in the next chapter) out of the fundamental spatial mode of the interferometer's various cavities, thus reducing the amount of power resonating in the instrument.

Furthermore, distortions not common to both arms will prevent perfect interference of the carrier field at the dark port, causing some nominal amount of carrier light to always leak out. This light contains no useful information about the differential length of the instrument, but carries noise nonetheless (the power P of this leakage carries the standard Poissonian intensity fluctuation $\delta P = \sqrt{2\hbar\omega P}$) which contributes to the instrument's total level of shot noise.

Loss of carrier and sideband power as well as unintended carrier power leaking out the dark port will all serve to increase the level of shot noise, hence lower the sensitivity of the instrument. Quantitatively, the impact of optical imperfections on phase sensitivity is calculated as [21, 11]:

$$\delta\Phi_{\text{shot}}(f) = \left(F_{ns} \sqrt{1 + R_{bs} \frac{1 - C}{4 \sin^2(2k_m \delta l)}} \right) \sqrt{\frac{2\hbar c}{\lambda} \mathbb{G}_{\text{PRC}}^{(c)} P_i} \quad (1.17)$$

where F_{ns} is the non-stationary correction factor ($\sqrt{3/2}$ for the current readout scheme, caused by the fact that real RF mixers used to demodulate the output signal would rather see square waves than sine waves), $k_m = 2\pi f_m/c$ is the modulation wavenumber for the resonant sidebands, R_{bs} is the ratio of TEM₀₀ carrier power to TEM₀₀ sideband power at the beamsplitter, $C \equiv (P_{\text{PRC}}^{(c)} - P_{\text{DP}}^{(c)})/(P_{\text{PRC}}^{(c)} + P_{\text{DP}}^{(c)})$ is the fringe contrast ($1 - C$ is the contrast loss), ω_c is the carrier light frequency, $P_{\text{PRC}}^{(c)}$ is the carrier power incident on the beamsplitter, and $P_{\text{DP}}^{(c)}$ is the total carrier

power leaking out of the dark port. For LIGO, $k_m \approx 1/\text{m}$ and $\delta l \approx 0.2\text{m}$ so that $4 \sin^2(2k_m \delta l) \approx 0.25$. If we further assume that the carrier power leaking out the anti-symmetric port is small compared to the TEM_{00} sideband power at the beamsplitter (i.e., $P_{\text{DP}}^{(c)}/P_{\text{PRC}}^{(sb0)} \ll 1$), we can reduce equation 1.17 to:

$$\delta\Phi_{\text{shot}}(f) \approx \left(\sqrt{\frac{3}{2}} \sqrt{1 + 8 \frac{P_{\text{DP}}^{(c)}}{P_{\text{PRC}}^{(sb0)}}} \right) \sqrt{\frac{2hc}{\lambda} \mathbb{G}_{\text{PRC}}^{(c)} P_i} \quad (1.18)$$



The shot noise limited gravitational wave sensitivity of Advanced LIGO in the presence of optical imperfections can thus be written:

$$\boxed{\delta h_{\text{shot}}(f) \approx \left(\sqrt{\frac{3}{2}} \sqrt{1 + 8 \frac{P_{\text{DP}}^{(c)}}{P_{\text{PRC}}^{(sb0)}}} \right) \times \frac{1}{2L\mathbb{G}_{\text{ARM}}} \frac{\lambda}{2\pi} \sqrt{\frac{2hc}{\lambda} \mathbb{G}_{\text{PRC}}^{(c)} P_i}} \quad (1.19)$$

From this, we note that increasing the input power P_i does increase the strain sensitivity of the instrument, provided that this increased power has no effect on the gains in the arm cavities or the power recycling cavity, as well as maintaining a reasonable amount of circulating sideband power and small amount of carrier power exiting the dark port. In the next chapter we'll see that absorption-induced optical distortions will create massive losses in circulating sideband power, diminished arm cavity gain, and increased carrier power leaking out the dark port, thus significantly decreasing the strain sensitivity of the instrument according to equation 1.19.

Chapter 2

The Problem and The Solution



In the previous section, we saw that increasing the probe light circulating in a laser interferometer gravitational wave detector increases the broadband length sensitivity of the instrument up to the point where test mass position perturbations induced by Poissonian fluctuations in the probe light power begin to dominate the shot noise viewed at the instrument's readout. In practice, however, the amount of allowable circulating probe power in a laser interferometer with transmissive elements is limited by the nonzero optical absorption in the substrate and coatings of the test masses and the beam splitter. The temperature dependent index of refraction and a nonzero thermal expansion coefficient of optical materials found suitable for laser interferometric gravitational wave detectors (primarily fused silica and sapphire) ensures that nonuniform temperature increases induced by the absorption of the Gaussian-profiled probe light will result in nonuniform optical path length distortions which will affect both the controllability and length sensitivity of the instrument.

This chapter examines the physical mechanisms by which wavefront distortions are induced by probe light absorption, the effects of these distortions on the performance of the interferometer, and the method of correcting these distortions via thermal actuation.

2.1 Optical Power Handling Issues in Advanced LIGO (The Problem)

First, let's examine what happens in an optical material when it's weakly heated (e.g., by optical absorption). Consider a small block of transmissive optical material with uniform index of refraction n , thickness h , and temperature T_0 . The optical path length through this block is simply $S_0 = nh$. If this small piece of material is slightly heated, we expect the temperature to increase and the block to expand. If the small block is mechanically constrained in some way (e.g., by adjacent blocks also trying to expand), then the material will strain. The increase in both strain and temperature each serve to alter the block's index of refraction by amounts Δn_E and

Δn_T , respectively, and the net change optical path length S through this block is:

$$S = (h + \Delta h)(n + \Delta n_T + \Delta n_E) \approx S_0 + n\Delta h + h\Delta n_T + h\Delta n_E$$

The three effects that we have just identified to modify the optical path through a heated optic are:

1. **Thermooptic Effect:** (also termed “thermal lensing” or “thermal blooming”): The local refractive index changes with temperature T , so that the optical path change over the path S through a heated optic is (see §3.1.2):

$$\Delta S_T = \int_S \Delta n_T ds = \int_S \frac{dn}{dT} \Delta T ds$$

where $\frac{dn}{dT}$ is the thermooptic coefficient.

2. **Thermoelastic Deformation:** The surfaces of a heated optic will, by virtue of a nonzero thermal expansion coefficient α , expand along the optical axis, the magnitude of which can be roughly expressed as (see §3.1.3):

$$\Delta u \approx \alpha \int_S \Delta T ds.$$

3. **Elastooptic Effect:** The local refractive index changes with mechanical strain resulting from thermal expansion, thus changing the optical path length through an optic by approximately (see §3.1.4):

$$\Delta S_E = \int_S \Delta n_E ds \approx -\alpha p_{11} \int_S \Delta T ds$$

where p_{11} is the component of the elastooptic tensor along the probe beam polarization axis.

where we have assumed that the observed optical path distortions are small enough that the shape of the original path S does not appreciably change (e.g., it remains a straight line through the substrate). Each of these effects are approximately pro-

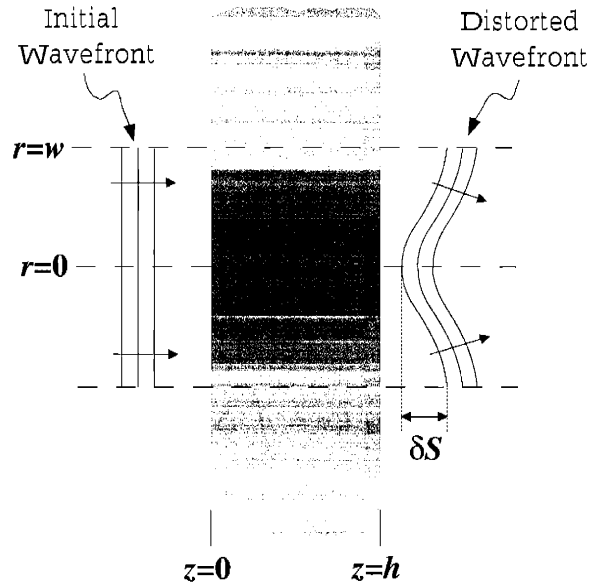


Figure 2-1: Wavefront distortion caused by mirror heating.

	Sapphire	Fused Silica
Thermooptic Effect	1	1
Thermal Expansion	0.8	0.06
Elasto-optic Effect	0.2	-0.01

Table 2.1: Approximate strengths of thermoelastic deformation and the elasto-optic effect relative to the thermo-optic effect. Materials parameters are taken from Tables A.1 and A.2.

portional to the quantity $\int_S \Delta T ds$, and we roughly calculate the relative magnitude of each of these effects in each optical material of interest in Table 2.1. For Fused Silica, thermoelastic deformation and the elasto-optic effect are seen to be negligible compared to the thermo-optic effect; thus these effects can, in general, be ignored when considering transmission through a weakly heated fused silica substrate. This assumption is not valid, however, in sapphire.

In the case where the slight heating is done by uniform optical absorption of a well-collimated beam of wavelength λ propagating in the z direction with nonuniform intensity profile $I(x, y)$, the quantity $\int_S \Delta T ds$ will vary across the optical aperture (x, y) . The three thermal effects discussed above will then impose the nonuniform phase distortion, related to the optical path length change by:

$$\phi(x, y) = \frac{2\pi}{\lambda} \Delta S(x, y)$$

on our originally well-collimated beam. To approximate how the quantity $\int_S \Delta T ds$ varies over the optical aperture, similar to the method of [47], consider a collimated beam of radius w (i.e., $I(r) \approx 0$ for $r > w$) propagating along the z -axis through the center of a cylindrical optic with thermal conductivity k , thickness h , and radius R considerably larger than w (see figure 2-1). If total power P_a is uniformly absorbed along the beam's path in the substrate, then, assuming radial heat flow only, the temperature drop between the radial points $r = 0$ and $r = w$ is approximately:

$$\delta\Gamma \equiv \Delta T|_{r=0} - \Delta T|_{r=w} \approx \frac{P_a}{k(2\pi wh)} w.$$

Integrating this along the optical path through the substrate then yields:

$$\int \delta\Gamma dz \approx \frac{P_a}{2\pi k}. \quad (2.1)$$

From this we see that the quantity $\frac{k}{P_a} \int_S \Delta T ds$ is approximately constant for any choice of substrate material. We can now compute the relative strengths of the thermo-optic effect, thermoelastic deformation, and the elasto-optic effect as we did

	Sapphire	Fused Silica
Thermooptic Effect	1	26
Thermal Expansion	0.8	1.6
Elastooptic Effect	0.2	-0.3

Table 2.2: Approximate strengths of thermoelastic deformation and the elastooptic effect relative to the thermooptic effect in sapphire. Materials parameters are taken from Tables A.1 and A.2.

above, but now between substrate materials. The strengths of these effects compared to the thermooptic effect in sapphire are shown in table 2.2. The distortions are all seen to be amplified (for equal amounts of power absorbed) in fused silica due to its $26\times$ smaller thermal conductivity. For transmitted beams, the distortion through silica is thus $\sim 20\times$ worse than in sapphire for equal amounts of power absorbed. For reflected beams, the thermoelastic “bump” in silica is actually $2\times$ larger than that in sapphire for equal amounts of power absorbed, despite the fact that sapphire has a $\sim 10\times$ larger thermal expansion coefficient.

In absolute terms, the approximate total optical path difference between $r = 0$ and $r = w$ is calculated to be approximately:

$$\begin{aligned}
 \delta S &\equiv \Delta S|_{r=0} - \Delta S|_{r=w} \approx P_a \frac{\beta}{2\pi k} \\
 &\approx 1 \mu\text{m} \left(\frac{1.4 \text{ W/m/}^\circ\text{K}}{k} \right) \left(\frac{\beta}{10 \text{ ppm/}^\circ\text{K}} \right) \left(\frac{P_a}{1 \text{ W}} \right) \\
 \text{where } \beta &= \begin{cases} \frac{dn}{dT} + \alpha - \alpha p_{11}, & \text{in transmission} \\ -2\alpha, & \text{on reflection} \end{cases}
 \end{aligned} \tag{2.2}$$

To lowest order, the thermal distortion δS will act to change an optical element’s effective radius of curvature R . Approximating in terms of a pure thermal radius of curvature R_t :

$$\begin{aligned}
 R_t &\approx \frac{w^2}{2\delta S} \approx \frac{\pi k w^2}{\beta P_a} \\
 &\approx 1 \text{ km} \left(\frac{k}{1.4 \text{ W/m/}^\circ\text{K}} \right) \left(\frac{10 \text{ ppm/}^\circ\text{K}}{\beta} \right) \left(\frac{w}{5 \text{ cm}} \right)^2 \left(\frac{1 \text{ W}}{P_a} \right)
 \end{aligned} \tag{2.3}$$

and the new curvature R is calculated through the relation:

$$\frac{1}{R} = \frac{1}{R_0} + \frac{1}{R_t}$$

where R_0 is the affected optic's original effective radius of curvature. This relation will, later in this chapter, help us to determine the stability of a thermally distorted cavity (recall the stability criterion discussed in §1.2.2).

An often useful way to picture the effect of a thermal distortion in a cavity is as a source of power loss for light resonating in a single spatial mode of that cavity. Cavity light incident on the distortion will have some component scattered into higher order modes which are not resonant in the cavity and thus promptly exit. Consider a cavity field E_0 is in terms of the space of transverse electromagnetic field (TEM) modes, discussed in §1.2.2. If E_0 is subjected to the phase distortion $\phi(x, y)$, then the field immediately afterwards can be written:

$$E = E_0 e^{i\phi(x,y)}$$

and the fractional amplitude remaining in the mode of the original field E_0 is:

$$\mathcal{A} = \frac{\langle E_0 | E \rangle}{\langle E_0 | E_0 \rangle} = \frac{\langle E_0 | e^{i\phi(x,y)} | E_0 \rangle}{\langle E_0 | E_0 \rangle} = \frac{\iint E_0^*(x, y) E_0(x, y) e^{i\phi(x,y)} dx dy}{\iint E_0^*(x, y) E_0(x, y) dx dy} \quad (2.4)$$

and the fractional power scattered out of the original mode is simply:

$$\mathcal{S} \equiv 1 - \mathcal{A}^* \mathcal{A}.$$

A phase distortion in a cavity, in effect, acts to scatter power out of the fundamental mode, and thus out of the cavity, and so can be viewed as a simple loss term. This additional loss will reduce the resonant intensity gain of the cavity, by virtue of equation 1.11. In the case of thermal distortions, which increase with the amount of power pumped into the cavity, it is clear that the resonant intensity gain will decrease with increased power. Assuming that E_0 is a well collimated Gaussian beam (as is

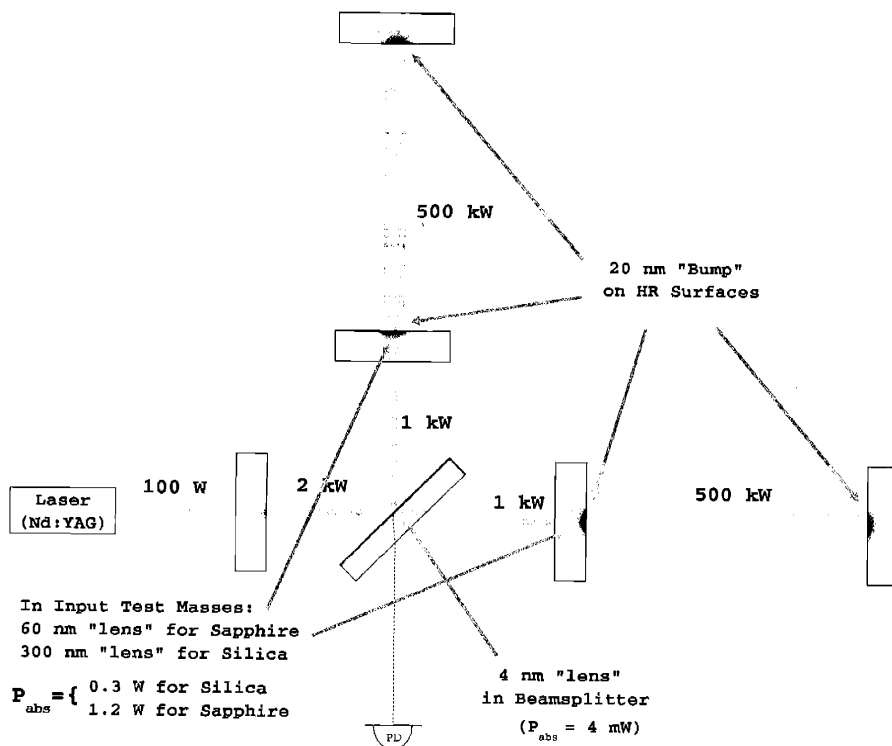


Figure 2-2: Mirror heating in advanced LIGO.

approximately true for the large diameter beams throughout all of LIGO):

$$E_0(x, y) = \sqrt{\frac{2}{\pi w^2}} e^{-\frac{x^2+y^2}{w^2}} \quad (2.5)$$

and inserting $\phi(r) \approx \frac{2\pi}{\lambda} \frac{r^2}{2R_t}$ into equation 2.4, we find that the fractional power scattered into higher order modes upon single pass through the cylindrically symmetric thermal distortion $\phi(r)$ is approximately:

$$\mathcal{S} = \frac{\left(\frac{w^2\pi}{2\lambda R_t}\right)^2}{1 + \left(\frac{w^2\pi}{2\lambda R_t}\right)^2} = \frac{\left(\frac{\beta P_a}{2\lambda k}\right)^2}{1 + \left(\frac{\beta P_a}{2\lambda k}\right)^2} \approx \left(\frac{\beta P_a}{2\lambda k}\right)^2 \quad \text{for } \frac{\beta P_a}{2\lambda k} \ll 1 \quad (2.6)$$

$$\approx 10 \text{ ppm} \left(\frac{1.4 \text{ W/m/}^\circ\text{K}}{k}\right)^2 \left(\frac{\beta}{10 \text{ ppm/}^\circ\text{K}}\right)^2 \left(\frac{1 \mu\text{m}}{\lambda}\right)^2 \left(\frac{P_a}{\text{mW}}\right)^2.$$

Figure 2-2 diagrams the anticipated distortions at the nominal operating point of advanced LIGO, calculated from equation 2.2. Now, how can these distortions

impact the sensitivity of the instrument? Recalling the expression for the instrument’s sensitivity from §1.2.2:

$$\delta h_{\text{shot}}(f) \approx \left(\sqrt{\frac{3}{2}} \sqrt{1 + 8 \frac{P_{\text{DP}}^{(c)}}{P_{\text{PRC}}^{(sb0)}}} \right) \times \frac{1}{2L\mathbb{G}_{\text{ARM}}} \frac{\lambda}{2\pi} \sqrt{\frac{2\hbar c}{\lambda} \frac{1}{P_{\text{PRC}}^{(c)}}} \quad (2.7)$$

we note that there are four possible ways by which wavefront distortions can decrease the gravitational wave strain sensitivity of the instrument:

1. **Decrease the arm cavity gain for the carrier.** The fewer “bounces” that light undergoes in the arm cavity means less sensitivity to the arm length. Recall the \mathbb{G}_{ARM} term in equation 2.7.
2. **Decrease the power recycling cavity gain for the carrier.** Fewer photons at the beamsplitter means less phase sensitivity. Recall the $P_{\text{PRC}}^{(c)}$ term in equation 2.7.
3. **Decrease the power recycling cavity gain for the sidebands.** Less power in the sidebands degrades the ability to readout the gravitational wave signal. Recall the $P_{\text{PRC}}^{(sb0)}$ term in equation 2.7.
4. **Increase the total carrier power leaking out of the dark port.** (i.e., increase the contrast defect). This light carries no information about the differential lengths of the instrument, but contributes shot noise nonetheless. Recall the $P_{\text{DP}}^{(c)}$ term in equation 2.7.

The following subsections analyze each of these potential pitfalls in detail. The first 3 in the list will be determined primarily by thermal distortions that are common between the arms, so a simplified common mode analysis is performed neglecting the beamsplitter. We’ll find that common mode thermal distortions cause item 3 to suffer the most, while item 1 is slightly affected and item 2 is hardly affected at all. Differential thermal distortions will have the most profound effect on item 4, and are analyzed in their own section. We’ll find that the thermal distortion in the beamsplitter isn’t significant enough to surpass static optical distortions that we can

expect, although a significant imbalance in absorption between the optical elements in the arms can serve to greatly increase the power leaking out of the dark port. We'll finish this section by calculating the effect on strain sensitivity versus power circulating in the power recycling cavity, and find the level of circulating optical power where thermal distortions prevent any further gain in strain sensitivity.

2.1.1 Common Mode Thermal Effects

Recall figure 1-3 which diagrams a simple overcoupled optical cavity (i.e., an arm cavity), and suppose that some small phase distortion $\phi(x, y)$ exists for transmission through the input coupler. For the carrier, which is resonant in the arm cavity, the reflected field may be written:

$$\begin{aligned} E_r^{(c)} &\approx r_1 e^{2i\phi(x,y)} E_i^{(c)} + it_1 e^{i\phi(x,y)} E_c^{(c)} \\ &\approx r_1 (1 + 2i\phi(x, y)) E_i^{(c)} - \frac{1 - r_1^2}{1 - r_1} (1 + i\phi(x, y)) E_i^{(c)} \\ &\approx -E_i^{(c)} \end{aligned}$$

where we have assumed $r_1 \sim 1$ and the input spatial mode E_i matches the cavity's fundamental mode E_c . The weak transmissive phase distortion in the input coupler is thus essentially invisible to the carrier. The antiresonant sidebands see the cavity as a simple mirror, however:

$$E_r^{(sb)} \approx e^{2i\phi(x,y)} E_i^{(sb)}$$

and are thus scattered into spatial modes other than that of the carrier.

The shape of the arm cavity surfaces (thermal distortions included) will serve to define the fundamental spatial mode of the interferometer, which will also determine the carrier mode resonant in the power recycling cavity. According to the simple analysis above, transmissive distortions in the input test mass will have little effect on the carrier in the power recycling cavity, thus the carrier gain in the power recycling cavity should remain unaffected. The sidebands, however, will fully experience the

transmissive thermal distortion in the input test mass, thus the power recycling gain for the sidebands should suffer a great deal. The three following sections examine in detail the carrier's behavior in the arms, its behavior in the power recycling cavity, and the sidebands' behavior in the power recycling cavity.

The carrier in the arm cavity




Figure 2-3: Thermal effects in an arm cavity.

First, let's examine the effects of thermal surface distortions on the light circulating in a single arm cavity, as is diagrammed in figure 2-3. The surfaces of the mirrors, originally concave with the same radius of curvature R_0 , will thermally expand and thus flatten. As the mirror surfaces become flatter, the lowest order Gaussian beam which is resonant in the cavity (recall §1.2.2) will become larger in diameter, which eventually becomes infinite in the limit where the test mass faces become parallel over the entire beam (i.e., the arm cavity becomes unstable). In other words, if R_t^{ITM} and R_t^{ETM} are the thermal radii of curvature changes in the input test mass and end test mass, the arm cavity becomes unstable approximately when:

$$\frac{1}{R_0} - \frac{1}{R_t^{ETM}} = - \left(\frac{1}{R_0} - \frac{1}{R_t^{ITM}} \right)$$

Arm Cavity Failure Points		
	Advanced LIGO Sapphire	Advanced LIGO Silica
$P_a^{\text{ITM}} + P_a^{\text{ETM}}$	2.6 W	1.1 W

Table 2.3: The total carrier power absorbed in both arm cavity test masses which causes the arm cavity to become unstable.

which, utilizing equation 2.3, is true when:

$$\boxed{\frac{P_a^{\text{ETM}} + P_a^{\text{ITM}}}{2} = \frac{\pi k w^2}{\beta R_0}} \quad (2.8)$$

where $P_a^{\text{ITM}} + P_a^{\text{ETM}}$ is the sum of the power absorbed in the input test mass and end test mass. Table 2.3 lists the mean test mass absorbed power where the arm cavity becomes unstable. Fortunately, this level of absorption is not exceeded in any advanced LIGO test mass (although it is very close in sapphire ITM's), and table A.6 shows the estimated effects of beam heating on the Gaussian mode resonant in LIGO arm cavities.

Although the resonant spatial mode will change slightly in the arms, the reflectivities of the mirror do not. Thus, we expect that the resonant intensity gain of the cavity (i.e., the mean number of bounces the trapped beam undergoes before exiting) will remain unchanged. However, we must consider how well the beam pumping the cavity (the power recycling cavity mode) matches the new cavity mode, as the only component of the input beam which will drive the arm cavity mode is that which matches the arm cavity mode. The effect on arm cavity gain will be seen when we write the carrier field in the power recycling cavity in terms of the field in the arm cavity in the following section.

The carrier in the power recycling cavity

To calculate how the arm cavity spatial mode affects the power recycling cavity spatial mode for the carrier, we'll assume that there exists a stable Gaussian mode which is resonant in the arm cavity, and then work backwards. Consider the optical

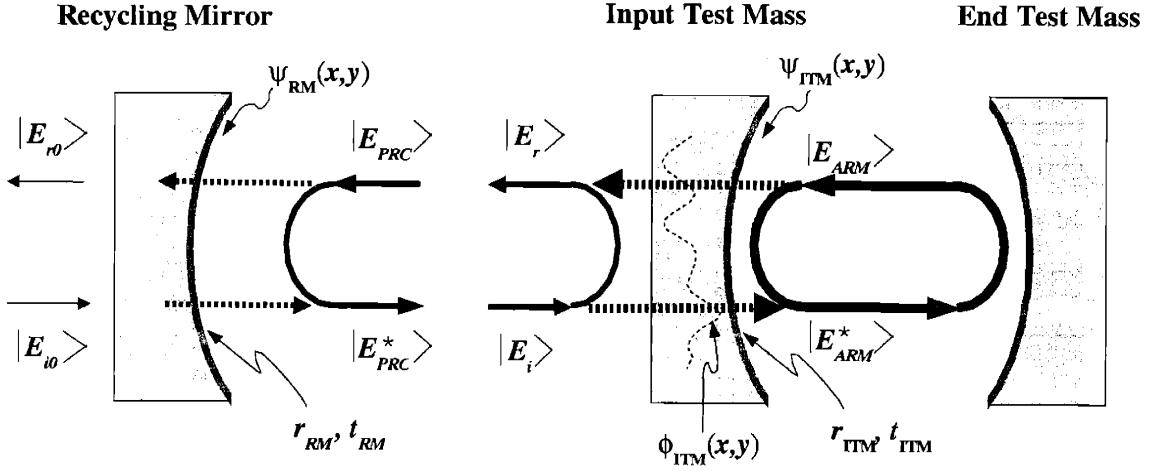


Figure 2-4: Electric fields in the Power Recycling Cavity.

fields around an imperfect Input Test Mass, diagrammed in figure 2-4, with total transmissive distortion ϕ_{ITM} , total surface curvature $\psi_{ITM}(x, y)$ (including both the nominal cold curvature, $\psi_{ITM}^{(0)}$ as well as the thermal change of curvature $\psi_{ITM}^{(t)}$), and nominal refractive index n . The carrier field reflected can be expressed as:

$$|E_r^{(c)}\rangle = \underbrace{r_{ITM} e^{2i\phi_{ITM}} e^{2in\psi_{ITM}^{(0)}} |E_i^{(c)}\rangle}_{\text{Initially reflected field}} + \underbrace{i \frac{t_{ITM}}{r_{ITM}} e^{i\phi_{ITM}} e^{in\psi_{ITM}^{(0)}} |E_{ARM}\rangle}_{\text{Cavity "leakage" field}}$$

where t_{ITM} is the test mass transmissivity and r_{ITM} is the input test mass reflectivity. Solving in terms of the arm cavity field:

$$i \frac{t_{ITM}}{r_{ITM}} |E_{ARM}\rangle = e^{-i\phi_{ITM}} e^{-in\psi_{ITM}^{(0)}} |E_r^{(c)}\rangle - r e^{i\phi_{ITM}} e^{in\psi_{ITM}^{(0)}} |E_i^{(c)}\rangle$$

If we assume that $|E_i^{(c)}\rangle$ and $|E_r^{(c)}\rangle$ are stable modes of the relatively very short power recycling cavity, we can propagate $|E_r^{(c)}\rangle$ off of the recycling mirror to attain $|E_i^{(c)}\rangle$:

$$|E_i^{(c)}\rangle = -e^{-2i\psi_{RM}} |E_r^{(c)}\rangle$$

G_{ARM} 10% Falloff Points		
	Sapphire	Fused Silica
P_a^{ITM}	1.3 W	0.070 W

Table 2.4: The amount of carrier power absorbed in an input test mass where the arm cavity gain decreases by 10%.

where ψ_{RM} is the nominal cold curvature of the recycling mirror. Assuming $(1 - r_{\text{ITM}}) \ll 1$, the reflected mode can now be written in terms of the cavity mode:

$$|E_r^{(c)}\rangle = i \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{e^{i\psi_{\text{RM}}}}{\cos((\psi_{\text{RM}} - n\psi_{\text{ITM}}^{(0)}) - \phi_{\text{ITM}})} |E_{\text{ARM}}\rangle.$$

If the instrument is designed such that the optics are cold curvature optimized, then $\psi_{\text{RM}} = n\psi_{\text{ITM}}^{(0)}$ and we arrive at the slightly simpler expression:

$$\boxed{|E_r^{(c)}\rangle = i \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{e^{in\psi_{\text{ITM}}^{(0)}}}{\cos(\phi_{\text{ITM}})} |E_{\text{ARM}}\rangle.} \quad (2.9)$$

From this we see that, in the absence of distortions, the beam resonant in the power recycling cavity is exactly the beam in the arm cavity after it accrues an additional amount of phase from passing through the curved Input Test Mass (a very weak negative lens) and the resonant intensity gain remains the same. For small thermal distortions, the curvature of the reflected carrier field exactly matches that of the arm cavity, the intensity profile only weakly changes, and the resonant intensity gain decreases very slightly. As the thermal distortion becomes larger, the field in the power recycling cavity begins to mismatch the arm cavity mode, and less power couples into the arm cavity (equation 2.11). Table 2.4 shows the powers absorbed in an ITM where the arm cavity gain decreases by 10%, as calculated using equations 2.6 and 2.11. As absorbed power increases further, and the net transmissive thermal curvature begins to approach $\pi/2$ in the optical aperture, interference fringes between the cavity mode and the power recycling mode begin to appear, so $|E_{\text{PRC}}\rangle$ must have a very large amplitude to counteract this interference and match the arm cavity field and equation 2.9 fails. Thus, there exist no stable Gaussian mode in

Carrier Failure Points		
Input Test Mass	Sapphire	Fused Silica
$\max_{(c)} P_a^{ITM}$	4.5^\dagger W	0.24 W

Table 2.5: The amount of carrier power absorbed in an input test mass where the thermal distortion exceeds $\lambda/4$, where no stable Gaussian mode exists in the power recycling cavity. Note that, from equation 2.10, these numbers are dependent on intrinsic material parameters *only*.

† Recall from table 2.3 that the arm cavity will fail before this level of absorption can even be reached in a Sapphire Advanced LIGO.

the power recycling cavity when the net thermal curvature approaches $\pi/2$ in the optical aperture. Utilizing equation 2.2, we find that the magnitude of the thermal distortion reaches this level, and hence the carrier's behavior begins to significantly degrade, when the total power absorbed in an input test mass is:

$$\boxed{\max_{(c)} P_a^{ITM} = \frac{2\pi k \lambda}{\beta} \frac{1}{4} \approx 220 \text{ mW} \left(\frac{k}{1.4 \text{ W/m/}^\circ\text{K}} \right) \left(\frac{10 \text{ ppm/}^\circ\text{K}}{\beta} \right) \left(\frac{\lambda}{1 \mu\text{m}} \right)} \quad (2.10)$$

and table 2.5 lists this value for sapphire and fused silica test masses.

Finally, equation 2.9 allows us to directly obtain the arm cavity gain (i.e., the amount of light resonating in the arm divided by the amount of light resonating in the power recycling cavity):

$$\boxed{\mathbb{G}_{\text{ARM}} \equiv \frac{P_{\text{ARM}}^{(c)}}{P_{\text{PRC}}^{(c)}} = \frac{\langle E_{\text{ARM}} | E_{\text{ARM}} \rangle}{\langle E_r^{(c)} | E_r^{(c)} \rangle} \approx \frac{4r_{\text{ITM}}^2}{t_{\text{ITM}}^2} \left(1 - \frac{\langle E_{\text{ARM}} | \phi_{\text{ITM}}^2 | E_{\text{ARM}} \rangle}{\langle E_{\text{ARM}} | E_{\text{ARM}} \rangle} \right)} \quad (2.11)$$

$$\approx \frac{4r_{\text{ITM}}^2}{t_{\text{ITM}}^2} (1 - 2\mathcal{S}_{\text{ITM}})$$

where we have again approximated the thermal distortion as purely spherical, and \mathcal{S}_{ITM} is the fractional power scattered by the ITM thermal distortion out of the cavity mode given by equation 2.6.

The sidebands in the power recycling cavity

In the case of the sidebands whose frequencies are chosen precisely to *not* couple into the arm cavities (i.e., they are anti-resonant), there is no arm cavity mode to couple to, and the ITM looks like a perfect mirror (or very nearly so). The power recycling cavity mode for the sidebands is thus defined by the curvature of the recycling mirror and the input test mass alone. The recycling mirror is made only slightly less concave by thermoelastic deformation, as the dominant absorption is on the surface here and the thermal expansion coefficient for fused silica is very small. The effective curvature of the input test mass will change much more, due to the fact that the power absorbed here is $> 100\times$ larger due to coating absorption of the large arm cavity power as well as substrate absorption of power recycling cavity light. Also, the effective curvature of the input test mass, which includes thermal lensing in addition to the thermoelastic distortion of the arm cavity surface, will become *more concave* as seen from the power recycling cavity. Thus, in contrast to the carrier mode, which becomes larger in the arms and is matched in the PRC, the mode for the sidebands will become *smaller*. Table A.7 shows the effects in the power recycling cavity.

Since the only component of the sidebands that concerns us is that which overlaps the carrier, which is in the mode defined by the cold power recycling cavity, we'll assume that the mirror curvatures don't change and consider the thermal distortions in terms of power scattered out of the cavity's fundamental mode, as in equation 2.6. Each time the sidebands reflect off one of the end mirrors, a fraction of the field's power is scattered into higher order modes which we do not need to consider further as these modes have no impact on the detected phase sensitivity. Once again, refer to figure 2-4. The sideband field leaving the recycling mirror in the power recycling cavity can be written:

$$|E_{\text{PRC}}^{*(sb0)}\rangle = it_{\text{RM}}e^{in\psi_{\text{RM}}}|E_{i0}^{(sb0)}\rangle + r_{\text{RM}}\mathcal{A}_{\text{ITM}}^2 \underbrace{e^{-2i\psi_{\text{RM}}}|E_{\text{PRC}}^{(sb0)}\rangle}_{=|E_{\text{PRC}}^{*(sb)}\rangle}$$

where \mathcal{A}_{ITM} is the fractional amplitude remaining in the cold cavity mode after trans-

mitting once through the distorted input test mass substrate (which passes through the substrate twice, hence is squared):

$$\mathcal{A}_{\text{ITM}} \equiv \frac{\langle E_{\text{PRC}} | e^{i\phi_{\text{ITM}}} | E_{\text{PRC}} \rangle}{\langle E_{\text{PRC}} | E_{\text{PRC}} \rangle} \approx \frac{1}{1 - i\frac{\pi}{\lambda} \frac{w^2}{2R}} \approx \frac{1}{1 - i\sqrt{\mathcal{S}_{\text{ITM}}}}$$

where we have written in terms of the fractional power scattered \mathcal{S}_{ITM} as defined in equation 2.6. Writing the cavity field in terms of the input field:

$$|E_{\text{PRC}}^{*(sb)}\rangle = i \frac{t_{\text{RM}}}{1 - r_{\text{RM}} \mathcal{A}_{\text{ITM}}^2} e^{in\psi_{\text{RM}}} |E_{i0}^{(sb0)}\rangle$$

we arrive at the resonant intensity gain for the cold cavity mode of the sidebands:

$$\begin{aligned} \mathbb{G}_{\text{PRC}}^{(sb)} &\equiv \frac{P_{\text{PRC}}^{(sb0)}}{P_{i0}^{(sb0)}} = \frac{t_{\text{RM}}^2}{|1 - r_{\text{RM}} \mathcal{A}_{\text{ITM}}^2|^2} \approx \frac{t_{\text{RM}}^2}{(1 - r_{\text{RM}})^2 + r_{\text{RM}} \mathcal{S}_{\text{ITM}}} \\ &= \mathbb{G}_{\text{PRC}}^{(c)} \left(\frac{1}{1 + \frac{r_{\text{RM}} \mathcal{S}_{\text{ITM}}}{(1 - r_{\text{RM}})^2}} \right). \end{aligned} \quad (2.12)$$

As the carrier power in the interferometer increases, the “thermal scattering” of the sidebands off of the input test mass increases, thus decreasing the resonant intensity gain for the sidebands in the power recycling cavity. Since \mathcal{S}_{ITM} increases with absorbed carrier power squared, we expect there to exist a critical value of \mathcal{S}_{ITM} where the amount of sideband power stored in the interferometer reaches the maximum attainable. To find this, we first note that the input sideband power is proportional to the input carrier power, which is proportional to the power absorbed in the ITM (since we saw before that the resonant intensity gain for the carrier does not appreciably change until thermal distortions in the input test mass prevent the interferometer from functioning at all), whose square is proportional to the thermal scatter \mathcal{S}_{ITM} for the sidebands (see equation 2.6), i.e.:

$$P_i^{(sb0)} \propto P_i^{(c)} \propto P_a \propto \sqrt{\mathcal{S}_{\text{ITM}}}.$$

Sideband Failure Points			
	Initial LIGO	Advanced LIGO Sapphire	Advanced LIGO Silica
$\max_{(sb)} P_a^{ITM}$	0.005 W	0.17 W	0.010 W

Table 2.6: The amount of carrier power absorbed in an input test mass where the stored sideband power reaches its maximum.

and we have:

$$P_{\text{PRC}}^{(sb0)} \propto \frac{\sqrt{\mathcal{S}_{\text{ITM}}}}{(1 - r_{\text{RM}})^2 + r_{\text{RM}} \mathcal{S}_{\text{ITM}}}.$$

Differentiating $P_{\text{PRC}}^{(sb0)}$ with respect to \mathcal{S}_{ITM} , and setting the result to zero, we find that the sideband power stored in the power recycling cavity reaches its maximum value when the thermal scatter for through the input test mass $\max_{(sb)} \mathcal{S}_{\text{ITM}}$ reaches the value:

$$\max_{(sb)} \mathcal{S}_{\text{ITM}} = \frac{(1 - r_{\text{RM}})^2}{r_{\text{RM}}} \quad (2.13)$$

which is about a 0.1% loss for the Advanced LIGO value of $r_{\text{RM}}^2 = 0.94$. Also, inserting 2.13 this into equation 2.12, we find that the sideband power recycling gain \mathbb{G}_{ITM} for $\mathcal{S}_{\text{ITM}} = \max_{(sb)} \mathcal{S}_{\text{ITM}}$ is one-half the nominal sideband power recycling gain. In terms of the power absorbed in the input test mass, utilizing equation 2.6, the sideband power stored in the power recycling cavity reaches its maximum when the carrier power absorbed in the input test mass is:

$$\boxed{\begin{aligned} \max_{(sb)} P_a^{ITM} &\approx \frac{2\lambda k}{\beta} (1 - r_{\text{RM}}) \\ &\approx 8.5 \text{ mW} \left(\frac{k}{1.4 \text{ W/m/}^\circ \text{ K}} \right) \left(\frac{\lambda}{1 \mu\text{m}} \right) \left(\frac{10 \text{ ppm}}{\beta} \right) \left(\frac{1 - r_{\text{RM}}}{0.03} \right) \end{aligned}} \quad (2.14)$$

the value of which is listed in table 2.6 for Initial LIGO and Advanced LIGO.

Summary

So far, we have seen how the sideband and carrier resonant intensity gains are influenced in both arm and power recycling cavities by thermal distortions, the most

significant of which is the transmissive distortion in the input test mass. Except for a decrease in the coupling efficiency between the arm and power recycling cavities, the carrier field remains relatively unaffected until either (1) the surfaces of the arm cavity distort past the point of stability (recall table 2.3), or (2) the peak-to-valley magnitude of the transmissive input test mass distortion approaches one-quarter wave (recall table 2.5). For Sapphire Advanced LIGO, condition (1) is reached first, while condition (2) is what limits a silica Advanced LIGO. The sidebands, since they do not couple into the arm cavities, are seriously affected by thermal distortions in the input test mass, and the total sideband power stored in a mode overlapping the carrier begins to decrease when the total power absorbed in the input test mass exceeds that given in equation 2.14.

To view all of this in terms of something more realizable in terms of the state of the instrument, mainly in terms of the carrier power circulating in the Power Recycling Cavity $P_{\text{PRC}}^{(c)}$, recall that the arm and power recycling cavity gains remain approximately constant for the carrier, and note that the power absorbed in any core optic is directly proportional to the carrier power in the power recycling cavity:

$$\begin{aligned}
P_a^{\text{ITM}} &= \frac{(2h_{\text{ITM}}a_s^{\text{ITM}} + \mathbb{G}_{\text{ARM}}a_c^{\text{ITM}})}{2} P_{\text{PRC}}^{(c)} \equiv \mathbb{A}_{\text{ITM}} P_{\text{PRC}}^{(c)} & (2.15) \\
P_a^{\text{ETM}} &= \frac{\mathbb{G}_{\text{ARM}}a_c^{\text{ETM}}}{2} P_{\text{PRC}}^{(c)} \equiv \mathbb{A}_{\text{ETM}} P_{\text{PRC}}^{(c)} \\
P_a^{\text{BS}} &= \frac{(2h_{\text{BS}}a_s^{\text{BS}} + 3a_c^{\text{BS}})}{2} P_{\text{PRC}}^{(c)} \equiv \mathbb{A}_{\text{BS}} P_{\text{PRC}}^{(c)}
\end{aligned}$$

where a_c represents the nominal coating absorption (typically about 0.5 ppm), a_s represents the nominal substrate absorption (typically about 1 ppm/cm), and h is the optic's thickness (typically about 10 cm). The constants \mathbb{A} are simply the power absorbed in each optic per unit carrier power circulating in the Power Recycling Cavity, and will allow us to replace the different P_a 's we frequently see with $P_{\text{PRC}}^{(c)}$. Table 2.7 shows the calculated values of the "power absorption parameters" \mathbb{A} for Initial LIGO and Advanced LIGO test masses and beamsplitters.

Utilizing these constants, we estimate in table 2.8 the power recycling cavity powers where Initial LIGO and Advanced LIGO experience sideband failure, arm

Power Absorption Parameters				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
\mathbb{A}_{ITM}	70	600	200	ppm
\mathbb{A}_{ETM}	33	200	200	ppm
\mathbb{A}_{BS}	20	3	3	ppm

Table 2.7: Total power absorbed per unit power at the beamsplitter of the Power Recycling Cavity, utilizing the cavity gains listed in table A.5 and absorptions listed in A.2 and A.1

cavity gain falloff, and carrier failure compared to the instrument’s nominal operating point. All instruments experience the relatively early sideband failure, and both Advanced LIGO instruments will experience some degree of arm cavity gain falloff. A Sapphire Advanced LIGO stays well below the quarter-wave ITM transmissive distortion limit (by virtue of its large thermal conductivity), but does come within $\sim 30\%$ of the absorption level where its arms become unstable. Only a Silica Advanced LIGO appears to exceed the carrier failure point, where the transmissive optical path distortion exceeds the quarter-wave limit in the optical aperture (the arm cavities don’t become unstable until twice this power is absorbed, due to the large magnitude of the transmissive thermal distortion compared to the reflective one in fused silica).

Critical Events						
	Initial LIGO		Advanced LIGO Sapphire		Advanced LIGO Silica	
(Watts)	P_a^{ITM}	$P_{PRC}^{(c)}$	P_a^{ITM}	$P_{PRC}^{(c)}$	P_a^{ITM}	$P_{PRC}^{(c)}$
Nominal Operation	0.028	400	1.20	2100	0.270	1300
Sideband Failure	0.005	70	0.17	280	0.010	50
Arm Cavity Gain Falloff	0.070	1000	1.30	2200	0.070	350
Carrier Failure	0.240	3400	2.00	3300	0.240	1200

Table 2.8: Summary and comparison of critical events with increasing power absorption, utilizing the instrument and material parameters listed in Appendix A.

2.1.2 Differential Mode Thermal Effects

We now turn our focus to differential thermal effects which, most significantly, prevent perfect interference at the beamsplitter thus causing carrier power to leak out of the dark port.




Figure 2-5: Electric fields at the Beam Splitter.

When considering both arms in tandem, the problem at hand becomes slightly more complicated, as the X and Y arm cavities can conceivably be different. Consider the electric fields around the beamsplitter, detailed in figure 2-5. The distortion of the beamsplitter will have an effect on the X path but little effect on the Y, and the macroscopic arm length asymmetry itself will prevent perfect interference at the beamsplitter (since one arm's return beam propagates slightly farther than the other arm's, the return beams' curvatures will be slightly mismatched). In reality, the X and Y cavities will have slightly different modes and gains, as the input test masses will not have exactly the same curvature and their coatings will not have exactly the same reflectivities. Also, the beamsplitter's reflective surface may not be

exactly 50/50, thus reflecting slightly more light into one arm than the other (more light in one arm means a larger thermal distortion in that arm). To account for all of this in the analysis, first assume that the arm cavity modes are identical except for a slightly different amplitude and curvature, i.e.,:

$$|E_{Y \text{ ARM}}\rangle = \left(1 - \sqrt{\frac{\delta P}{P}}\right) e^{-i\frac{\delta\psi_{\text{ITM}}}{2}} |E_{\text{ARM}}\rangle$$

$$|E_{X \text{ ARM}}\rangle = \left(1 + \sqrt{\frac{\delta P}{P}}\right) e^{i\frac{\delta\psi_{\text{ITM}}}{2}} |E_{\text{ARM}}\rangle.$$

Let ϕ_X and ϕ_Y denote the thermal phase distortion accrued between the beamsplitter surface and arm cavity surface for the respective arm, and define the common and differential thermal distortions, ϕ_c and ϕ_d as:

$$\phi_c = \frac{\phi_X + \phi_Y}{2} \quad \phi_d = \frac{\phi_X - \phi_Y}{2}.$$

If the differential distortion is small compared to the common distortion, then the field incident on the beamsplitter will be, by the analysis of the previous section:

$$|E_i^{(c)}\rangle = e^{-iK_c} \left(i \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{e^{-i\psi_{\text{RM}}}}{\cos(\phi_c)} \right) |E_{\text{ARM}}^*\rangle$$

where K_c represents the curvature change the field will experience over the common length from the beamsplitter to an input test mass, l_c . Note that the power at the beamsplitter is simply:

$$P_{\text{BS}}^{(c)} = \langle E_i^{(c)} | E_i^{(c)} \rangle = \frac{t_{\text{ITM}}^2}{4r_{\text{ITM}}^2} \langle E_{\text{ARM}} | \frac{1}{\cos^2(\phi_c)} | E_{\text{ARM}} \rangle \approx \frac{t_{\text{ITM}}^2}{4r_{\text{ITM}}^2} P_{\text{ARM}} (1 + 2\mathcal{S}_c)$$

where \mathcal{S}_c is the fractional TEM_{00} scatter due to the common mode distortion (approximately the average distortion in the ITM's). Immediately after the beamsplitter

surface, the outgoing fields are:

$$|E_{Yi}^{(c)}\rangle = r_{\text{BS}}|E_i^{(c)}\rangle \quad |E_{Xi}^{(c)}\rangle = it_{\text{BS}}|E_i^{(c)}\rangle$$

The carrier field returning from the Y arm to the vacuum-coating beamsplitter surface will be:

$$|E_{Yr}^{(c)}\rangle = r_{\text{BS}}e^{i(K_c - K_d - \delta\psi_{\text{ITM}} - 2\phi_d)} \left(1 - \sqrt{\frac{\delta P}{P}}\right) \left(i \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{e^{i\psi_{\text{RM}}}}{\cos(\phi_c)}\right) |E_{\text{ARM}}\rangle$$

where K_d is the curvature change the field has experienced by traveling the differential length l_d , and we have used the relation $\phi_Y - \phi_c = -\phi_d$. Similarly, the field returning to the substrate-coating beamsplitter surface from the X arm is:

$$|E_{Xr}^{(c)}\rangle = it_{\text{BS}}e^{i(K_c + K_d + \delta\psi_{\text{ITM}} + 2\phi_d)} \left(1 + \sqrt{\frac{\delta P}{P}}\right) \left(i \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{e^{i\psi_{\text{RM}}}}{\cos(\phi_c)}\right) |E_{\text{ARM}}\rangle$$

where the relation $\phi_X - \phi_c = \phi_d$ has been used. The field exiting the dark port is now:

$$\begin{aligned} |E_{\text{DP}}^{(c)}\rangle &= it_{\text{BS}}|E_{Yr}^{(c)}\rangle - r_{\text{BS}}|E_{Xr}^{(c)}\rangle \\ &= -2it_{\text{BS}}r_{\text{BS}}e^{i(\psi_{\text{RM}} + K_{lc})} \frac{t_{\text{ITM}}}{2r_{\text{ITM}}} \frac{\sin(K_d + \delta\psi_{\text{ITM}} + 2\phi_d)}{\cos(\phi_c)} |E_{\text{ARM}}\rangle \end{aligned} \quad (2.16)$$

and the total carrier power leaking out the dark port is:

$$\begin{aligned} P_{\text{DP}}^{(c)} &= \langle E_{\text{DP}}^{(c)} | E_{\text{DP}}^{(c)} \rangle \\ &\approx 4t_{\text{BS}}^2 r_{\text{BS}}^2 \frac{t_{\text{ITM}}^2}{4r_{\text{ITM}}^2} \langle E_{\text{ARM}} | \frac{\sin^2(K_d + \delta\psi_{\text{ITM}} + 2\phi_d)}{\cos^2(\phi_c)} | E_{\text{ARM}} \rangle \\ &\approx \langle E_0 | (K_d + \delta\psi_{\text{ITM}} + 2\phi_d)^2 | E_0 \rangle P_{\text{BS}}^{(c)} \end{aligned} \quad (2.17)$$

where we have assumed that $r_{\text{BS}}^2 = t_{\text{BS}}^2 = \frac{1}{2}$ and $|E_0\rangle$ is defined as:

$$|E_{\text{ARM}}\rangle \equiv \sqrt{P_{\text{ARM}}} |E_0\rangle.$$

We'll now examine each of the distortions that contribute in equation 2.17.

K_d This term is the curvature accrued by the probe beam by propagating the differential length l_d (typically about 0.2 meters). In general, when a TEM₀₀ field $|E_0\rangle$ of initial curvature R_0 propagates a distance $l_d \ll R_0$, the curvature of the field will change by some amount $\Delta R \ll R_0$, which can be expressed as:

$$\frac{1}{R} = \frac{1}{R_0 + \Delta R} \approx \frac{1}{R_0} \left(1 - \frac{\Delta R}{R_0} \right)$$

and we may write the propagated field as:

$$|E\rangle = e^{iK_d} |E_0\rangle$$

where
$$K_d \equiv -\frac{2\pi}{\lambda} \frac{r^2}{R_0} \left(\frac{\Delta R}{R_0} \right) \tag{2.18}$$

Since the detector operates exclusively within the Rayleigh range of the probe beam ($z_R \equiv \frac{\pi w_0^2}{\lambda}$, which is the distance a beam of waist w_0 stays well-collimated), the curvature change over a short distance d can be written as:

$$\Delta R \approx \frac{R_0^2}{z_R^2} l_d = \frac{R_0^2 \lambda^2}{\pi^2 w_0^4} l_d$$

and the empty space propagator over the beam waist w becomes:

$$K_d \approx -\frac{2\lambda l_d}{\pi w^2} \approx -0.05 \text{ mrad} \left(\frac{\lambda}{1 \mu\text{m}} \right) \left(\frac{5 \text{ cm}}{w} \right)^2 \left(\frac{l_d}{20 \text{ cm}} \right).$$

Recall also that the arm cavity field curvature will weaken slightly with input power, as the test mass surfaces heat and thus flatten. Thus, this term, which is very small to begin with, becomes even smaller as absorbed power squared (since R_0 changes inversely with absorbed power), and we will thus ignore this term further.

$\delta\psi_{\text{ITM}}$ This term represents the field mismatch due to static imperfections and curvature mismatch of mirrors between the X and Y arm. In Initial LIGO, the curvature

difference varies by $\sim 2\%$ from test mass to test mass [6]. Over the beam waist, the phase distortion is approximately:

$$\begin{aligned}\delta\psi_{\text{ITM}} &= \frac{2\pi}{\lambda} \frac{w^2}{R_X R_Y} (R_Y - R_X) \approx \frac{2\pi}{\lambda} \frac{w^2}{R_0^2} \Delta R_0 \\ &\approx 21 \text{ mrad} \left(\frac{1 \mu\text{m}}{\lambda} \right) \left(\frac{w}{5 \text{ cm}} \right)^2 \left(\frac{15 \text{ km}}{R_0} \right)^2 \left(\frac{\Delta R_0}{300 \text{ m}} \right)\end{aligned}$$

This term, which has a dependence on R_0 which is identical to the previous term, will also become slightly smaller with absorbed power squared.

ϕ_d This term represents the difference in the total thermal distortion experienced between the beamsplitter surface and the arm cavity surface for the X and Y paths. Nominally, this only includes the thermal distortion of the beamsplitter. In reality, a difference in the transmissive distortion in each test mass will exist, due to different (or inhomogeneous) absorption in the bulk and surface of the input test masses as well as a different amounts of power incident and circulating in each arm. In general, the differential thermal distortion can be written in terms of the beamsplitter thermal distortion and the differential test mass thermal distortion:

$$2\phi_d = \phi_{\text{BS}} + (\phi_{\text{ITM } X} - \phi_{\text{ITM } Y}) \equiv \phi_{\text{BS}} + \delta\phi_{\text{ITM}}.$$

The beamsplitter thermal distortion will be significantly smaller than the input test mass thermal distortion, as the beamsplitter does not suffer coating absorption of the high arm cavity powers. Utilizing the approximation 2.2, the phase distortion over the beam waist is approximately:

$$\phi_{\text{BS}} = \frac{2\pi}{\lambda} \frac{\beta}{2\pi k} P_a^{\text{BS}} \approx 7 \text{ mrad} \left(\frac{1 \mu\text{m}}{\lambda} \right) \left(\frac{\beta}{10 \text{ ppm}/^\circ\text{K}} \right) \left(\frac{1.4 \text{ W}/\text{m}/^\circ\text{K}}{k} \right) \left(\frac{P_a^{\text{BS}}}{1 \text{ mW}} \right).$$

Comparing to the previous paragraph discussing the static curvature mismatch term, $\delta\psi_{\text{ITM}}$, we see that the beamsplitter distortion surpasses the static test mass curvature mismatch at $\sim 3\text{mW}$ of total optical absorption (approximately the operating point for all three interferometers we're considering, see table A.7).

For the differential thermal distortion in the input test masses, there is a great deal of uncertainty. The beamsplitter's 50/50 coating is 50/50 to within a 1% tolerance, so the fractional power between the arms can vary by at least as much. In addition, Initial LIGO input test mass bulk absorption has been measured to vary by up to a factor of 2 from test mass to test mass (in particular, WA4K-X and WA4K-Y) [6], the absorption in sapphire could be extremely inhomogeneous [2], and the very small, yet very critical coating absorption on the arm cavity mirror surfaces may very well vary by a factor of 2. If the total optical absorption varies by a factor of 2 between test masses, the input test mass differential distortion over the beam waist is approximately:

$$\begin{aligned}\delta\phi_{\text{ITM}} &= \frac{2\pi}{\lambda} \frac{\beta}{2\pi k} \frac{\delta P_a^{\text{ITM}}}{P_a^{\text{ITM}}} P_a^{\text{ITM}} \\ &\approx 7 \text{ mrad} \left(\frac{1 \mu\text{m}}{\lambda} \right) \left(\frac{\beta}{10 \text{ ppm}/^\circ\text{K}} \right) \left(\frac{1.4 \text{ W/m}/^\circ\text{K}}{k} \right) \left(\frac{\delta P_a^{\text{ITM}}}{P_a^{\text{ITM}}} \right) \left(\frac{P_a^{\text{ITM}}}{1 \text{ mW}} \right)\end{aligned}$$

While absorption in the beamsplitter alone is not sufficient to eclipse typical static optical imperfections, only a small amount of additional differential optical absorption in the test masses is necessary to completely dominate the power scattered out of the dark port. As such, we'll disregard static defects in the rest of our toy analysis.

2.1.3 Thermal Effects on Sensitivity

The goal of this section is to utilize what has been found so far to determine strain sensitivity versus the carrier power at the beamsplitter $P_{\text{PRC}}^{(c)}$. We'll then use this relation to find the maximum operating power where improvements in sensitivity can still be realized (differentiate h_{shot} with respect to $P_{\text{PRC}}^{(c)}$ and set this equal to zero).

Finding $P_{\text{DP}}^{(c)}$ vs. $P_{\text{PRC}}^{(c)}$

First, to express the carrier power leaking out of the dark port in terms of the carrier power circulating in the recycling cavity, note that the differential thermal distortion

ϕ_d is proportional to the power in the instrument, so that we may write:

$$P_{\text{DP}}^{(c)} \approx \langle E_0 | \underbrace{\left(\frac{2\phi_d}{P_{\text{PRC}}^{(c)}} \right)^2}_{\equiv \mathcal{C}_d} | E_0 \rangle (P_{\text{PRC}}^{(c)})^3 \quad (2.19)$$

where \mathcal{C}_d , hereafter termed the “differential distortion parameter”, defined as

$$\mathcal{C}_d \equiv \langle E_0 | \left(\frac{2\phi_d}{P_{\text{PRC}}^{(c)}} \right)^2 | E_0 \rangle = \frac{2}{\pi w^2} \iint \left(\frac{2\phi_d(x, y)}{P_{\text{PRC}}^{(c)}} \right)^2 e^{-2\frac{x^2+y^2}{w^2}} dx dy$$

is independent of $P_{\text{PRC}}^{(c)}$, and is, in effect, a measure of the power scattered out of the interferometer’s main mode by differential distortions at unit power in the recycling cavity.

Finding $P_{\text{PRC}}^{(sb)}$ vs. $P_{\text{PRC}}^{(c)}$

To express the sideband power at the beamsplitter in terms of the total carrier power at the beamsplitter, recall the expression for the sideband behavior in the power recycling cavity (equation 2.12) and note that we can relate the sideband input field to the resonant carrier field:

$$\mathbb{G}_{\text{PRC}}^{(c)} P_{i0}^{(sb)} = \mathbb{G}_{\text{PRC}}^{(c)} \frac{\Gamma^2}{2} P_{i0}^{(c)} = \frac{\Gamma^2}{2} P_{\text{PRC}}^{(c)}$$

where Γ is the modulation depth of phase modulation on the input laser light. Finally, we can express the sideband power at the beamsplitter in terms of the carrier power in the recycling cavity:

$$P_{\text{PRC}}^{(sb)} = \frac{\frac{\Gamma^2}{2} P_{\text{PRC}}^{(c)}}{\left(1 + \frac{r_{\text{RM}} \mathcal{S}_c}{\underbrace{1 - r_{\text{RM}} (P_{\text{PRC}}^{(c)})^2}_{\equiv \mathcal{C}_c}} (P_{\text{PRC}}^{(c)})^2 \right)^2} \quad (2.20)$$

where the constant \mathcal{C}_c , defined as:

$$\mathcal{C}_c \equiv \frac{r_{\text{RM}}}{1 - r_{\text{RM}}} \frac{\mathcal{S}_c}{(P_{\text{PRC}}^{(c)})^2}$$

is hereafter termed the “common thermal distortion parameter”. This parameter is independent of $P_{\text{PRC}}^{(c)}$ and represents a measure of the distortion seen in an input test mass (the common distortion) at unit power in the recycling cavity.

Finding $\delta\tilde{\phi}_{\min}$ vs. $P_{\text{PRC}}^{(c)}$

We may now rewrite the equation for the strain sensitivity (equation 2.7) in terms of the carrier power in the power recycling cavity:

$$\delta h_{\text{shot}} \propto \sqrt{\frac{1}{P_{\text{PRC}}^{(c)}} + 16 \frac{\mathcal{C}_d}{\Gamma^2} P_{\text{PRC}}^{(c)} \left(1 + \mathcal{C}_c (P_{\text{PRC}}^{(c)})^2\right)^2}. \quad (2.21)$$

which will have a global minimum for some power $P_{\text{PRC}}^{(c)} > 0$. The equation for this minimum is found to be quite lengthy and unrevealing, so it is omitted here, although figure 2-6 shows the results of a numerical minimization of $P_{\text{PRC}}^{(c)}$ versus \mathcal{C}_d and \mathcal{C}_c .

The common distortion parameter is fairly straightforward to calculate from its definition, utilizing the approximation of pure thermal curvature (equation 2.3) and the power absorption parameters (equation 2.15):

$$\mathcal{C}_c \equiv \frac{r_{\text{RM}}}{1 - r_{\text{RM}}} \frac{\mathcal{S}_c}{(P_{\text{PRC}}^{(c)})^2} \approx \frac{r_{\text{RM}}}{1 - r_{\text{RM}}} \left(\frac{\beta_{\text{ITM}} \mathbb{A}_{\text{ITM}}}{2\lambda k_{\text{ITM}}} \right)^2.$$

The differential distortion parameter is more difficult, but if we assume a pure radius of curvature mismatch R_t :

$$\begin{aligned} \mathcal{C}_d &\equiv \langle E_0 | \left(\frac{2\phi_d}{P_{\text{PRC}}^{(c)}} \right)^2 | E_0 \rangle \approx \langle E_0 | \left(\frac{2\pi r^2}{\lambda R_t} \frac{1}{P_{\text{PRC}}^{(c)}} \right)^2 | E_0 \rangle \\ &\approx 8 \frac{\mathcal{S}_d}{(P_{\text{PRC}}^{(c)})^2} = 8 \left(\frac{\beta \mathbb{A}}{2k\lambda} \right)^2. \end{aligned}$$

where \mathcal{S}_d is the TEM_{00} scatter calculated from the net differential distortion ϕ_d . If the

beamsplitter is the only source of differential thermal curvature, then the parameters of the beamsplitter must be used. If the arm power asymmetry and/or input test mass absorption asymmetry dominate, then the parameters above must use input test mass parameters.




Figure 2-6: Maximum operating power versus differential and common distortion parameters. The blocks represent ranges of the common and differential distortion parameters for the configurations indicated, calculated via the approximations derived in this chapter and using the parameters listed in Appendix A. A modulation depth of $\Gamma = 0.5$ is assumed in the placement of the boxes.

Table 2.9 details the values of the distortion parameters for LIGO and Advanced LIGO, and figure 2-6 shows where these values lie in terms of the approximate maximum useful power at the beamsplitter. Initial LIGO is seen to be relatively immune to thermal effects, when considering that the nominal power at the beamsplitter does not exceed 0.4kW (in reality, however, static optical defects heavily dominate over dif-

ferential thermal defects at the relatively low powers of Initial LIGO, so this estimate is excessively optimistic). The common distortion parameter for fused silica increases between Initial LIGO and Advanced LIGO, by virtue of the fact that the arm cavity gain is increased significantly in Advanced LIGO, and is increased to the point that, even if the input test mass distortions are identical, lensing in the beamsplitter alone is enough to prevent the instrument from reaching its nominal operating point of 1.4 kW. Furthermore, even if the lens in the beamsplitter were small enough to permit operation at 1.4 kW, the distortion in a fused silica test mass reaches the $\lambda/4$ breakdown point at about 1.1 kW. Thus, thermal distortions are a serious problem in a silica Advanced LIGO. In a sapphire Advanced LIGO, however, the common distortion is significantly smaller due to its $20\times$ larger thermal conductivity, and the nominal operating point is well below that where the test mass distortion is expected to become near $\lambda/4$ (about 8 kW in the recycling cavity). Thus, thermal effects may only marginally effect a sapphire interferometer's performance. However, sapphire's bulk absorption is known to be inhomogeneous [2], the degree of which is currently unknown and could significantly increase the differential distortion and seriously affect the performance of the instrument, as indicated by the broken line and question mark in figure 2-6.

Distortion Parameters				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
\mathcal{C}_c	2.4×10^{-6}	3.3×10^{-7}	1.3×10^{-5}	$1/W^2$
\mathcal{C}_d (BS Only)	2.8×10^{-10}	4.7×10^{-12}	4.7×10^{-12}	$1/W^2$
\mathcal{C}_d (ITM Dominated)	1.9×10^{-9}	5.3×10^{-10}	2.1×10^{-8}	$1/W^2$

Table 2.9: Approximate distortion parameters for Initial LIGO and Advanced LIGO, calculated via the approximations derived in this chapter and using the parameters listed in Appendix A.

2.1.4 Summary

We have roughly examined the effects of thermal distortions on the performance of a power recycled Michelson interferometer with Fabry-Perot arms (i.e., LIGO's layout). Provided there exists a stable Gaussian mode in the arm cavities, large common distortions in the input test mass cause a significant loss in the sideband power stored in the power recycling cavity, while only slightly affecting the carrier until the wavefront distortion becomes so large (one quarter wave) that interference between the arm and the power recycling cavity prevents any stable spatial mode from existing in both cavities. Differential distortions in the beamsplitter and input test masses cause carrier power to couple out of the dark port of the interferometer, thus contributing shot noise yet yielding no information on the state of the interferometer. This additional carrier power, coupled with the loss of sideband power which is used to read out the interferometer's differential length signal, results in an increase in the observed phase noise at the dark port to the point that no further sensitivity is gained by increasing the power circulating in the instrument. The net effect on the performance of the instrument in the presence of thermal distortions can be well characterized by the common and differential distortion parameters \mathcal{C}_c and \mathcal{C}_d . While thermal distortions do not appreciably affect Initial LIGO, they will seriously hinder a fused silica Advanced LIGO, and a sapphire Advanced LIGO may run into significant problems for appreciable inhomogeneous absorption.

2.2 Thermally Adaptive Optics (The Solution)

To fix an absorption-induced wavefront distortion in the optics of a gravitational wave detector, where one cannot "touch" the core optics for fear of compromising the instrument's length sensitivity, we may use some external source of radiative heating to homogenize the absorption-induced temperature fields, thus flattening the initial wavefront distortion (see figure 2-7). Examples of potential external sources are a carefully shaped laser beam at a wavelength which is strongly absorbed in the optical material's surface or bulk, or a ring of resistive wire suspended away from the optic's




Figure 2-7: Thermally compensating absorption-induced thermal distortions.

face (i.e., a toaster). More generally, as localized heating is a mechanism by which the optical path length through an optic can be changed, one may make a passive optic adaptive through tailored radiative heating by a controlled external source. In principle, one can arbitrarily change the transmissive optical path through an optic by tailoring the external radiative heating pattern across the optic's face, which hereafter is termed the “compensation pattern” $I_c(x, y)$.

A very obvious means of compensating an absorption pattern $I_a(x, y)$ is to insulate the edge of the optic so heat does not escape radially, and then radiatively heat the optic with a compensation pattern $I_c(x, y)$ such that the total intensity absorbed is uniform over the entire optic, i.e.:

$$I_a(r) + I_c(r) = I_0$$

where I_0 is constant over all (x, y) . The compensation pattern with the least power input will be where $I_0 = \max_{(x,y)} I_a(x, y)$, so that:

$$I_c(x, y) = \max_{(x,y)} I_a(x, y) - I_a(x, y)$$

and the total power absorbed by the optic is $P_{tot} = \pi R^2 I_0$. All core optics in LIGO are suspended by thin wires in ultra-high vacuum, thus the only route for heat to escape is through blackbody radiation. Using the Stefan-Boltzmann law linearized for temperature increases small compared to ambient (more on this in the next chapter), the mean surface temperature increase $\overline{\Delta T}$ is related to the total power absorbed by:

$$P_{tot} = \pi R^2 I_0 = 4\epsilon\sigma T_\infty^3 \overline{\Delta T} (2\pi R^2)$$

where σ is the Stefan-Boltzmann constant, ϵ is the emissivity of the optic's surface, and T_∞ is the ambient temperature (the temperature of the walls of the vacuum enclosure). The mean surface temperature increase under ideal compensation of an absorber with maximum intensity $\max_{(x,y)} I_a(x,y)$ is found to be:

$$\overline{\Delta T} = \frac{\max_{(x,y)} I_a(x,y)}{8\epsilon\sigma T_\infty^3} \quad (2.22)$$

Consider the case where the absorber we wish to compensate is a Gaussian beam of waist radius w :

$$I_a(r) = \frac{2P_a}{\pi w^2} e^{-2\frac{r^2}{w^2}}$$

where P_a is the total power absorbed out of the beam. The ideal compensator's heating pattern will then take the form:

$$I_c(r) = \max I_a(r) - I_a(r) = \frac{2P_a}{\pi w^2} (1 - e^{-2\frac{r^2}{w^2}})$$

and the uniform intensity absorbed in the optic will be:

$$I_0 = \frac{2P_a}{\pi w^2}$$

Using equation 2.22, the mean surface temperature increase for ideal thermal compensation of power P_a absorbed by a Gaussian beam of waist radius w at room

temperature is then:

$$\overline{\Delta T} = \frac{P_a}{4\pi w^2 \epsilon \sigma T_\infty^3} \approx 21 \text{ }^\circ\text{K} \left(\frac{0.05 \text{ m}}{w} \right)^2 \left(\frac{1}{\epsilon} \right) \left(\frac{300 \text{ }^\circ\text{K}}{T_\infty} \right)^3 \left(\frac{P_a}{1 \text{ W}} \right) \quad (2.23)$$

which is completely independent of the compensated optic's size and material (the emissivity will be set by the optical coating). This equation represents a standard comparison for the efficiency of any particular means of thermal compensation.

This work examines two specific means of correcting absorption-induced wavefront distortions in the core optics of laser interferometric gravitational wave detectors:

1. **Heating Ring Thermal Compensation:** For uniform absorption of Gaussian beam, a heating ring with some strategically placed reflectors is designed to compensate the anticipated cylindrically symmetric distortion. The design and optimization of the heating ring compensator are discussed in chapter 4.
2. **Scanning Laser Thermal Compensation:** In the case where absorption is nonuniform and perhaps unknown, a scanning heating laser is used to tailor the compensation pattern to the thermal distortion. We'll use a Carbon Dioxide laser, which, at a wavelength of $10 \mu\text{m}$, is strongly absorbed in the surfaces of both fused silica and sapphire. This method is discussed in chapter 5.

2.3 Displacement Noise Induced by Thermal Actuation

By virtue of the fact that our thermal actuation scheme will actively alter the optical path length through different components of the interferometer, unchecked intensity fluctuations of the thermal actuator will couple directly into displacement noise. Here we discuss a simple method by which we can calculate displacement noise given the intensity noise in a thermal actuator.

Consider an optic illuminated by a nominally steady heating source with a power

ripple of amplitude $\mathcal{R}(\omega)$ at frequency ω :

$$I_c(x, y, \omega t) = I_c(x, y) (1 + \mathcal{R}(\omega)e^{-i\omega t})$$

and let a be the characteristic width of the pattern $I_c(x, y)$. In general, the optical path distortion through an optic is a linear function of the heater's intensity (we'll see this in the next chapter), so we can break the problem into two separate parts: (1) the steady source I_c acts alone, and (2) the oscillating component $I_c(x, y)\mathcal{R}(\omega)e^{-i\omega t}$ acts alone, and the full solution is simply the sum of the two. If the source has been shining on the optic for larger than the optic's full time constant, then the optical path distortion induced by the steady component does not change with time. The oscillating source, however, deposits no net energy into the optic, but rather adds and, in effect, subtracts energy with frequency ω . If the period of oscillation $\frac{2\pi}{\omega}$ is much smaller than the local time constant of the optic, $\tau_l \simeq \frac{a^2\rho c}{k}$ (where ρ is the optical material's density, c is its heat capacity, and k is the thermal conductivity, and is about 10s for $a = 1\text{cm}$ on sapphire), then radial diffusion of absorbed energy will be negligible (the net energy deposited in the first half of the source's cycle is removed in the second half quicker than the energy can diffuse radially). The energy per unit area contained in a column through the optic as a function of time is then given by:

$$E(x, y, \omega t) = \frac{-1}{i\omega} I_c(x, y)\mathcal{R}(\omega)e^{-i\omega t}$$

However, if the temperature field is known, the energy per unit area in the column can be calculated as:

$$E(x, y, \omega t) = \rho c \int_0^h T(x, y, z, t) dz$$

Equating the previous two equations, we get:

$$\int_0^h T(x, y, z, t) dz = \frac{i}{\rho c \omega} I_c(x, y)\mathcal{R}(\omega)e^{-i\omega t}$$

The optical path distortion induced by the ripple in the thermal actuation will be

approximately:

$$\delta x_t(x, y, \omega t) \simeq \frac{\beta i}{\rho c \omega} I_c(x, y) \mathcal{R}(\omega) e^{-i\omega t}$$

where β is defined as in equation 2.2 for reflected and transmitted beams.

Given the target sensitivity of the instrument and the surface of actuation (recycling cavity vs. arm cavity), this calculation puts an upper limit on the tolerable intensity noise of thermal actuators, $I_c(x, y) \mathcal{R}(\omega)$. If the surfaces of the arm cavity are actuated, then the measured change in arm length is simply the overlap of the the fluctuating optical path distortion δx_t with the intensity profile of the probe beam (there is no detection where there is no probe light):

$$\delta x_{\text{ARM}}(\omega t) = \frac{\alpha \mathcal{R}(\omega) i}{\rho c \omega} e^{i\omega t} \iint I_c(x, y) \frac{2}{\pi w^2} e^{-2\frac{x^2+y^2}{w^2}} dx dy \quad (2.24)$$

where w is the waist of the LIGO beam at the optic under actuation. If, instead, the surfaces within the recycling cavity are actuated, where the interferometer's sensitivity to length changes is diminished by the carrier gain of the arm cavity \mathbb{G}_{ARM} (~ 500), then intensity noise in the actuator will create a resulting noise in the measured arm length

$$\delta x_{\text{PRC}} = \frac{1}{\mathbb{G}_{\text{ARM}}} \delta x_{\text{ARM}}.$$

For example, consider the so-called “photo-thermal noise” [7], generated by shot noise on laser light absorbed (i.e. detected) on the arm cavity surface of a test mass. With an arm cavity power of P_{ARM} and a coating absorption of a_c , the total amount of absorbed (detected) power is $P_a = a_c P_{\text{ARM}}$, and the intensity fluctuation due to shot noise alone on the detected power P_a is given by:

$$\delta P_a = \sqrt{2\epsilon a_c P_{\text{ARM}}} = \sqrt{2\epsilon P_a}$$

where ϵ is the energy of a single arm cavity photon. So, plugging the intensity

fluctuation function given by:

$$I_c(x, y)\mathcal{R}(\omega) = \left(\frac{2P_a}{\pi w^2} e^{-2\frac{x^2+y^2}{w^2}} \right) \left(\frac{\delta P_a}{P_a} \right)$$

into equation (2.24), we get a measured amplitude of length fluctuation which is nearly identical to the rigorous solution [7]:

$$\begin{aligned} |\delta x_{\text{PRC}}| &= \frac{\alpha \sqrt{2\epsilon P_a}}{\pi w^2 \rho c \omega} \approx 7 \times 10^{-22} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{\alpha}{5 \text{ ppm}/^\circ\text{K}} \right) \left(\frac{\epsilon}{2 \times 10^{-19} \text{ J}} \right)^{\frac{1}{2}} \left(\frac{a_c}{0.5 \text{ ppm}} \right)^{\frac{1}{2}} \\ &\quad \times \left(\frac{P_{\text{ARM}}}{500 \text{ kW}} \right)^{\frac{1}{2}} \left(\frac{5 \text{ cm}}{w} \right)^2 \left(\frac{4 \text{ g/cm}^3}{\rho} \right) \left(\frac{770 \text{ J/kg}/^\circ\text{K}}{c} \right) \left(\frac{100 \text{ Hz}}{\omega} \right). \end{aligned}$$

Finally, consider the perfect compensation pattern discussed in the previous section, compensating absorbed power P_a in an input test mass and actuating on the power recycling cavity side of the mirror:

$$I_c(x, y) = \frac{2P_a}{\pi w^2} \left(1 - e^{-2\frac{x^2+y^2}{w^2}} \right).$$

The observed displacement noise induced by the compensator is then (by equation 2.24):

$$|\delta x_{\text{PRC}}| = \frac{1}{\mathbb{G}_{\text{ARM}}} \frac{\alpha \mathcal{R}(\omega) P_a}{\pi w^2 \rho c \omega}.$$

If we wish that the compensator contribute no more displacement noise than the ever-present photothermal noise in the arm cavity, the broadband ripple on the thermal actuator power cannot exceed:

$$\begin{aligned} \mathcal{R}(\omega) &\leq \mathbb{G}_{\text{ARM}} \frac{\sqrt{2\epsilon P_a}}{P_a} \\ &\approx 3 \times 10^{-7} / \sqrt{\text{Hz}} \left(\frac{\mathbb{G}_{\text{ARM}}}{500} \right) \left(\frac{\epsilon}{2 \times 10^{-19} \text{ J}} \right)^{\frac{1}{2}} \left(\frac{1 \text{ W}}{P_a} \right)^{\frac{1}{2}}. \end{aligned} \tag{2.25}$$

If we instead decide to actuate on the arm cavity surface of the test mass, then we simply set $\mathbb{G}_{\text{ARM}} = 1$ in this equation and attain a more stringent limit on the broadband ripple \mathcal{R} .

The Radiometer Effect

Another effect, which we will consider no further than the present discussion, is the “radiometer effect” in an extremely dilute gas, and pertains to a net force exerted on a disk in a near vacuum when one face of the disk is brought to a temperature higher than the opposite face (as will be the case for a thermally compensating a cylindrical optic suspended in LIGO’s vacuum enclosure). For moderate vacuums (where the mean free path of the molecules in the dilute gas is much smaller than the size of the vane, and is where novelty Crookes Radiometers operate), the dominant “pushing” comes from the hydrodynamic effect of thermal transpiration, caused by the flow of gas around the edges from the cooler face to the warmer face of the disk¹. The vacuum envelope of the LIGO interferometer nominally contains $\sim 10^{-9}$ torr of molecular hydrogen ($\sim 1\text{\AA}$ effective diameter) at room temperature (300 °K), leading to a mean free path of 10^6 m; hence, all hydrodynamic effects can be ignored here. In this case, the dominant “pushing” is due to adsorption of room temperature gas molecules on the warmer face which are quickly re-emitted at an elevated temperature².

Assuming that all molecules incident on the disk experience this adsorption and subsequent thermalization (yielding a “worst-case” estimate), the effective pressure on the warmer side of the disk can be written

$$P_{eff} = nk_B \frac{T_\infty + T}{2} = P + nk_B \frac{\Delta T}{2}$$

where P is the true pressure of the gas at ambient temperature T_∞ and ΔT is the temperature increase of the warmer face above ambient. The net force acting on the disk is then:

$$F = \pi R^2 (P_{eff} - P) = P\pi R^2 \frac{\Delta T}{2T_\infty}$$

where R is the radius of the disk. This net force consists of a large number of discrete, uncorrelated momentum “kicks” when each molecule leaves the face of the optic with

¹See, e.g., problem 5.6 of [17]

²See, e.g., problem 5.5 of [17]

mean energy $2k_B T$:

$$p_i \approx \sqrt{2k_B T m_{\text{H}_2}}.$$

As such, the fluctuation of the net force is Poissonian:

$$\begin{aligned} \delta F &= \sqrt{p_i F} \approx \sqrt{(2k_B T m_{\text{H}_2}) \left(P \pi R^2 \frac{\Delta T}{2T_\infty} \right)} \\ &\approx 10^{-17} \frac{\text{N}}{\sqrt{\text{Hz}}} \left(\frac{P}{10^{-9} \text{ torr}} \right) \left(\frac{R}{15 \text{ cm}} \right) \left(\frac{\Delta T}{1^\circ \text{K}} \right) \end{aligned}$$


and the resulting displacement noise as a function of angular frequency ω is:

$$\delta x(\omega) = \frac{\delta F}{m\omega^2} \approx 10^{-23} \frac{\text{m}}{\sqrt{\text{Hz}}} \left(\frac{40 \text{ kg}}{m} \right) \left(\frac{100 \text{ Hz}}{\omega} \right)^2 \left(\frac{P}{10^{-9} \text{ torr}} \right) \left(\frac{R}{15 \text{ cm}} \right) \left(\frac{\Delta T}{1^\circ \text{K}} \right)$$

where m is the mass of the disk. Thus, a temperature gradient of greater than 1000°K is necessary to make this effect appear in the detection band of Advanced LIGO. Since such temperature increases are intolerable by thermal noise considerations alone, this effect will play no part in any even remotely practical method of thermal compensation.

Chapter 3

The Numerical Tools



From the rough analysis of the previous chapter, it is now clear that thermal wavefront distortions, in particular the distortion seen in the Input Test Mass (ITM) and the Beamsplitter (BS), have serious effects on the sensitivity and controllability of a laser interferometric gravitational wave antenna. We now discuss the numerical tools by which we are able to precisely calculate thermal optical path distortions in individual optics, as well as their resulting effects on the performance of laser interferometric gravitational wave detectors.

3.1 Numerical Models of a Single Optic

In general, the total optical path length for a collimated probe beam transmitting through a heated optic is:

$$\begin{aligned}
\phi_{\text{trans}}(x, y, t) &= \frac{2\pi}{\lambda} \int_0^{h(r)} n(x, y, z, t) dz \\
&\simeq \frac{2\pi}{\lambda} \left(n_0 h_0 + \int_0^{h_0} \Delta n_T(x, y, z, t) dz + \int_0^{h_0} \Delta n_E(x, y, z, t) dz \right) \\
&\quad + \frac{2\pi}{\lambda} \left(n_0 u_z(x, y, h_0, t) - (n_0 - 1) u_z(x, y, 0, t) \right) \\
&= \phi_0 + \phi_T(x, y, t) + \phi_E(x, y, t) + \psi(x, y, t)
\end{aligned}$$

where h_0 is the nominal distance through the optic, $u_z(x, y, z, t)$ is the net local axial displacement of the material under thermoelastic deformation at the coordinates (x, y, z) at time t , Δn_T is the local refractive index change induced by temperature change (the thermo-optic effect), Δn_E is the local refractive index change induced by local thermal strain (the elasto-optic effect), and we have defined:

$$\begin{aligned}
\phi_T(x, y, t) &\equiv \frac{2\pi}{\lambda} \int_0^{h_0} \Delta n_T(x, y, z, t) dz & \phi_E(x, y, t) &\equiv \frac{2\pi}{\lambda} \int_0^{h_0} \Delta n_E(x, y, z, t) dz \\
\phi_0 &\equiv \frac{2\pi}{\lambda} n_0 h_0 & \psi(x, y, t) &\equiv \frac{2\pi}{\lambda} \left(n_0 u_z(x, y, h_0, t) - (n_0 - 1) u_z(x, y, 0, t) \right)
\end{aligned}$$

For beams reflecting directly off of the $z = 0$ surface, the total optical path length

change is simply:

$$\psi_{\text{ref}}(x, y, t) = -2\frac{2\pi}{\lambda}u_z(x, y, 0, t)$$

where $u_z(x, y, 0, t)$ is the net axial displacement of the $z = 0$ surface under thermoelastic deformation.

The finite element method [41, 43] gives us a general, precise, and rapid means of solving the partial differential equations which govern heat transfer [41] and thermal expansion [22] in a single optic. We construct a finite element model of the optic to calculate the temperature field $T(x, y, z, t)$, which allows us to compute the thermooptic effect ϕ_T . We then use the calculated temperature field T to load a mechanical finite element model, which numerically calculates the induced mechanical displacements $u_z(x, y, z, t)$ and $u_r(x, y, z, t)$ and allows us to fully calculate both surface distortions ψ and the elasto-optic effect ϕ_E . Finally, we evaluate the accuracy of our finite element models by comparing to an analytic solution which has been found for cylindrically symmetric heating of a cylindrical optic [15, 16].

3.1.1 The Finite Element Model

The finite element models developed in this work have been constructed using a pair of commercial add-ons for the technical computing package MATLAB[®], available from The Mathworks, Inc. For situations where we can reduce the governing partial differential equations to two spatial dimensions (i.e., cylindrical symmetry), thus greatly reducing the necessary computation time as well as increasing the accuracy of the solution, we use the Partial Differential Equation Toolbox (available also from The Mathworks, Inc.). For the general case with three spatial dimensions, we use the FEMLAB[®] add-on for MATLAB[®] (available from Comsol, Inc.). Both packages provide a finite element solver which addresses a generic elliptic system of N partial differential equations in M spatial dimensions (M is 2 for the PDE toolbox and 3 for FEMLAB[®]):

$$\left(d\frac{\partial}{\partial t} - \vec{\nabla} \cdot (\underline{c} \otimes \vec{\nabla}) \right) \mathbf{u} = \mathbf{f} \quad (3.1)$$

with generalized Neumann boundary conditions:

$$\left(\vec{n} \cdot \left(\underline{\mathbf{c}} \otimes \vec{\nabla}\right) + \underline{q}\right) \mathbf{u} = \mathbf{g} \quad (3.2)$$

where \mathbf{u} is the unknown we wish to obtain (an N -dimensional column vector); \mathbf{f} and \mathbf{g} are N -dimensional column vectors; \underline{q} is an $N \times M$ matrix; \vec{n} is the unit surface normal (an M -dimensional row vector) determined by the geometry we specify; d is a scalar; and $\underline{\mathbf{c}}$ is an $N \times N \times M \times M$ tensor. All of these inputs (and, of course, the output \mathbf{u}), are free to be functions of the M spatial variables or time. By the notation $\vec{\nabla} \cdot \left(\underline{\mathbf{c}} \otimes \vec{\nabla}\right) \mathbf{u}$, we mean an N -dimensional column vector with i th component:

$$\left(\vec{\nabla} \cdot \left(\underline{\mathbf{c}} \otimes \vec{\nabla}\right) \mathbf{u}\right)_i = \sum_{k=1}^M \sum_{j=1}^N \sum_{l=1}^M \left(\frac{\partial}{\partial x_k} c_{ijkl} \frac{\partial}{\partial x_l}\right) u_j.$$

We leave the general discussion of the finite element method to references [41] and [43].

One key characteristic of the generic elliptic partial differential equation 3.1 with boundary condition 3.2 is that it is fully linear in both the solution \mathbf{u} as well as the loads \mathbf{f} , \mathbf{g} , and \underline{q} . This means that if \mathbf{u}_1 is the unique solution for loads \mathbf{f}_1 , \mathbf{g}_1 , and \underline{q}_1 and \mathbf{u}_2 is the unique solution for loads \mathbf{f}_2 , \mathbf{g}_2 , and \underline{q}_2 on an identical geometry, then $\mathbf{u} = \mathbf{u}_1 + \mathbf{u}_2$ must be the unique solution for loads $\mathbf{f} = \mathbf{f}_1 + \mathbf{f}_2$, $\mathbf{g} = \mathbf{g}_1 + \mathbf{g}_2$, and $\underline{q} = \underline{q}_1 + \underline{q}_2$.

For each physical situation (i.e., either heat transfer or stress-strain), we will examine the basic physics involved to show that equations 3.1 and 3.2 apply, and then determine \mathbf{f} , \mathbf{g} , d , \underline{q} , and $\underline{\mathbf{c}}$.

3.1.2 The Thermooptic Effect

For temperature increases small compared to the ambient temperature (300 °K in our experiment), the thermooptic effect may be written:

$$\Delta n_T = \frac{dn}{dT}(T - T_\infty) \quad (3.3)$$

where $\frac{dn}{dT}$ is the refractive index derivative with respect to temperature, typically $\sim 10\text{ppm}/^\circ\text{K}$, and T_∞ is the ambient temperature ($\sim 300^\circ\text{K}$). In §3.1.2, we discuss the equations governing T , which we then reduce to two dimensions under cylindrical symmetry in §3.1.2 (which greatly reduces the computational time involved in the numerical calculation), and then examine the equations to extract the functional dependence of T on the material parameters we wish to measure in §3.1.2. The knowledge of this functional behavior of the solutions will not only give us physical insight, but will also allow us to efficiently fit materials parameters to data taken in the experimental part of this work, discussed further in Chapter 6.

The Heat Equation

Consider a 3 dimensional body occupying some closed volume V with temperature distribution $T(\vec{r})$. The amount of heat $Q(t)$ contained in this region at a time t is then, in general:

$$Q(t) = \iiint_V c\rho T(\vec{r}, t) d^3r$$

where $c = c(\vec{r})$ is the heat capacity and $\rho = \rho(\vec{r})$ is the density of the object. The net change in thermal energy per unit time is then:

$$\frac{dQ}{dt} = \iiint_V c\rho \frac{\partial T}{\partial t} d^3r$$

Suppose there is local heating of the substrate (i.e. net energy deposit per unit volume) of the form $H_s(\vec{r}, t)$. The total heat flux dP through an area $d\vec{a}$ is determined by the local thermal conductivity k_i in each direction i :

$$dP = \left(k_x \hat{x} \frac{\partial T}{\partial x} + k_y \hat{y} \frac{\partial T}{\partial y} + k_z \hat{z} \frac{\partial T}{\partial z} \right) \cdot d\vec{a} = (\underline{k} \otimes \vec{\nabla} T) \cdot d\vec{a}$$

where the thermal conductivity is expressed as a diagonal 3×3 matrix. By the conservation of energy, the net change in heat must be equal to the net heat flux through the boundary of V plus the net input of energy due to the substrate heating

H_s :

$$\frac{dQ}{dt} = \oint_{\partial V} (\underline{k} \otimes \vec{\nabla} T) \cdot d\vec{a} + \iiint_V H_s d^3r.$$

Equating the previous two expressions involving Q and using the divergence theorem, we see:

$$\iiint_V c\rho \frac{\partial T}{\partial t} d^3r = \iiint_V H_s d^3r + \iiint_V \vec{\nabla} \cdot (\underline{k} \otimes \vec{\nabla} T) d^3r$$

Now, since V is an arbitrary domain, we arrive at the Heat Equation (a special case of the Fourier/Diffusion Equation):

$$\rho c \frac{\partial T}{\partial t} - \vec{\nabla} \cdot (\underline{k} \otimes \vec{\nabla} T) = H_b \quad (3.4)$$

If the body interacts only with a thermal radiation bath of temperature T_∞ (i.e. it is suspended in a vacuum enclosure with walls at temperature T_∞), the heat loss on the boundary is governed by the Stefan-Boltzmann law:

$$\hat{n} \cdot (\underline{k} \otimes \vec{\nabla} T) = -\sigma\epsilon(T^4 - T_\infty^4) + H_c$$

where \hat{n} is the unit normal on the boundary, ϵ is the emissivity, σ is the Stefan-Boltzmann constant, and $H_c = H_c(x, y, z, t)$ is the rate of surface (coating) heating due to an external source (power input per unit area). If the temperature increase is sufficiently small, so that $\frac{T-T_\infty}{T_\infty} \ll 1$, we may linearize the boundary conditions so that:

$$\hat{n} \cdot (\underline{k} \otimes \vec{\nabla} T) = -4\sigma\epsilon T_\infty^3 (T - T_\infty) + H_c. \quad (3.5)$$

Hence, heat transfer for $\frac{T-T_\infty}{T_\infty} \ll 1$ is governed by the elliptic partial differential

equation 3.1 with boundary equation 3.2 with coefficients:

$$\begin{array}{l}
 N = 1 \\
 M = 3 \\
 d = \rho c
 \end{array}
 \begin{array}{l}
 \mathbf{f} = H_s(x, y, z, t) \\
 \mathbf{g} = H_c(x, y, z, t) \\
 \underline{q} = 4\sigma\epsilon T_\infty^3 \\
 \mathbf{u} = (T(x, y, z, t) - T_\infty)
 \end{array}
 \quad \underline{\mathbf{c}} = \begin{pmatrix} k_x & 0 & 0 \\ 0 & k_y & 0 \\ 0 & 0 & k_z \end{pmatrix} \quad (3.6)$$

It is interesting to note that, in the steady-state, equation 3.4 is simply Poisson's Equation, which also governs the static scalar potential in electromagnetism. We suspect that an analogous form of Gauss' Law applies here, where scalar potential corresponds to temperature and charge corresponds to power input over thermal conductivity. This becomes clear when we recall how we used the divergence theorem in the derivation of the Heat Equation above:

$$\iiint_D \vec{\nabla} \cdot (\underline{\mathbf{k}} \otimes \vec{\nabla} T) d^3r = \oiint_{\partial D} (\underline{\mathbf{k}} \otimes \vec{\nabla} T) \cdot d\vec{\mathbf{a}}$$

where D is a closed region contained within the closed body V , and we arrive at the integral form of the steady-state Heat Equation:

$$-\oiint_{\partial D} (\underline{\mathbf{k}} \otimes \vec{\nabla} T) \cdot d\vec{\mathbf{a}} = \iiint_D H_s d^3r = P_D \quad (3.7)$$

where P_D is the total power input to the closed region D . In the case where the substrate heating depends only on the radial cylindrical coordinate (i.e., $H_b(r, \theta, z) = H_b(r)$), and the entire body V is much larger than the spatial extent of the substrate heating, we ignore the boundaries and construct a Gaussian cylinder of radius r and length h , thus allowing us to solve equation 3.7 as:

$$-2\pi r h k_r \frac{\partial T}{\partial r} = 2\pi h \int_0^r H_b(r') r' dr'$$

and the radial temperature field is given by:

$$T(r) - T(0) = -\frac{1}{k_r} \int_0^r \frac{dr'}{r'} \int_0^{r'} H_b(r'') r'' dr'' = -\frac{1}{2\pi k_r} \int_0^r \frac{P(r') dr'}{r'}.$$

where $P(r')$ is the total power per unit axial length contained within $r < r'$. If there is no heat input for radii greater than R , $P(r > R) = P_R$ is constant and, for $r > R$:

$$T(r) - T(R) = -\frac{P_R}{2\pi k_r} \log\left(\frac{r}{R}\right).$$

If the total thickness of the optical material is h , then the optical path distortion per unit absorbed power due to the thermo-optic effect is:

$$\frac{\phi_T(r) - \phi_T(R)}{hP_R} = \frac{dn}{dT} \frac{1}{2\pi k_r} \begin{cases} \int_r^R \frac{P(r') dr'}{P_R r'} & \text{for } r < R \\ -\log\left(\frac{r}{R}\right) & \text{for } r \geq R. \end{cases} \quad (3.8)$$

Although this equation is of extremely limited utility, as the large LIGO beams are comparable in size to the dimensions of the optics which they heat, two subtle points in the behavior of the thermo-optic effect: (1) in the area outside of a region where heat is input, the thermo-optic optical path distortion will behave identically for all spatial patterns of the heat input into the region, and (2) a different thermal conductivity does not alter the spatial pattern of the thermo-optical optical path distortion, but rather serves to change its amplitude.

We'll use the three dimensional finite element model to calculate situations where there is a non-cylindrical anisotropy in the materials parameters (e.g., M-axis sapphire test masses), or where the heating terms are not on-axis (when we perform the necessary calculations for the scanning laser actuator in Chapter 5).

Two Dimensional Form

It is more efficient and accurate to solve in two spatial dimensions, if the symmetry of the situation will allow it. For a cylindrical optic with cylindrically symmetric material parameters (either isotropic, like fused silica, or anisotropic in the radial

and axial directions, as is the case for C-axis sapphire) and loads then the heat equation 3.4 can be readily reduced to two dimensions:

$$\rho c r \frac{\partial T}{\partial t} - \vec{\nabla}_2 \cdot (r \underline{k} \otimes \vec{\nabla}_2 T) = r H_s \quad (3.9)$$

where we have written $\vec{\nabla}_2 \equiv \hat{r} \frac{\partial}{\partial r} + \hat{z} \frac{\partial}{\partial z}$. This is simply the Heat Equation for a plate in the right half plane with density ρr , heat capacity c , and thermal conductivity $r \underline{k}$.

Similarly, the boundary conditions are rewritten in two-dimensional form as:

$$\hat{n} \cdot (r \underline{k} \otimes \vec{\nabla} T) = -4\sigma r \epsilon T_\infty^3 (T - T_\infty) + r H_s. \quad (3.10)$$

and in the notation of §3.1.1, we now have:

$N = 1$	$\mathbf{f} = r H_s(r, z, t)$	$\underline{\mathbf{c}} = \begin{pmatrix} r k_r & 0 \\ 0 & r k_z \end{pmatrix}$	(3.11)
$M = 2$	$\mathbf{g} = r H_c(r, z, t)$		
$d = \rho c r$	$\underline{\mathbf{q}} = 4r\sigma\epsilon T_\infty^3$		
	$\mathbf{u} = (T(r, z, t) - T_\infty)$		

Functional Dependences

Although we now have the means to numerically find T , and thus ϕ_T , it is useful to see how the thermo-optic optical path distortion behaves as a function of the material parameters of the problem. The temperature gradients within the optic will be primarily dependent on the thermal conductivity, and the mean temperature increase of the optic \bar{T} depends entirely on its surface area A , emissivity ϵ , and power input P_{in} . If we write the temperature increase in terms of spatial AC and DC components:

$$(T - T_\infty) = \delta T + \bar{T} \quad \text{where} \quad \bar{T} \equiv \frac{P_{in}}{4\sigma\epsilon T_\infty^3 A}.$$

the heat equation becomes:

$$\begin{aligned} \left(\rho c \frac{\partial}{\partial t} - \vec{\nabla} \cdot (\underline{k} \otimes \vec{\nabla}) \right) \delta T &= H_b && \text{inside} && (3.12) \\ \hat{n} \cdot (\underline{k} \otimes \vec{\nabla}) \delta T &\approx -\frac{P_{in}}{A} + H_c && \text{on the boundary.} \end{aligned}$$

where we have utilized the approximation that the radiation by the AC component of the temperature field is dominated by the DC component at locations where there is no surface absorption ($H_c = 0$). This is not necessarily true in regions where $H_c > 0$, but it is necessarily true here that here the heating term itself will dominate the boundary condition (otherwise there would be no net heat input into the material).

For an isotropic material, it is now easy to see the functional dependence of the temperature field on the material parameters. Making the change of variables:

$$t \rightarrow \beta t' \quad \delta T \rightarrow \frac{F}{k}$$

where $\beta \equiv \frac{\rho c}{k}$ is the inverse diffusivity, and the heat equation now becomes:

$$\begin{aligned} \left(\frac{\partial}{\partial t'} - \vec{\nabla}^2 \right) F &= H_b && \text{inside} \\ \hat{n} \cdot \vec{\nabla} F &\approx -\frac{P_{in}}{A} + H_c && \text{on the boundary.} \end{aligned}$$

The function F is seen to be purely dependent on the geometry of the optic as well as the shape and magnitude of the heating terms, but *not* significantly on any of the isotropic material parameters. Thus, the AC component of the temperature field can be written as:

$$\delta T(x, y, z, t) \simeq \frac{1}{k} F(x, y, z, t/\beta) \quad (3.13)$$

where the function F is dependent only on the heating terms and the geometry.

For the anisotropic case, where the thermal conductivity is free to take different values in different directions, the argument above doesn't quite hold. Assuming cylindrical symmetry, the radial temperature gradients near the center of the optic, hence the phase distortion, will be governed by the radial thermal conductivity as well as

the shape and intensity of the heating pattern. If we integrate equation 3.12 over the optical axis z , applying the change of variables $t \rightarrow \beta_r t'$, with $\beta_r \equiv \frac{\rho c}{k_r}$, we arrive that:

$$\begin{aligned} \left(\frac{d}{dt'} - \frac{1}{r} \frac{d}{dr} r \frac{d}{dr} \right) \phi_T(r, t') &\approx \frac{2\pi}{\lambda} \frac{dn}{dT} \frac{1}{k_r} \left(H(r, t) - \frac{2P_{in}}{A} \right) && \text{on } (0, R) && (3.14) \\ \frac{d\phi_T}{dr} &= 0 && \text{at } r = 0 \\ \frac{d\phi_T}{dr} &\approx -\frac{2\pi}{\lambda} \frac{dn}{dT} \frac{1}{k_r} \frac{hP_{in}}{A} && \text{at } r = R \end{aligned}$$

where the heating term $H(r)$ is the total heating per unit area due to surface heating of both surfaces as well as bulk heating through the optic at coordinate r . Thus, the solution for the thermo-optic optical path distortion may be written as:

$$\boxed{\phi_T(r, t) \approx \frac{dn}{dT} \frac{1}{k_r} f(r, t/\beta_r)} \quad (3.15)$$

where the function $f(r, t)$ is dependent only on the geometry and shape of the heating terms.

Comparison to an Analytic Model

Although a general analytic solution has been found by P. Hello and J. Vinet for cylindrically symmetric heating in a cylindrical optic [15], it is, in general, quite complicated and must be numerically evaluated anyway (a number of transcendental equations must be solved and a slowly converging infinite series must be summed). The advantage of the two-dimensional finite element model is the speed at which the solution can be computed, and the obvious advantage of the three dimensional model is that it can handle situations where cylindrical symmetry does not hold (although the three dimensional model is very computationally intensive). Here we use the solution of Hello and Vinet [15], as implemented by D. Djambazov and D. Coyne for Gaussian beam heating in a Initial LIGO Input Test Mass [9], as a spot-check comparison for the accuracy of our finite element models.

Traditionally, the solution for coating heating and substrate heating are solved

separately, and linearity is exploited to find the solution for any arbitrary combination of coating and substrate heating. Figure 3-1 shows the results of the two dimensional finite element calculation for 1 Watt of uniform bulk absorption in a Initial LIGO input test mass, and compares the calculated optical path distortion with that of the analytic solution, as well as that of the unbounded solution of equation 3.8. The distortion predicted by the two dimensional finite element model is seen to agree with the analytic model over the entire optic to within 0.5%, while the unbounded calculation differs by approximately 10%. Figure 3-2 shows the same comparison for 1 Watt of coating absorption. Note that, although the temperature fields are quite different, the thermo-optical path distortion for coating absorption looks nearly identical in shape and magnitude to that for uniform bulk absorption.

Figure 3-3 shows the comparison of the three dimensional finite element model to the analytic model for both coating and substrate absorption. The agreement here is excellent over the entire optic for coating absorption (less than 0.5%), and is, as is also the case of the two dimensional model, slightly less so for bulk absorption (about 2%).

In general, the finite element model is found to closely correspond to the analytic model for beams comparable to the size of the optic. We note here, however, that the accuracy deteriorates (not surprisingly) when the beam becomes small enough to become comparable to the average size of the individual elements of the finite element mesh.




Figure 3-1: Temperature field (above) and thermo-optic optical path distortion (below) as calculated by the 2D finite element model for 1 Watt of substrate absorption in a Initial LIGO Input Test Mass.





Figure 3-2: Temperature field (above) and thermo-optic optical path distortion (below) as calculated by the 2D finite element model for 1 Watt of coating absorption in a Initial LIGO Input Test Mass.

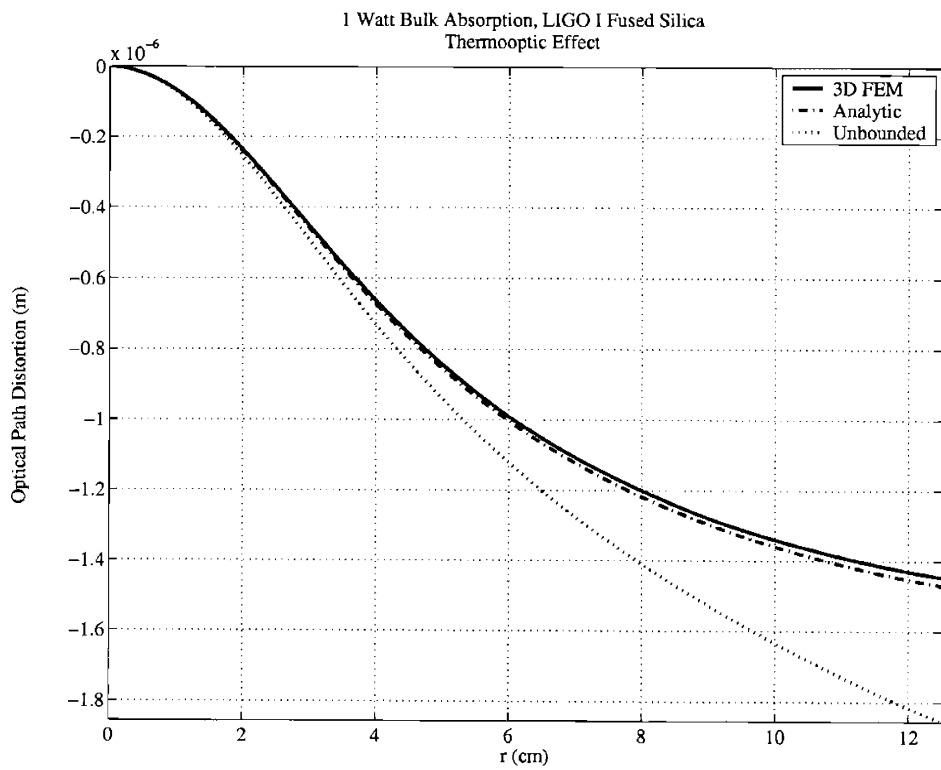
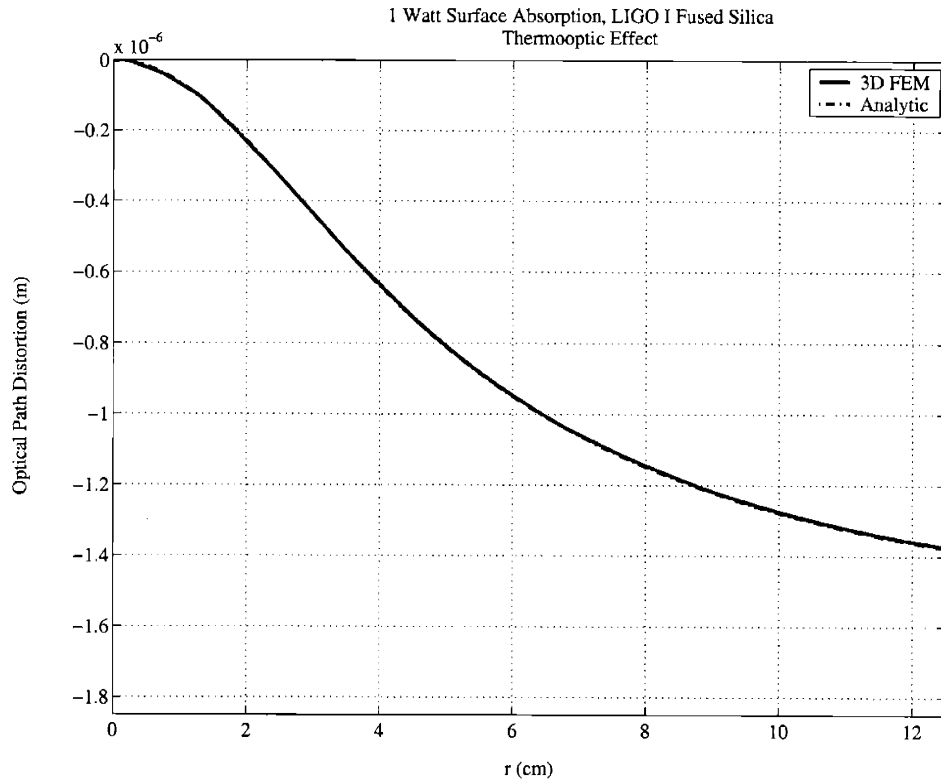


Figure 3-3: Comparison of the 3D finite element model to the analytic solution for 1 Watt of coating absorption (above) and 1 Watt of substrate absorption (below) in a Initial LIGO Input Test Mass.

3.1.3 Thermoelastic Deformation

Once we are able to determine the temperature field $T(x, y, z, t)$, we can use this to compute the local mechanical loads due to thermal expansion, and then use another finite element model to numerically compute the resulting deformation field $\vec{u}(x, y, z, t)$. Since the time constant governing the evolution of temperature (typically minutes or hours) is much smaller than the sound crossing time for the materials we are concerned with (typically tens of microseconds), we explicitly assume that \vec{u} evolves in an adiabatic manner; i.e., inertial terms are not considered in the equations governing \vec{u} . §3.1.3 discusses the differential equations governing the deformation field $\vec{u}(x, y, x, t)$, which we then reduce to two dimensional form in §3.1.3. Similar to §3.1.2, we extract the functional dependence of u_z on the quantities we wish to measure in §3.1.3, and then compare to the two dimensional analytic solution in the final section.

The Stress-Strain Relations (Hooke's Law)

In general, in the limit of small distortions $\vec{u} = (u_x, u_y, u_z)$, strains in a material are completely described by local derivatives of the deformation components:

$$\epsilon_{ij} = \frac{1}{2} \left(\frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j} \right)$$

This 3×3 “strain tensor” is symmetric, so it is traditional [22] and somewhat simpler to redefine in terms of a 6-element strain vector $\vec{\epsilon} \equiv (\epsilon_{xx}, \epsilon_{yy}, \epsilon_{zz}, 2\epsilon_{yz}, 2\epsilon_{xz}, 2\epsilon_{xy})$. Similarly, we define the 6-element stress vector as $\vec{\sigma} \equiv (\sigma_{xx}, \sigma_{yy}, \sigma_{zz}, \sigma_{yz}, \sigma_{xz}, \sigma_{xy})$ and use Hooke's law to relate stresses (forces) to strains (physical displacements) via the elastic modulus matrix $\underline{\mathcal{C}}_3$:

$$\vec{\sigma} = \vec{\epsilon} \underline{\mathcal{C}}_3 - (T - T_\infty) \vec{\gamma} \quad (3.16)$$

where we have not considered inertial terms, $\vec{\gamma}$ is the stress temperature modulus vector, related to the thermal expansion coefficient vector $\vec{\alpha} = (\alpha_x, \alpha_y, \alpha_z, 0, 0, 0)$ by:

$\vec{\gamma} = \vec{\alpha}\underline{c}_3$ and the symmetric elastic modulus matrix \underline{c}_3 has elements:

$$\underline{c}_3 \equiv \begin{pmatrix} c_{11} & c_{12} & c_{13} & 0 & 0 & 0 \\ c_{12} & c_{22} & c_{23} & 0 & 0 & 0 \\ c_{13} & c_{23} & c_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & c_{66} \end{pmatrix}$$

For an isotropic material with Young's modulus E and Poisson ratio ν , the elastic modulus matrix elements are simply:

$$c_{11} = c_{22} = c_{33} = \frac{E(1 - \nu)}{(1 + \nu)(1 - 2\nu)}$$

$$c_{12} = c_{13} = c_{23} = \frac{E\nu}{(1 + \nu)(1 - 2\nu)}$$

$$c_{44} = c_{55} = c_{66} = \frac{E}{2(1 + \nu)}$$

To solve for the displacements \vec{u} , we use the stress equilibrium conditions in the steady state (i.e., the net force in every direction must vanish):

$$\frac{\partial \sigma_{ii}}{\partial x_i} + \frac{\partial \sigma_{ji}}{\partial x_j} + \frac{\partial \sigma_{ki}}{\partial x_k} = 0 \quad (3.17)$$

Inserting equation 3.16 into equation 3.17 and simplifying, we find the differential equation governing $\vec{u}(x, y, z)$:

$$\vec{\nabla} \cdot (\underline{C}_3 \otimes \vec{\nabla}) \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix} = \begin{pmatrix} \gamma_x \frac{\partial}{\partial x} \\ \gamma_y \frac{\partial}{\partial y} \\ \gamma_z \frac{\partial}{\partial z} \end{pmatrix} (T - T_\infty) \quad (3.18)$$

where the rank four tensor $\underline{\mathbf{C}}_3$ has elements:

$$\underline{\mathbf{C}}_3 = \left(\begin{array}{ccc} \begin{pmatrix} c_{11} & 0 & 0 \\ 0 & c_{66} & 0 \\ 0 & 0 & c_{44} \end{pmatrix} & \begin{pmatrix} 0 & c_{12} & 0 \\ c_{66} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 0 & c_{13} \\ 0 & 0 & 0 \\ c_{44} & 0 & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & c_{66} & 0 \\ c_{12} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} c_{66} & 0 & 0 \\ 0 & c_{22} & 0 \\ 0 & 0 & c_{55} \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & c_{23} \\ 0 & c_{55} & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & 0 & c_{44} \\ 0 & 0 & 0 \\ c_{13} & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & c_{55} \\ 0 & c_{23} & 0 \end{pmatrix} & \begin{pmatrix} c_{44} & 0 & 0 \\ 0 & c_{55} & 0 \\ 0 & 0 & c_{33} \end{pmatrix} \end{array} \right) \quad (3.19)$$

On the boundaries, the normal forces on must vanish, as well as components of shear that are parallel to the surface. This is equivalent to the Neumann boundary condition:

$$\vec{n} \cdot (\underline{\mathbf{C}}_3 \otimes \vec{\nabla}) \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix} = \begin{pmatrix} \gamma_x n_x \\ \gamma_y n_y \\ \gamma_z n_z \end{pmatrix} (T - T_\infty) \quad (3.20)$$

where n_i is the component of the surface normal in the i th direction.

So, the coefficients we insert into the finite element model (in the convention of §3.1.1) are:

$N = 3$	$\mathbf{f} = \begin{pmatrix} \gamma_x \frac{\partial}{\partial x} \\ \gamma_y \frac{\partial}{\partial y} \\ \gamma_z \frac{\partial}{\partial z} \end{pmatrix} (T - T_\infty)$	$\underline{\mathbf{c}} = \underline{\mathbf{C}}_3$ (equation 3.19)	(3.21)
$M = 3$	$\mathbf{g} = \begin{pmatrix} \gamma_x n_x \\ \gamma_y n_y \\ \gamma_z n_z \end{pmatrix} (T - T_\infty)$		
$d = 0$	$\underline{\mathbf{q}} = 0$		
$\mathbf{u} = \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix}$			

Two Dimensional Form

As in the case of the calculation of temperature, for silica and C-axis sapphire with cylindrically symmetric boundaries and loads we can reduce to two dimensions and greatly reduce the required computation time for the numerical model. In the limit of small, cylindrically symmetric distortions $\mathbf{u} = \begin{pmatrix} u_r \\ u_z \end{pmatrix}$, strains are now described in cylindrical coordinates by [22]:

$$\begin{aligned} \epsilon_{rr} &= \frac{\partial u_r}{\partial r} & \epsilon_{rz} &= \frac{1}{2} \left(\frac{\partial u_r}{\partial z} + \frac{\partial u_z}{\partial r} \right) \\ \epsilon_{\theta\theta} &= \frac{u_r}{r} & \epsilon_{\theta z} &= 0 \\ \epsilon_{zz} &= \frac{\partial u_z}{\partial z} & \epsilon_{\theta r} &= 0 \end{aligned}$$

The 3×3 strain tensor has 4 nonzero elements, so we define a 4-element strain vector as $\vec{\epsilon} \equiv (\epsilon_{rr}, \epsilon_{\theta\theta}, \epsilon_{zz}, 2\epsilon_{rz})$. Similarly, we define the 4-element stress vector as $\vec{\sigma} \equiv (\sigma_{rr}, \sigma_{\theta\theta}, \sigma_{zz}, \sigma_{rz})$ and again use Hooke's law to relate stresses to strains via a modified elastic modulus matrix \underline{c}_2 :

$$\vec{\sigma} = \vec{\epsilon} \underline{c}_2 - (T - T_\infty) \vec{\gamma} \quad (3.22)$$

where the stress temperature modulus vector $\vec{\gamma}$ is related to the thermal expansion coefficient vector $\vec{\alpha} = (\alpha_r, \alpha_r, \alpha_z, 0)$ by $\vec{\gamma} = \vec{\alpha} \underline{c}_2$ so that

$$\gamma_r = (c_{11} + c_{12})\alpha_r + c_{13}\alpha_z \quad \gamma_z = 2c_{13}\alpha_r + c_{33}\alpha_z$$

and the modified elastic modulus matrix \underline{c}_2 has elements:

$$\underline{c}_2 \equiv \begin{pmatrix} c_{11} & c_{12} & c_{13} & 0 \\ c_{12} & c_{11} & c_{13} & 0 \\ c_{13} & c_{13} & c_{33} & 0 \\ 0 & 0 & 0 & c_{44} \end{pmatrix}$$

To solve for the displacements \mathbf{u} , we use the stress equilibrium conditions in the steady state (i.e., the net force in every direction must vanish):

$$\begin{aligned} \frac{\partial \sigma_{rr}}{\partial r} + \frac{\partial \sigma_{rz}}{\partial z} + \frac{\sigma_{rr} - \sigma_{\theta\theta}}{r} &= 0 \\ \frac{\partial \sigma_{rz}}{\partial r} + \frac{\partial \sigma_{zz}}{\partial z} + \frac{\sigma_{rz}}{r} &= 0 \end{aligned} \quad (3.23)$$

Inserting equation 3.22 into equation 3.23 and simplifying, we find the differential equation governing $\mathbf{u}(r, z)$:

$$\vec{\nabla} \cdot (\underline{\mathbf{C}}_2 \otimes \vec{\nabla}) \begin{pmatrix} r u_r \\ u_z \end{pmatrix} = \begin{pmatrix} \gamma_r \frac{\partial}{\partial r} \\ \gamma_z r \frac{\partial}{\partial z} \end{pmatrix} (T - T_\infty) \quad (3.24)$$

where the rank four tensor $\underline{\mathbf{C}}_2$ has elements:

$$\underline{\mathbf{C}}_2 = \begin{pmatrix} \begin{pmatrix} \frac{c_{11}}{r} & 0 \\ 0 & \frac{c_{44}}{r} \end{pmatrix} & \begin{pmatrix} 0 & c_{13} \\ c_{44} & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & c_{44} \\ c_{13} & 0 \end{pmatrix} & \begin{pmatrix} c_{44} r & 0 \\ 0 & c_{33} r \end{pmatrix} \end{pmatrix}$$

On the boundaries of the optic, stress in the direction of the normal must vanish. In our two dimensional framework, this is found to be equivalent to a generalized Neumann condition of the form:

$$\left(\vec{n} \cdot (\underline{\mathbf{C}}_2 \otimes \vec{\nabla}) + \begin{pmatrix} -\frac{c_{11}-c_{12}}{r^2} n_r & 0 \\ 0 & 0 \end{pmatrix} \right) \begin{pmatrix} r u_r \\ u_z \end{pmatrix} = \begin{pmatrix} \gamma_r n_r \\ \gamma_z n_z r \end{pmatrix} (T - T_\infty) \quad (3.25)$$

where n_i is the component of the surface normal in the i th direction.

So, the coefficients we insert into the finite element model (in the convention of

§3.1.1) are found to be:

$$\begin{array}{l}
 N = 2 \\
 M = 2 \\
 d = 0 \\
 \mathbf{u} = \begin{pmatrix} ru_r \\ u_z \end{pmatrix}
 \end{array}
 \quad
 \begin{array}{l}
 \mathbf{f} = \begin{pmatrix} \gamma_r \frac{\partial}{\partial r} \\ \gamma_z r \frac{\partial}{\partial z} \end{pmatrix} (T - T_\infty) \\
 \mathbf{q} = \begin{pmatrix} -\frac{c_{11}-c_{12}}{r^2} n_r & 0 \\ 0 & 0 \end{pmatrix} \\
 \mathbf{g} = \begin{pmatrix} \gamma_r n_r \\ \gamma_z r n_z \end{pmatrix} (T - T_\infty)
 \end{array}
 \quad
 \underline{\mathbf{c}} = \left(\begin{array}{cc|cc}
 \frac{c_{11}}{r} & 0 & 0 & c_{13} \\
 0 & \frac{c_{44}}{r} & c_{44} & 0 \\
 0 & c_{44} & c_{44}r & 0 \\
 c_{13} & 0 & 0 & c_{33}r
 \end{array} \right)
 \end{array}
 \quad (3.26)$$

Functional Dependences

The functional dependence here becomes more complicated, as the deformation field will depend on the characteristics of the temperature field, as well as the possibly anisotropic thermal expansion coefficient. A constant (small) increase in temperature has the trivial solution of uniform thermal expansion in all directions, thus is of no concern to us here. Recalling the discussion of the dependence of the temperature field in §3.1.2: As the mean temperature increase is primarily determined by the optic's emissivity and area, there will be only a weak dependence of the relative surface distortion on this quantity. The magnitude of the gradients of the temperature field (but not the spatial pattern) are determined primarily by the thermal conductivity, thus we expect the magnitude (but not the shape) of the surface distortions to also be primarily governed by thermal conductivity. The time evolution of the temperature field is governed by the diffusivity $\beta^{-1} = \frac{k}{\rho c}$, so, since the thermal distortion field evolves in an adiabatic way, we expect the time evolution of the surface distortion to be governed in an identical manner. For an isotropic material, then, similar to equation 3.13 for temperature:

$$u_z(x, y, z, t) \approx \frac{\alpha}{k} G(x, y, z, t/\beta)$$

where the function G is dependent only on the heating terms, the geometry, and the elastic modulus matrix elements.

In an anisotropic material, however, the thermal conductivity (hence diffusivity) and thermal expansion coefficient are different in different directions. Again assuming cylindrical symmetry, and examining the thermal expansion equations 3.24 and 3.25, it is clear the solution u_z may be written in terms of two components: a component which solves the thermal expansion equations with $\gamma_r = 0$ and $\gamma_z = 1$, representing axial expansion caused by local axial thermal expansion, and a component which solves the equations with $\gamma_r = 1$ and $\gamma_z = 0$, representing axial expansion caused by local radial thermal expansion coupling into the axial direction (i.e., buckling). The dependence of axial thermal expansion on radial and axial thermal conductivities is not obvious, however, although through experimentation with the two dimensional finite element model, we have observed the heuristic behavior:

$$\boxed{u_z(r, z, t) \approx \frac{\gamma_r}{k_z} g_r(r, z, t/\beta_z) + \frac{\gamma_z}{k_r} g_z(r, z, t/\beta_r)} \quad (3.27)$$

where g_z is the component of u_z which solves the thermal expansion equation with $\gamma_r = 0$ and $\gamma_z = 1$, g_r solves it for $\gamma_r = 1$ and $\gamma_z = 0$, and both function g_r and g_z are dependent only on the heating terms, the geometry, and the elastic modulus matrix elements.

Comparison to an Analytic Model

Here we use the solution of Hello and Vinet [16], as implemented (and corrected) by R. Beausoleil for Gaussian beam heating in any LIGO Core Optic [3], as a spot-check comparison for the accuracy of our finite element models.

The top of figure 3-4 illustrates the temperature fields and deformations (highly exaggerated) calculated using the two dimensional finite element model for both substrate and coating absorption. The bottom of the figure shows the results of the two dimensional finite element calculation for 1 Watt of uniform bulk absorption in a Initial LIGO input test mass, and compares the calculated surface distortion with that

of the analytic solution. The two dimensional numerical model is seen to behave the worst near the center of the optic, mainly due to the fact that the differential equations governing thermal expansion are poorly behaved near zero (some terms approach zero while others approach infinity; recall \underline{c} in equation 3.26). Using a polynomial fitting technique to remove the ill-behaved region near zero, a correspondence with the analytic model to within about 2% is achieved over the entire optic.

Figure 3-5 shows the same comparison for 1 Watt of coating absorption in a Initial LIGO input test mass, but this time showing the axial distortion on both the front (where the absorption is occurring) as well as the back of the optic. On the front, the same poor behavior of the solution is seen near zero, which we attempt to address with our polynomial smoothing technique, and we get a correspondence with the analytic model of about 3% over the entire optic. The rear distortion is seen to correspond extremely well with the analytic model.

Figures 3-6 and 3-7 detail the same information for the three dimensional finite element model as seen in figures 3-4 and 3-5 for the two dimensional finite element model. The 3D model does not behave poorly near $r = 0$, for obvious reasons, but the overall accuracy is less than that of the 2D model. Overall, the 3D model corresponds to the analytical model to within about 5% on all surfaces for both modes of absorption.

In general, as is the case for the numerical temperature solution, the finite element model is found to correspond well to the analytic model for beams comparable to the size of the optic. Here too, the accuracy of the numerical solution deteriorates when the beam becomes small enough to become comparable to the average size of the individual elements of the finite element mesh.





Figure 3-4:





Figure 3-5:





Figure 3-6:

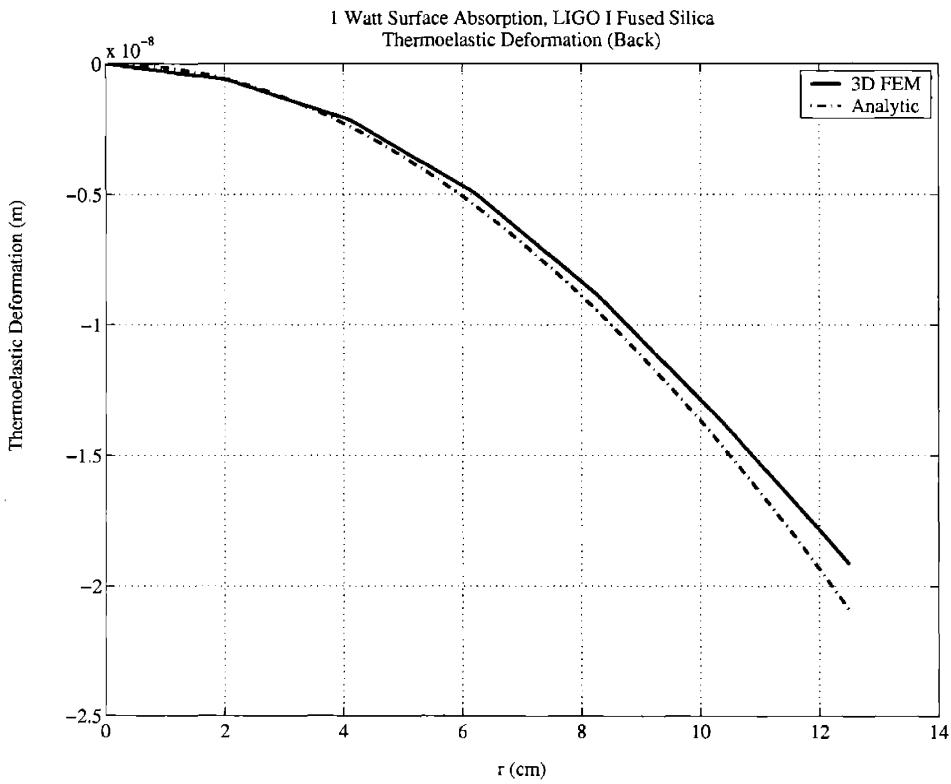
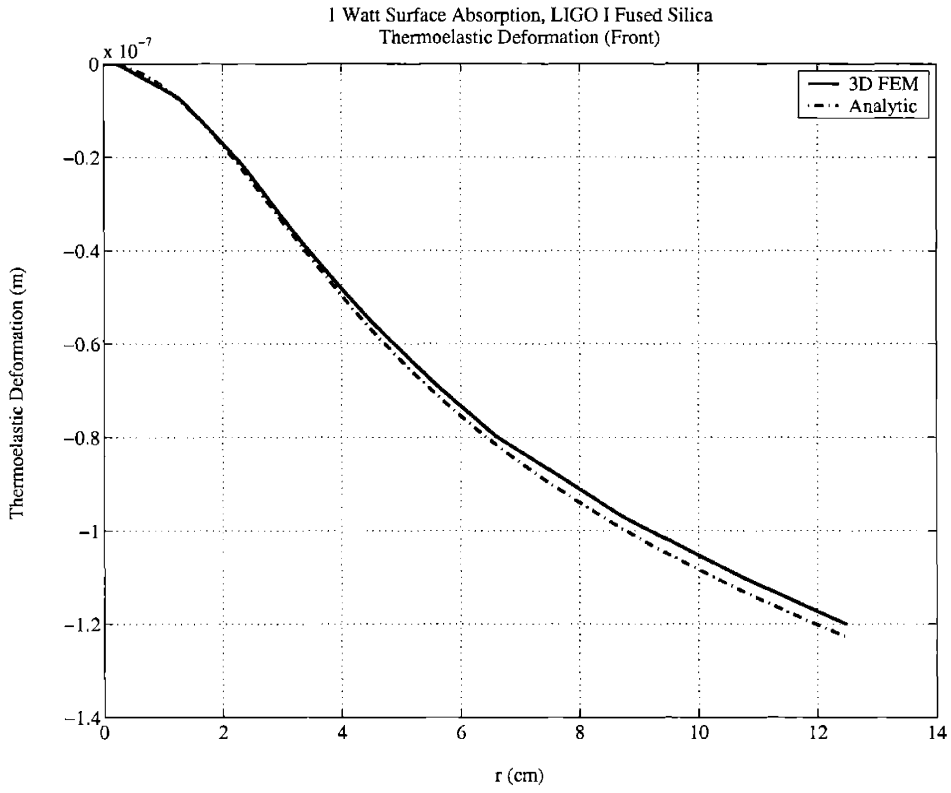


Figure 3-7: 1 Watt of coating absorption (above) and 1 Watt of substrate absorption (below) in a Initial LIGO Input Test Mass.

3.1.4 The Elasto-optic Effect

This effect is simply stress-induced birefringence, where local strains in an optical material will cause a local change in the index of refraction in a polarization dependent manner. The effect is described by the elasto-optic tensor p_{ijkl} , which relates the change in the impermeability tensor (sometimes called the “indicatrix” or the “index ellipsoid”):

$$\Delta \left(\frac{1}{n^2} \right)_{ij} \approx -\frac{2}{n^3} \Delta n_{ij}$$

to the local strain tensor ϵ_{ij} . Assuming the principal axes of the impermeability tensor are aligned with our coordinate axes, and using the traditional 6 dimensional reduced notation as in §3.1.3, we write [32]:

$$\begin{pmatrix} \Delta n_x \\ \Delta n_y \\ \Delta n_z \\ \Delta n_{yz} \\ \Delta n_{xz} \\ \Delta n_{xy} \end{pmatrix} = -\frac{n^3}{2} \begin{pmatrix} p_{11} & p_{12} & p_{31} & 0 & 0 & 0 \\ p_{12} & p_{11} & p_{31} & 0 & 0 & 0 \\ p_{31} & p_{31} & p_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & p_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & p_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2}(p_{11} - p_{12}) \end{pmatrix} \begin{pmatrix} \epsilon_x \\ \epsilon_y \\ \epsilon_z \\ 2\epsilon_{yz} \\ 2\epsilon_{xz} \\ 2\epsilon_{xy} \end{pmatrix}$$

where we have written this with the 3rd coordinate axis (the z axis) as the crystal symmetry axis. So, for light polarized along the j th coordinate axis, the optical path distortion due to the elasto-optic effect is:

$$\phi_E(x, y) = \frac{2\pi}{\lambda} \int_0^h \Delta n_j dz = -\frac{n^3}{2} \frac{2\pi}{\lambda} \int_0^h \sum_{m=1}^3 p_{jm} \epsilon_m(x, y, z) dz \quad (3.28)$$

To find the approximate the magnitude of this effect (an expression which we used in the previous chapter), assume cylindrical symmetry with a probe polarization along the x -axis. Note that this symmetry constrains $\epsilon_m(0, 0, z) = 0$. Along the x -axis, the strain can be approximated:

$$\epsilon_x = \frac{\partial u_x}{\partial x} \approx \frac{\partial}{\partial x} \alpha x \Delta T \approx \alpha \Delta T$$

and the approximate elasto-optical optical path distortion is:

$$\Delta S_E \approx -\frac{n^3}{2} p_{xx} \int_0^h \epsilon_x(x, y, z) dz \approx -p_{xx} \alpha \int_S \Delta T ds$$

Although we will not utilize it any further in this work (we will use the 3D version above for M-axis sapphire test masses), it is useful to note that one can calculate the thermo-optic effect from the cylindrically symmetric distortions u_r and u_z calculated in our cylindrically symmetric model. Changing to rectangular coordinates, we see that:

$$\epsilon_x = \frac{\partial u_x}{\partial x} = \cos^2 \theta \frac{\partial u_r}{\partial r} + \sin^2 \theta \frac{u_r}{r} \quad \epsilon_y = \frac{\partial u_y}{\partial y} = \sin^2 \theta \frac{\partial u_r}{\partial r} + \cos^2 \theta \frac{u_r}{r}$$

$$\epsilon_z = \frac{\partial u_z}{\partial z}$$

and the elasto-optic effect can be computed from equation 3.28.

The Elasto-optic Effect in a Advanced LIGO Sapphire Test Mass

For M-axis sapphire input test masses, as will most likely be the case for Advanced LIGO (sapphire is very difficult to grow along the C-axis), the crystal symmetry axis is parallel to the optic's face, hence we must utilize the full three dimensional finite element model to calculate the temperature and deformation fields. Figure 3-8 shows the results of this calculation. The elasto-optic tensor for sapphire is approximately "diagonal" (see Appendix A), meaning that the index of refraction for polarization along a principal direction is only appreciably affected by the strain along that same direction. For central heating by a Gaussian beam, the temperature gradients, hence the strains, are primarily in the radial direction, so we expect that the effect will be most significant where the local polarization is aligned with the radial direction. This behavior is clearly indicated in figure 3-8, although the numerical data are slightly noisy due to the fact that numerical differentiation is performed over an irregular finite element mesh to compute the local strains ϵ_{ij} . The noisiness is slightly more-so for coating absorption, as the gradients of the deformation field are sharper compared

to that of bulk absorption.

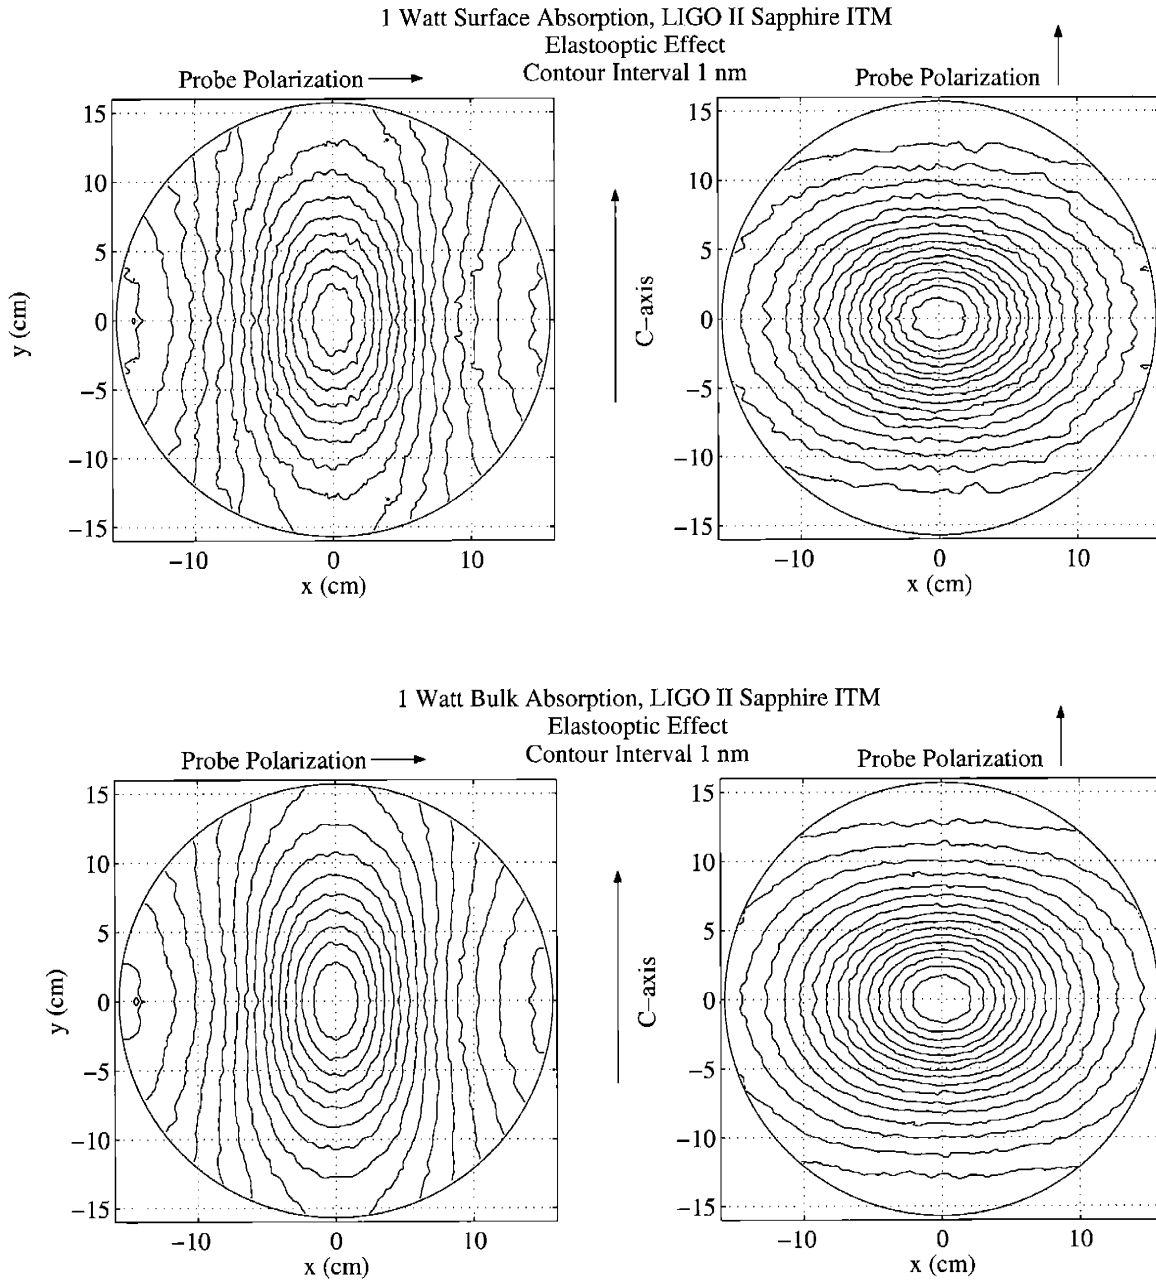


Figure 3-8: The elasto-optic effect in a Advanced LIGO sapphire input test mass for 1 Watt of coating absorption (above) and 1 Watt of substrate absorption (below). The C-axis is aligned with the Y-axis in both diagrams, and the probe polarization is indicated on each graph.

3.2 Modeling the Entire Interferometer

To calculate the effects that thermal distortions have on the circulating optical fields of a laser interferometric gravitational wave antenna, we use the Melody¹ modal model of gravitational wave antennae, constructed and maintained by Raymond Beausoleil of Stanford University/Hewlett-Packard Laboratories [3]. In the analyses performed below, the model uses the 66 lowest order Hermite-Gauss modes (TEM_{*nm*} modes with $n + m \leq 10$).

Melody is a hierarchical object-oriented model of gravitational wave interferometers based in MATLAB[®] which, at the lowest level, treats carrier and sideband optical fields in terms of a finite number of Hermite-Gauss modes (the `laser_field` object class) as well as these fields' interaction with imperfect mirrors (the `mirror` object class) and beamsplitters (the `beamsplitter` class). Transmission through and reflection off of finite-sized, imperfect mirrors and beamsplitters is interpreted in terms of modal scattering matrices, where, given the modal content of the input beam, the modal content of the output beam is determined through simple matrix multiplication. Thermal distortion matrices in `mirror` objects are calculated from the Hello and Vinet analytic solution discussed above [15, 16], while the thermal distortion matrices in the `beamsplitter` object are computed using the three dimensional finite element model discussed developed in this work and discussed in the previous section. On top of these base objects are constructed higher level classes which model the behavior of Fabry-Perot cavities (the `fpi` class, which contains a `laser_field` circulating between two spatially separated `mirror`'s), a Michelson interferometer with Fabry-Perot arm cavities (the `michelson_fp` class, containing a `laser_field` propagating off of a `beamsplitter`, and reflecting off of `fpi`'s for the end mirrors), and a LIGO class (a `laser_field`, circulating between a `mirror`, and a `michelson_fp`). The model numerically determines the steady-state fields in a perturbed object by propagating the initial state `laser_field` object everywhere throughout the perturbed object until a stable modal content is achieved. In order to guarantee the convergence of the fields

¹available at the URL <http://www.phys.ufl.edu/LIGO/LIGO/STAIC/SOFT/>

in the perturbed model, the model is iterated from near-zero power, where there are no thermal perturbations and the fields everywhere can be accurately determined analytically, over small steps of increasing input power. The entire Melody model is written as a collection of simple MATLAB® scripts, which allow the user to carefully tailor the simulation to the particular instrument design as well as the quantities of interest.

In our case, we are concerned with simulating Initial LIGO, Sapphire Advanced LIGO, and Silica Advanced LIGO with the instrument parameters listed in tables A.3 and A.4. The quantities of interest, as found from the analysis of the previous chapter, are the following:

$P_{\text{PRC}}^{(sb0)}$ The sideband power circulating in the recycling cavity which overlaps the carrier is what the gravitational wave signal is read out with, and has a direct impact on the detected phase noise.

$P_{\text{PRC}}^{(c)}$ The carrier power leaking out the dark port will contribute shot noise, but carrier no information on the differential length of the instrument.

$\tilde{\phi}_{min}$ The detected phase noise is directly calculated from the previous two items, and is directly related to gravitational wave strain sensitivity.

G_{ARM} The arm cavity gain determines how many “bounces” a photon undergoes in each arm cavity, and is thus directly related to the instrument’s gravitational wave strain sensitivity.

Additionally, the power recycling cavity gain for the carrier is of interest to us, since this tells us how much power from the laser actually couples into the instrument. From what we found in the previous chapter, it should not be influenced by thermal effects; however, through the Melody simulations we have run, there is a case where it is affected fairly severely in an indirect manner, which we’ll discuss further below.

3.2.1 Initial LIGO

Figure 3-9 shows the results for a Melody model of an “idealized” Initial LIGO, where all of the optics’ curvatures are perfect and optimized for distortion-free operation, the beamsplitter is perfect, and the arms are identical. The dotted line across all graphs indicated the instrument’s nominal operating point, and the straight dashed lines in the top three graphs represent the instrument’s performance in the absence of thermal distortions. Recall from the rough analysis of the previous chapter (table 2.8) that we expect the sideband power in the power recycling cavity to reach a maximum at about 70 Watts of carrier power at the beamsplitter, the power recycling cavity gain for the carrier should not change, and the the arm cavity gain should fall off at 1000 Watts (well above the nominal operating point). The distortion in the beamsplitter is not seen to eclipse the power leaking out of the dark port due to static imperfections (a contrast loss of 10^{-3} is assumed in the simulation), and this large amount of carrier power drives up the phase noise significantly earlier than that predicted by ignoring static defects.

In reality, the Initial LIGO interferometers were designed with thermal distortions in mind, and the recycling mirror was deliberately ground with less curvature so that it would match the effective curvature of the input test mass when thermally loaded. Such a “hot optimized” instrument performs poorly at low powers, as the flattened recycling mirror makes the power recycling cavity unstable for the sidebands here, but should improve as the instrument thermally deforms. Figure 3-10 shows the results for a Melody model of the Livingston 4km interferometer from measured curvatures, lengths, and levels of optical absorption [6] of all the core optics. The total TEM_{00} sideband power in the recycling cavity is seen to reach its maximum slightly farther than the desired operating point (absorption in the optics is lower than what was anticipated), as is the minimum of the phase noise. The arm cavity gain is slightly different between the two arms, as there is about a 1% difference between the transmission and reflection of the beamsplitter as well as a slightly different mode in each arm due to different test mass curvatures.

LIGO I (Idealized)




Figure 3-9: Melody model of an idealized Initial LIGO.

LIGO I (Livingston)




Figure 3-10: Melody model of the Livingston interferometer.

3.2.2 Sapphire Advanced LIGO

Figure 3-11 shows the results for an ideal Advanced LIGO with sapphire test masses, where all of the optics' curvatures are perfect and optimized for distortion-free operation, the beamsplitter is perfect, and the arms are identical. The dotted line across all graphs again indicates the instrument's nominal operating point, and the straight dashed lines in the top three graphs represent the instrument's performance in the absence of thermal distortions. Recall from the rough analysis of the previous chapter (table 2.8) that we expect the sideband power in the power recycling cavity to reach a maximum at about 300 Watts of carrier power at the beamsplitter, the power recycling cavity gain for the carrier should not change, and the the arm cavity gain should fall off at about 2000 Watts (approximately the nominal operating point). The sideband power in the power recycling cavity peaks where expected, but it does not continuously decrease as we originally anticipated. This is due to the fact that the power recycling gain for the sidebands will never decrease to identically zero, but rather to a level equal to the amplitude transmissivity of the recycling mirror, which is about $1 - 0.94^2 = 0.12$ here (sideband power always enters the interferometer in the proper mode, hence there is always at least this much incident on the beamsplitter, even for heavily distorted input test masses). The power exiting the dark port does not deviate from that expected from a static contrast loss (assumed to be 3×10^{-4} here), and caused the phase noise to deviate from the ideal slightly before the nominal operating point. The arm cavity gain decreases as expected, although the carrier gain in the power recycling cavity is seen to decrease at high power. Further experimentation with the model indicates that this is due to increased diffraction losses in the arms (making the mirrors larger makes this feature disappear). As the mirror surfaces of the arms become less concave due to thermal expansion, the spatial size of the mode resonant in the arms increases from a radius of 6.1 to about 6.9 cm (which, incidentally, corresponds well to that which we estimated in table A.6). The larger mode means more light "spills" off the edge of the finite-sized test masses each time it is reflected (about 800 times). Although this loss is small and has little effect on

the arm cavity gain, it has quite an impact on the power recycling cavity gain, since light reentering the power recycling cavity from the arm cavity has experienced this loss 800 times. Although larger beams in the arm cavities are more desirable from the standpoint of thermoelastic noise, there is a tradeoff in that more power must be input into the instrument to compensate this spillover loss and reach the instrument's nominal operating point.

Figure 3-11 shows the results for a Advanced LIGO with sapphire test masses as before, except now $1.5\times$ more power is absorbed in the Y input test mass and $0.5\times$ less power is absorbed in the X input test mass. Thus, in the language of the previous chapter, the common mode thermal distortion is the same, while the differential mode distortion is equal to half the common mode distortion. The performance of the sidebands, which we expected to be primarily dependent on the common mode distortion, is found to be nearly identical to the previous case. The carrier power leaking out the dark port dominates over the static contrast loss and increases as P^3 , as anticipated. Because of the quickly increasing amount of carrier power at the dark port, the phase noise is seen to bottom out at slightly over 1 kW in the power recycling cavity, which is slightly less than the 1.3 kW we expected from figure 2-6. The arm cavity gain is observed to split at about 100 Watts at the beamsplitter, and actually increases for the X arm (where the thermal distortion is smaller), while the mean of the two is identical to the arm cavity gain in the previous symmetric case. This is actually a "modelism", in that the model uses the X arm to define the basis for the rest of the interferometer as well as the mode of the input light, so while power has no trouble coupling into the X arm, it does not couple as well into the more heavily distorted Y arm, and the additional reflected power is then available to couple into the X arm upon reflection from the recycling mirror. Finally, the performance of the power recycling cavity gain is identical to the previous symmetric case.

LIGO II Sapphire (Ideal)




Figure 3-11: Melody model of an idealized Advanced LIGO with sapphire test masses.

LIGO II Sapphire (Asymmetric Absorption)




Figure 3-12: Melody model of a Advanced LIGO with sapphire test masses with differential mode absorption.

3.2.3 Silica Advanced LIGO

Figure 3-13 shows the results for an ideal Advanced LIGO with fused silica test masses, where all of the optics' curvatures are perfect and optimized for distortion-free operation, the beamsplitter is perfect, and the arms are identical. The dotted line across all graphs again indicates the instrument's nominal operating point, and the straight dashed lines in the top three graphs represent the instrument's performance in the absence of thermal distortions. The input power in this particular simulation is increased until Melody cannot converge on a stable mode in the instrument, which is seen to happen at about 1.6 kW at the beamsplitter. Recall from the rough analysis of the previous chapter (table 2.8) that we expect the sideband power in the power recycling cavity to reach a maximum at about 50 Watts of carrier power at the beamsplitter, the power recycling cavity gain for the carrier should not change, the the arm cavity gain should begin to fall off at about 400 Watts, and the interferometer should completely fail at about 1200 Watts. The sideband power peaks as predicted, and then rises again for the same reason discussed in the previous section (fresh sideband power is always entering the instrument, even if it is immediately scattered by the massive thermal distortions in the input test masses). The power leaking out the dark port is again dominated by the static contrast loss (again assumed to be 3×10^{-4}) which affects the phase sensitivity slightly earlier than anticipated. The instrument is not found to fail until about 1.6 kW at the beamsplitter, but note that the arm cavity gain falls off significantly before the operating point. Since the dominant form of absorption in a fused silica test mass is in the coating, this decreased arm cavity gain means that less power is absorbed in the input test masses per watt in the power recycling cavity, hence the carrier failure point is removed to the higher power (the 20% decrease in arm cavity gain means the failure point increases by 20%). The power recycling cavity gain for the carrier is seen to remain approximately constant for the entire range of operation and, although the mode size in the arms increases in size as the sapphire case, the larger fused silica mirrors prevent diffraction losses from appreciably affecting the power recycling gain for the

carrier.

Figure 3-13 shows the results for a Advanced LIGO with fused silica test masses as before, except now $1.5\times$ more power is absorbed in the Y input test mass and $0.5\times$ less power is absorbed in the X input test mass. Again the simulation is run until Melody cannot find a stable mode in the interferometer, which is seen to occur at about 800 Watts at the beamsplitter (which, not coincidentally, is 1.5 times smaller than the 1200 Watt predicted failure point and corresponds to where the Y test mass thermal distortion reaches one quarter wave). The sideband performance is identical to the previous symmetric case and the carrier power leaking out the dark port is significantly larger with a P^3 dependence. The phase noise bottoms out at a mere 300 Watts, as predicted by our rough analysis in figure 2-6.

LIGO II Silica (Ideal)




Figure 3-13: Melody model of an idealized Advanced LIGO with fused silica test masses.

LIGO II Silica (Asymmetric Absorption)






Figure 3-14: Melody model of a Advanced LIGO with fused silica test masses with differential mode absorption.

Chapter 4

Heating Ring Thermal Compensation






Figure 4-1: Coordinates for the simple heating ring method of thermal compensation.

For uniform optical absorption we can precisely calculate the resulting transmissive and reflective wavefront distortions on any absorbing optic, and we've seen that the transmissive cylindrically-symmetric distortions of the Input Test Masses have the most significant impact on the sensitivity and controllability of LIGO-like gravitational wave detectors. Given a cylindrically-symmetric distortion caused by absorption of the Gaussian-profiled probe beam, we can conceivably design a simple cylindrically-symmetric radiative heater to remove radial temperature gradients and thus remove the distortion. This chapter discusses the design, optimization, and performance of a fixed, cylindrically symmetric radiative compensator to address thermal distortions in the Power Recycling Cavity of Advanced LIGO.

4.1 Simple Heating Ring

The simplest geometry conceivable is a ring of resistive wire, mounted away from the face of the optic under actuation, centered on the optical (symmetry) axis with a radius larger than the optic's (to avoid hard-edge diffraction scattering out of the probe beam), as diagrammed in figure 4-1. Suppose the resistive ring carries an electrical current such that a total power P_r is radiated, and P_r is sufficiently small

so that the wavelength of the blackbody radiation lies in a region where the optical material is opaque ($\gtrsim 5 \mu\text{m}$ for sapphire, $\gtrsim 3 \mu\text{m}$ for fused silica). Qualitatively, the resultant heating of the optic face is greater for larger radii; so, as we increase P_r from zero, we may raise the optical path distortion at the periphery of the optical aperture enough to effectively “smooth out” the region in the center of the optic, through which most of the probe beam’s power travels.

Recall from Chapter 2 that the performance of the interferometer is well quantified in terms of the common and differential thermal distortion parameters \mathcal{C}_c and \mathcal{C}_d , which both were proportional to the TEM₀₀ power scattered out of the probe beam by common and differential thermal distortions through optics in the Power Recycling Cavity, respectively. As a figure of merit for the **quality** of a given correction through a single optic, we’ll use the TEM₀₀ power scattered out of the probe beam (\mathcal{S}) normalized to the original uncompensated TEM₀₀ scatter (\mathcal{S}_0) and call it the “correction parameter” \mathbb{C} . More precisely:

$$\mathbb{C} \equiv \frac{\mathcal{S}}{\mathcal{S}_0} \equiv \frac{1 - |\langle E_0 | \exp(i(\phi_0 + P_r \phi_c)) | E_0 \rangle|^2}{1 - |\langle E_0 | \exp(i\phi_0) | E_0 \rangle|^2} \quad (4.1)$$

where $\phi_0(r)$ is the original uncompensated phase distortion, P_r is the power radiated by the ring compensator, and $\phi_c(r) = \phi_c(r; R_r, H_r)$ is the phase distortion induced by the ring compensator of radius R_r and height H_r radiating unit power. For a given distortion ϕ_0 , note that the correction parameter is a function of three parameters: ring radius, height, and power. Recalling figure 2-6 from Chapter 2, which diagrams the maximum useful operating power versus the common and differential distortion parameters, we see that thermal compensation must provide a correction parameter of about 10^{-2} to be able to repair the worst-case scenario for a fused silica Advanced LIGO.

Additionally, we must be concerned about the power required of the compensator and the resulting temperature increase of the optic under actuation. As a figure of merit for the **efficiency** of compensation, we’ll examine the mean temperature increase of the optic under compensation of 1 Watt of probe absorption. Recall from

section 2.2 that the mean temperature increase of an optic under “perfect” thermal compensation is given by:

$$\frac{\overline{\Delta T}}{P_a} = \frac{1}{4\pi w^2 \epsilon \sigma T_\infty^3} \approx 21 \text{ }^\circ\text{K/W} \left(\frac{0.05 \text{ m}}{w} \right)^2 \left(\frac{1}{\epsilon} \right) \left(\frac{300 \text{ }^\circ\text{K}}{T_\infty} \right)^3.$$

where ϵ is the optic’s emissivity, w is the waist of the probe beam whose absorption we are attempting to compensate, σ is the Stefan-Boltzmann constant, and T_∞ is the characteristic temperature of the surrounding photon bath.

To compute the compensation-induced phase distortion ϕ_c , and hence the correction parameter $\mathbb{C}(R_r, H_r, P_r)$, we’ll apply the radiation pattern of the simple heating ring to the boundaries of our two-dimensional finite element model (recall section 3.1.2). The power absorbed on the optic’s boundaries are found in Appendix B to be:

$$H_s(r, 0) = \frac{P_r H_r}{2\pi^3} \int_{-\pi}^{\pi} \frac{((R_r - r \cos \phi)^2 + H_r^2)^{\frac{1}{2}}}{((R_r^2 + r^2 + H_r^2) - 2R_r r \cos \phi)^2} d\phi \quad (4.2)$$

on the optic’s face, and:

$$H_s(R_o, z) = \frac{P_r}{2\pi^3} \int_{-\arccos(\frac{R_o}{R_r})}^{\arccos(\frac{R_o}{R_r})} \frac{(R_r \cos \phi - R_o) ((R_r - R_o \cos \phi)^2 + (H_r - z)^2)^{\frac{1}{2}}}{((R_r^2 + R_o^2 + (H_r - z)^2) - 2R_r R_o \cos \phi)^2} d\phi \quad (4.3)$$

on the optic’s outer edge.

4.1.1 Practical Limits on Heating Ring Design

The parameter space of heating ring radii, heights, and powers has practical limits which will allow us to bound the region over which we minimize the correction parameter $\mathbb{C}(R_r, H_r, P_r)$.

In order to prevent the compensator and its associated supports from occluding the full optical aperture, we restrict the radius of the ring R_r to be larger than the radius of the optic. The radius of the ring cannot be excessively large, either, as the core optics are suspended about 50cm off of their optical tables. We anticipate

that regions of large ring radius, however, will be found undesirable since the ring's radiation pattern on the optic's boundaries will be less acute.

Similar to the ring's radius, the ring's height H_r cannot be excessively large, but this should be bounded for the same reason as the radius. The parameter space with respect to the ring's height will be symmetric about the plane of the optic's mid-point ($r, z = -H_o/2$), so we will use this as a lower bound on our search. Although we will probe the region $z < 0$ (where the ring actuates on the outer edge of the optic only), we note that placement of a heating ring here may be difficult and the finite element calculation imprecise due to suspensions, actuators, and restraints mounted in close proximity to the optic's outer edge.

The ring's emitted power P_r is limited by the wavelength at which the material under actuation becomes transmissive. The ring's radiated power is related to its temperature through the Stefan-Boltzmann law:

$$P_r = \epsilon_r \sigma (T_r^4 - T_\infty^4) (\pi D_r) (2\pi R_r)$$

where ϵ_r is the emissivity and D_r is the thickness of the ring. As the ring's power is increased, the temperature is increased, and more power radiated at shorter wavelengths according to the Planck distribution function, the maximum of which is given by Wien's displacement law [24]:

$$\lambda_{\max} T \approx 2.9 \text{ mm}^\circ\text{K}.$$

Thus, the wavelength where sapphire begins to become transmissive, $\sim 5 \mu\text{m}$, corresponds to a ring temperature of $T \approx 600^\circ\text{K}$. The corresponding ring power at this temperature is then approximately:

$$\frac{P_r}{2\pi R_r} \approx 11 \frac{\text{W}}{\text{cm}} \left(\frac{D_r}{0.5 \text{ cm}} \right) \epsilon_r$$

which places a practical upper limit on the ring's radiated power per linear unit length.

4.1.2 Optimum Ring Compensation of Input Test Masses

The correction parameter $\mathbb{C}(R_r, H_r, P_r)$ can now be numerically minimized against the distortion we wish to correct. Since P_r can be easily changed by altering the amount of current flowing through the ring, we examine the behavior of \mathbb{C} over a grid of ring radii and heights, minimizing over P_r at each point. Figure 4-2 shows the correction attainable, the power required for the correction, and the corresponding mean temperature increase of the Input Test Mass for a Advanced LIGO Sapphire Input Test Mass absorbing 1 Watt of probe power. We see that locating the ring along a line defined by:

$$H_r \approx 0.6 R_r - 0.06 \text{ m} \quad (4.4)$$

results in the best correction of about 10^{-2} with a corresponding ITM temperature increase of nearly 100°K per Watt compensated (compared to the $\sim 20^\circ\text{K}$ per Watt compensated for “perfect” thermal compensation of beams of this size). At ring heights just above this trough we find that the ring provides no correction whatsoever, primarily due to the fact that the intensity pattern of the ring on the face of the optic becomes nearly flat here. At heights below the trough we get a lesser correction which is more efficient (a smaller ITM temperature increase and less required ring power). Figure 4-3 shows the corrected transmitted optical path distortion through a compensated ITM for optimal placement of the ring compensator (in the trough), as well as the resulting temperature field for compensating 1 Watt of probe absorption.

Placing the ring anywhere along the edge of the optic (i.e., for any $H_r \leq 0$) yields a very modest correction (about 0.1) with small required power and temperature increase of the optic, so it could be a viable option for correcting the moderate thermal distortion in sapphire (recall that despite sapphire’s large amount of optical absorption, the magnitude of the distortion is greatly reduced by its large thermal conductivity). In general, this mode of correction results in a lower quality correction for larger optic radii as compared to the probe beam waist, as it relies on large temperature gradients on the periphery of the optical aperture to smooth the distortion in the center. As such, this mode of correction on fused silica test masses, already

insufficient to get the two orders of magnitude or so of correction necessary to address the very large thermal lens on account of silica’s low thermal conductivity, provides an even poorer correction as silica test masses can (and will) be made larger to reduce diffraction losses in the arm cavities as well as maintaining the amount of mass necessary to keep radiation pressure noise under control (recall that silica is half as dense as sapphire, so silica test masses must be made physically larger).

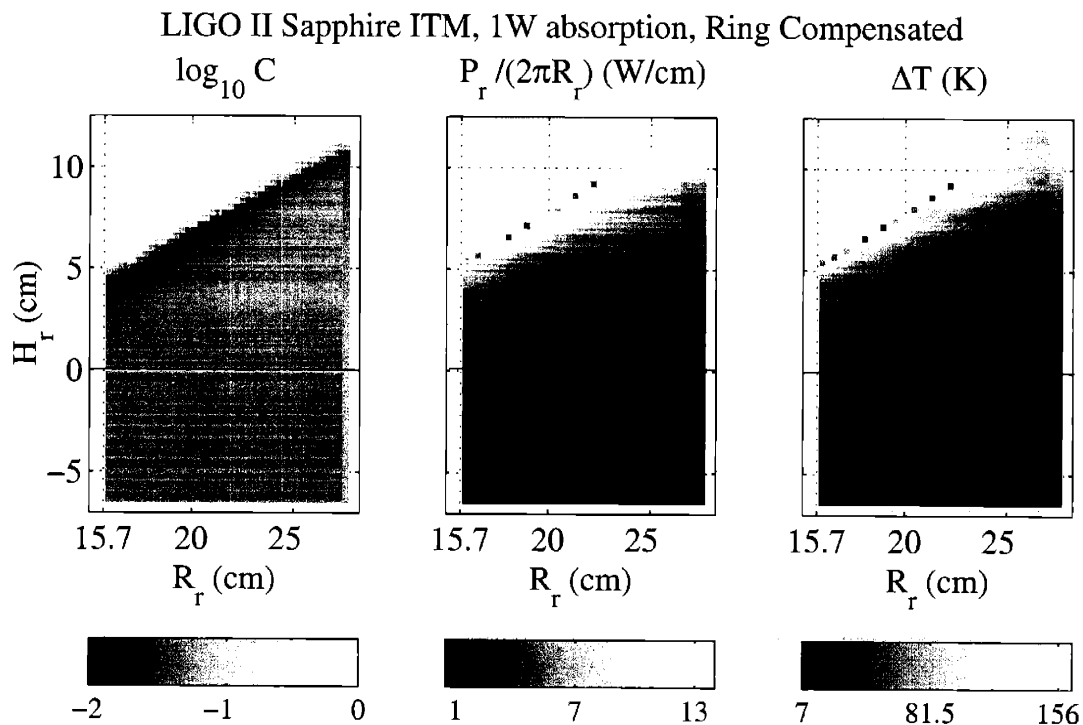


Figure 4-2: The quality and efficiency of simple ring compensation of a Advanced LIGO Sapphire Input Test Mass as a function of ring radii and heights. The black line at $z = 0$ shows the “horizon” of the optic’s face.

Although it is not displayed here, the results for a simple heating ring actuating on a fused silica Advanced LIGO ITM are essentially identical to the sapphire case that we’ve just discussed, except for a slightly lower quality of correction for rings placed along the outer edge of the optic, and the optimal ring placement is still defined by equation 4.4. In fact, in looking at various possible probe beam waists w on LIGO-sized test masses, we heuristically find that the barrier above which ring

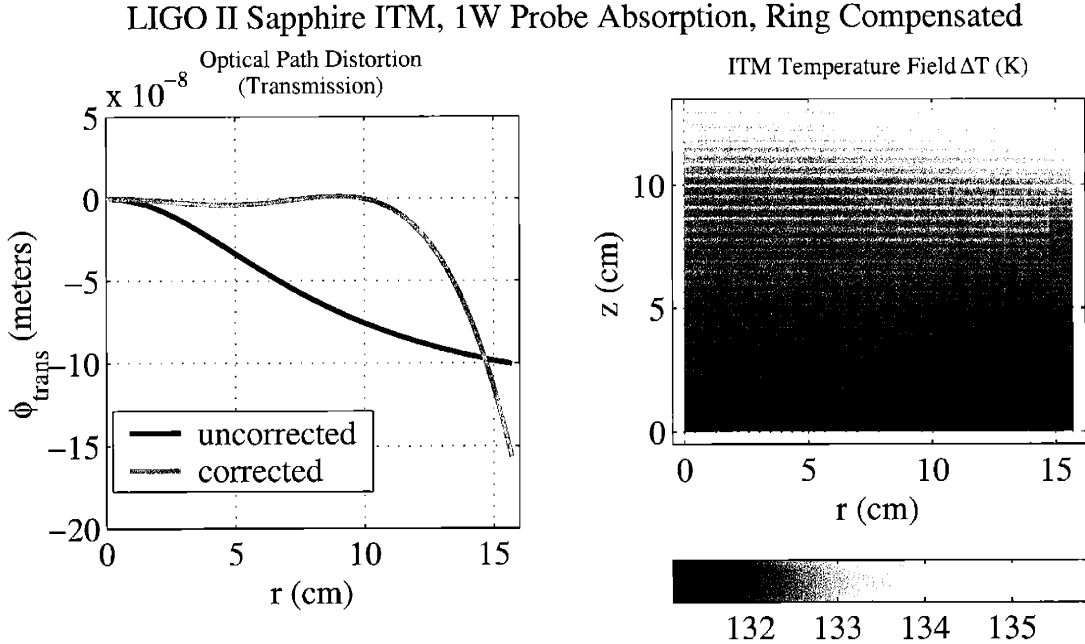


Figure 4-3: The optimal heating ring compensation of a Advanced LIGO Sapphire ITM with 1 Watt of probe absorption. The ring parameters are $R_r = 21.7$ cm, $H_r = 7.6$ cm, and $P_r/(2\pi R_r) = 14.0$ W/cm with a resulting correction parameter of $10^{-2.1}$.

thermal compensation does no good is approximately the line:

$$H_r \approx 0.6R_r - w \quad (4.5)$$

4.1.3 Conclusions

For a Sapphire Advanced LIGO, where we expect about 1.2 Watts of probe power absorbed, we cannot operate the ring at its optimum placement due to the resulting 150°K increase in the ITM's mean temperature. Actuating on the edge is a potentially viable option, due to the low ring power and small ITM temperature increase, although the correction is modest and the degree of accessibility of the edge is low due to the various hardware (mounting, retraining, etc.) that will be arrayed at the edge of the optic. Furthermore, the actual amount of bulk absorption achievable in sapphire remains unknown, and the marginal correction provided by a ring located at the edge of the test mass might not be sufficient if the absorption is larger than

the 30 ppm/cm that we have assumed here.

For a Fused Silica Advanced LIGO we need a correction of at least 10^{-2} (recall figure 2-6), so we are forced operate the bare ring at its optimum placement. With 0.3 Watts of probe power absorbed, a corresponding ITM temperature increase of 50°K would have to be tolerated, which is uncomfortable when considering the increase in thermal noise. Furthermore, operation of the ring at it's optimum location will require very tight tolerances on the dimensions and placement of the ring, due to the sharpness of the parameter space here (recall figure 4-2).

In the end, while the simple heating ring does serve to correct absorption-induced optical path distortions to a fair degree, the quality and efficiency of compensation, as well as the high sensitivity to ring placement, deems simple ring thermal compensating inadequate for either version of Advanced LIGO.

Finally, while it is not discussed further in this thesis, we simply note that using multiple rings to tailor the heating pattern is generally seen to increase the quality of correction; however, the efficiency of correction (the resulting temperature increase of the optic under actuation) is unimproved and remains far from that which we calculated in section 2.2 to be theoretically possible.

Ring Compensation, 1W Probe Absorption					
	R_r	H_r	$\log_{10} \text{C}$	$\frac{P_r}{2\pi R_r}$	ΔT
Highest quality	21.7 cm	7.7 cm	-2.1	14.0 W/cm	134 °K
Most efficient	$r > R_o$	$z < 0$	-1.1	1 W/cm	10 °K

Table 4.1: Optimum ring compensation of Sapphire Advanced LIGO ITM's. The results for Fused Silica Advanced LIGO ITM's are essentially identical.

4.2 Shielded Heating Ring

The clear contributor to the inefficiency of the simple heating ring is the fact that power is deposited in the center of the optic, where the probe absorption which we are attempting to compensate occurs. The obvious method of eliminating this deleterious central heating by the ring is to introduce a cylindrical shield, with a radius equal to




Figure 4-4: Coordinates for the shielded heating ring method of thermal compensation.

the optic's and a height such that the central portion of the optic is shielded from the ring's radiation. Figure 4-4 diagrams this method of compensation. Consider a small area dA on the face of our optic a radius r which is just outside the region completely shadowed by the coaxial shield. This area will see only a sliver of the heating ring, symmetric about the radially outward direction at dA . As we increase r , our small area will see increasingly larger sections of the ring. With this in mind, the calculation of the radiation pattern on the optic's face is identical for that of a simple ring, except that the limits of integration on the angular coordinate in equation 4.2 are now dependent on the radial coordinate r . In Appendix B, we find that for a ring of radius R_r and height H_r , shielded by a coaxial sheath of height H_s and radius R_o , the limits $\phi_l(r)$ of angular integration are found by numerically solving the equation (see Appendix B for the derivation):

$$\frac{R_o^2 + r^2 - 2R_o r \cos \phi'}{R_r^2 + r^2 - 2R_r r \cos \phi_l} = \frac{H_s^2}{H_r^2} \quad (4.6)$$

where $\cos \phi'$ is defined as:

$$\cos \phi' = \frac{r R_r^2}{R_o \rho^2} \sin^2 \phi_l \pm \sqrt{\left(\frac{r R_r^2}{R_o \rho^2} \sin^2 \phi_l\right)^2 - \left(\frac{r}{R_o} + \frac{R_o}{r}\right) \frac{r R_r^2}{R_o \rho^2} \sin^2 \phi_l + 1}.$$

and ρ is defined as:

$$\rho \equiv R_r^2 + r^2 - 2R_r r \cos \phi_l.$$

Clearly, the heating pattern on the optic's outer edge ($r = R_r, z$) is not affected by the coaxial shield.

4.2.1 Optimum Shielded Ring Compensation

Once again, we minimize the correction parameter $\mathbb{C}(R_r, H_r, H_s, P_r)$ against the distortion we wish to correct. To bound the search on the shield height H_s , it is somewhat simpler to view it in terms of the size of the shadow it casts on the face of the optic. We define a dimensionless height parameter h_s as:

$$h_s = \frac{H_s R_r}{R_o H_r}$$

which means that the shield casts no shadow at $h_s = 1$ and fully eclipses the ring at $h_s = 0$. We will minimize $\mathbb{C}(R_r, H_r, h_s, P_r)$ over the ring heights and powers noted in the previous section, now with the additional dimension h_s bounded by $0 < h_s < 1$.

Here we find corrections which are equal in quality to the bare ring, but nearly an order of magnitude more efficient in terms of test mass temperature increase. We further find that the placement of the heating ring is relatively unimportant, provided that it actuates primarily on the face (i.e., $H_r > R_r - R_o$) and the shield is placed at a height such that $0.8 < h_s < 1$. Figure 4-5 shows the optimal correction for a Advanced LIGO Sapphire Input Test Mass.

Now that we've improved the efficiency of the correction, we turn our attention to improving the quality. The limiting factor on the quality of thermal correction is the dissipation of heat out the outer edge of the optic ($r = R_r, z$) which causes the radical downturn in the corrected optical path distortion for large radii (recall the "corrected"

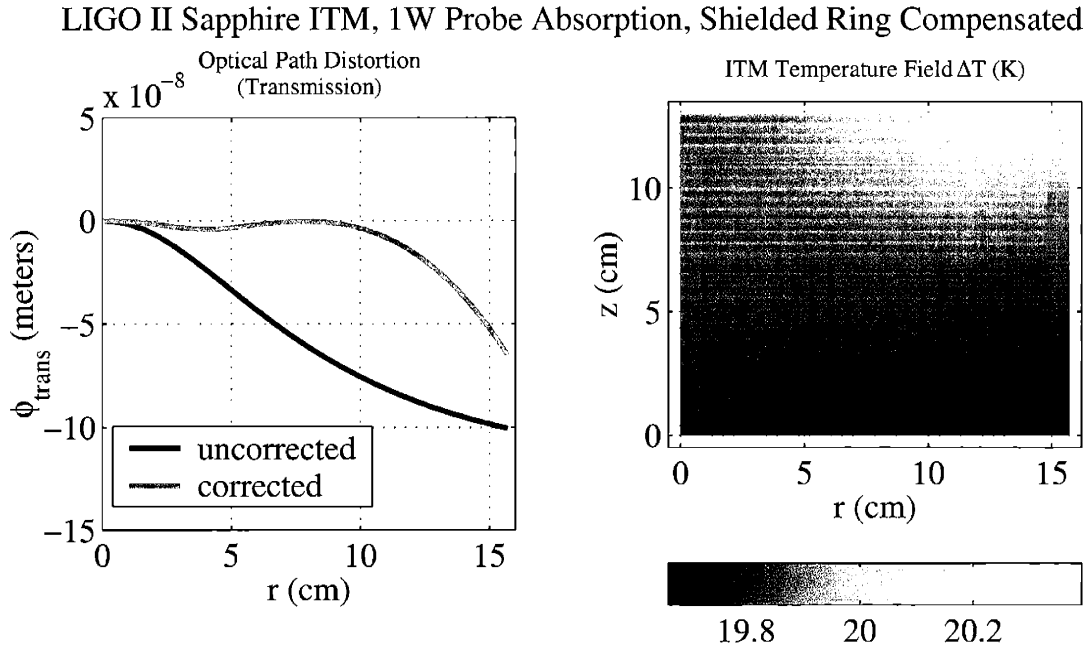


Figure 4-5: The optimal shielded heating ring compensation of a Advanced LIGO Sapphire ITM with 1 Watt of probe absorption. The ring parameters are $R_r = 22.3$ cm, $H_r = 16.0$ cm, $h_s = 0.89$, and $P_r/(2\pi R_r) = 4.1$ W/cm with a resulting correction parameter of $10^{-2.2}$.

curves in figures 4-3, 4-5, and 4-6). We can conceivably repair this by reducing the emissivity of this surface, by wrapping it in a material with low emissivity at room temperature (aluminum and gold are such materials). Coating the edge of a test mass is not an option, due to the corresponding increase in internal loss, hence an increase in test mass thermal noise (remember that more internal loss means more thermal motion of the test mass surfaces at frequencies away from the internal mechanical resonances of the test mass). Suspending an aluminum sheath close to the outer edge without contacting the optic should provide sufficient insulation (energy emitted from the the optic's edge is simple reflected back) while not increasing internal thermal noise. In any case, preventing radiation exchange with the external photon bath means that the boundary condition on the edge of the optic ($r = R_r, z$) is simply $\frac{\partial T}{\partial r} = 0$.

Optimizing \mathbb{C} once again, we find the results shown in figure 4-6. The efficiency of the correction is marginally improved over the uninsulated case, although the quality

LIGO II Sapphire ITM, 1W Probe Absorption, Shielded Ring on Insulated ITM




Figure 4-6: The optimal shielded heating ring compensation of a Advanced LIGO Sapphire ITM with 1 Watt of probe absorption and its outer edge insulated (e.g., zero emissivity). The ring parameters are $R_r = 16.3$ cm, $H_r = 20.5$ cm, $h_s = 0.99$, and $P_r/(2\pi R_r) = 3.9$ W/cm with a resulting correction parameter of $10^{-4.0}$.

of correction is increased significantly. Also note that the optimum placement of the ring is now closer to the optical axis (i.e., smaller ring radius), due to the fact that we no longer need to counteract heat escaping from the outer edge of the optic.

Shielded Ring Compensation, 1W Probe Absorption						
	R_r	H_r	H_s	$\log_{10} \mathbb{C}$	$P_r/(2\pi R_r)$	ΔT
Bare	22.3 cm	16.0 cm	0.89	-2.2	4.1 W/cm	20 °K
Insulated	16.3 cm	20.5 cm	0.99	-4.0	3.9 W/cm	23 °K

Table 4.2: Optimum shielded ring compensation of Sapphire Advanced LIGO ITM's. The results for Fused Silica Advanced LIGO ITM's are essentially identical.

4.2.2 Optimum Shielded Ring Compensation acting on a Compensator Plate

Until now we have only considered actuating directly on a thermally aberrated optic. Due to the linearity of absorption-induced wavefront distortions and the fact that the

beams propagating in the power recycling cavity of the interferometer are very large and well-collimated (the Rayleigh range is on the order of the 4 kilometer arm length), as well as the fact that the wavefront distortions addressing here are comparatively very small (smaller than $\lambda/4$ over the entire aperture), we can conceivably suspend a separate, purely transmissive “compensation plate” in each arm of the Power Recycling Cavity on which we thermally actuate to correct transmissive distortions seen in the input test masses. In essence, the compensation plates serve as transmissive “eyeglasses” to repair the transmissive thermal distortions in each arm of the Power Recycling Cavity (test mass and beamsplitter). See figure 4-7 for a diagram of thermally actuating on compensation plates in Advanced LIGO. The primary benefit of using a compensation plate is that the Input Test Mass remains untouched: we don’t need to worry about perturbing the test mass with intensity fluctuations from our thermal compensator, nor do we need to worry about increasing the temperature, and hence thermal surface fluctuations, of the test mass. The main drawback of using compensator plates is that the arm cavity surfaces of the ITM’s are free to thermally distort and thus change the spatial mode that is resonant in the arm cavity (recall that it is made larger). As we saw in the case of a sapphire Advanced LIGO in Chapter 2, this only becomes a problem when the arm cavity mode becomes large enough that the power which “spills” over the test mass edges upon each of the 800 or so bounces light undergoes in the arm cavity becomes comparable to all other losses in the interferometer, and can be fixed by simply making the test mass mirrors larger in diameter.

To choose the size of the fused silica compensator plate, first note that we don’t want to increase diffraction losses, so we’ll make the compensation plate’s radius equal to the radius of the recycling mirror (remember that we only need very large radii for test masses due to the fact that light bounces ~ 800 times off of each test mass before returning to the recycling cavity). We also don’t want substrate absorption in the compensator plates to contribute significantly to the distortion in each arm (a distortion which we’ll have to remove with our ring compensator), so we’ll want to make them as thin as possible. Recall from Chapter 2 that thickness of an absorbing

optic has little effect on the transmissive distortion due to a fixed amount of power, so we expect that a thin optic will, under actuation by a shielded ring, yield a very similar corrective optical path distortion as compared to a thicker optic. The main change in the corrective optical path distortion will be due to the obvious fact that less heat is able to escape out the outer edge of a thinner optic than a thicker one. Since less heat escaping out the edge was previously seen to be beneficial for the quality of thermal correction, we conclude that a thinner compensation plate is better overall. The thinness of the compensator plate is limited by the ability to grind and polish both surfaces to high precision, which diminishes greatly for very thin optics. The necessary precision has been demonstrated for Initial LIGO beamsplitters which are 4 cm thick, so we will use this as our compensation plate thickness.




Figure 4-7: Shielded ring thermal compensation in Advanced LIGO.

Figure 4-8, shows the optimum correction for a distorted Advanced LIGO Sapphire

LIGO II Sapphire ITM, 1W Probe Absorption Shielded Ring on Compensation Plate

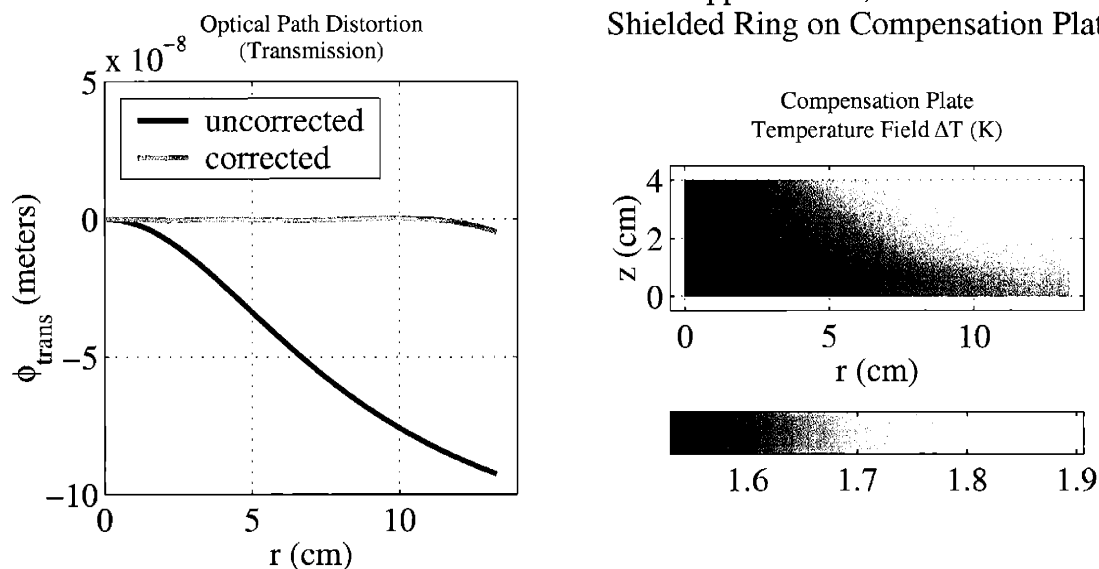


Figure 4-8: The optimal shielded heating ring compensation of a Advanced LIGO Sapphire ITM with 1 Watt of probe absorption by actuating on a fused silica compensation plate. The ring parameters are $R_r = 24.5$ cm, $H_r = 14.0$ cm, $h_s = 0.90$, and $P_r/(2\pi R_r) = 0.17$ W/cm with a resulting correction parameter of $10^{-4.1}$.

Input Test Mass achieved by actuating on a fused silica compensation plate of radius 13.3 cm and height 4 cm. The efficiency of the correction much improved due to the much lower thermal conductivity of fused silica. Table 4.3 shows the results for the optimum shielded heating ring acting on a fused silica compensator plate to correct the transmissive thermal distortion in both sapphire and fused silica Advanced LIGO ITM's. The ring parameters are slightly different for fused silica due to the dominance of the thermal lens over thermoelastic deformation, and the efficiency is not improved (but is still excellent) due to the fact that the compensation plate is also made out of fused silica.

Shielded Ring on Silica Compensation Plate, 1W Probe Absorption						
	R_r	H_r	H_s	$\log_{10} \mathcal{C}$	$P_r/(2\pi R_r)$	ΔT
Sapphire	24.5 cm	14.0 cm	0.90	-4.1	0.17 W/cm	2 °K
Fused Silica	21.3 cm	12.0 cm	0.92	-3.5	2.1 W/cm	17 °K

Table 4.3: Optimum shielded ring compensation of Advanced LIGO ITM's by actuating on a Fused Silica compensation plate.




Figure 4-9: The correction achieved by displacing the heating ring from its optimum for the shielded ring actuating on a silica compensation plate. The height of the shield H_s is kept constant and the ring power P_r is optimized at each ring radius R_r and height H_r .

There most certainly will be some error in the actual placement of the ring heater, and the finite width of the heating element may degrade the quality of the correction (a heating element of finite width can be seen the sum of many heaters with infinitesimal width located at the outside boundary of the heating element). To see how well we must place the ring, as well as how thick the ring can be, we fix the shield height H_s at the optimum and vary the infinitesimal ring radius R_r and height H_r over a fixed grid about the global optimum, optimizing \mathbb{C} over P_r at each point (R_r, H_r) . Figure 4-9 shows the resulting surface $\log_{10} \mathbb{C}$ versus ring coordinate (R_r, H_r) , and we see that displacing the heating ring position about 1 cm of the optimum placement retains the ability to provide a correction of $\mathbb{C} < 10^{-2.5}$.

4.2.3 Conclusions

The shielded ring provides a correction which is higher in both efficiency and quality than that of the bare ring, due to the fact that the shielded ring does not add heat to the areas where the optical absorption we wish to correct occurs. Insulating the

radial edge of the optic provides an even higher quality correction, as it prevents the radial escape of heat from the optic. Finally, actuating on purely transmissive fused silica compensator plates in the Power Recycling Cavity allows us to repair thermal distortions in that cavity without perturbing the Input Test Masses in any way, as well as providing an extremely efficient correction in a sapphire interferometer due to the lower thermal conductivity of the fused silica. Furthermore, radial insulation of the compensator plates is unnecessary to attain the highest quality correction as we are free to make them relatively thin, thus hindering the deleterious radiative heat exchange on the radial edge of the plate.

Heating Ring Thermal Compensation Comparison			
Method	$\log_{10} \text{C}$	$P_r/(2\pi R_r)$	ΔT
Bare Ring (most efficient)	-1.1	1 W/cm	10 °K
Bare Ring (highest quality)	-2.1	14 W/cm	130 °K
Shielded Ring	-2.2	4.1 W/cm	20 °K
Shielded Ring, Insulated Optic	-4.0	3.9 W/cm	23 °K
Shielded Ring on Thin Silica Compensation Plate	-4.1	0.2 W/cm	2 °K

Table 4.4: Comparison of the various methods of heating ring thermal compensation to address the transmissive thermal distortion in a Advanced LIGO sapphire Input Test Mass.

4.3 Effects on the Entire Interferometer

To predict the effects of thermal compensation on the Advanced LIGO interferometers (Sapphire and Fused Silica versions), we again turn to the Melody modal model, which we used in Chapter 2 to examine the performance of each instrument as the input laser power, hence optical absorption, is increased. To insert thermal compensation into Melody, we simply calculate the scattering matrices for the thermal-compensation-induced wavefront distortion (i.e., the wavefront correction) necessary to compensate unit probe absorption and add them onto the existing matrix elements representing the original wavefront distortion in each test mass due to unit probe absorption. In

what follows, we'll examine the use of compensator plates in the power recycling cavity to remove the transmissive thermal distortions through input test masses. Due to the relatively very short length of the power recycling cavity, we neglect the propagation between the compensator plates and the input test masses, and simply add the compensation plate matrix elements for total transmission (taking into account the thermo-optic effect as well as thermoelastic deformation of the plate) onto the input test mass matrix elements for the thermo-optic effect. We don't alter the matrix elements for thermoelastic deformation in the test masses, since we are not actuating on them (hence not distorting the surfaces), although we do take into account the effect that thermoelastic deformation in the test masses has transmission when calculating the required wavefront correction in the compensation plates.

4.3.1 Sapphire Advanced LIGO

Figure 4-10 shows the results for thermal compensation by a shielded ring actuating on a fused silica compensation plate in the power recycling cavity of an ideal Advanced LIGO with sapphire test masses, as diagrammed in figure 4-7. As in §3.2, where we considered the uncompensated performance of the identical interferometer, all of the optics' curvatures are perfect and optimized for distortion-free operation, the beamsplitter is perfect, and the arms are identical. Note that the performance of the instrument is now nearly perfect in all regards, except for the drop in power recycling cavity gain due to the additional diffraction losses of the enlarged modes in the arms caused by absorption-induced thermal expansion of the test mass faces. An additional amount of compensation (about 5% more) has been added into the X arm to compensate the beamsplitter thermal lens, and the power leaking out of the dark port is seen to be reduced by ~ 100 . There is a slight upturn in the arm cavity gain due to the fact that the model optimized the input field for the X arm cavity, and arm cavity gain is measured by comparing the carrier power stored in the arm with the carrier power stored in the power recycling cavity, which is reduced due to the increased diffraction losses in the arm.

Figure 4-11 shows the results for a thermally compensated Advanced LIGO with

sapphire test masses as before, except now $1.5\times$ more power is absorbed in the Y input test mass and $0.5\times$ less power is absorbed in the X input test mass (i.e., the common mode thermal distortion is the same, while the differential mode distortion is equal to half the common mode distortion). Again, the performance is near perfect, except for the drop in power recycling gain that we discussed before, and the large amount of power leaking out the dark port, only slightly less than that of the uncompensated case. The mode of each arm is defined by the interior surfaces of the test masses which are now distorting differently between the arms, and the surfaces of the test masses are unaffected by this method of thermal compensation (we're actuating on completely separate compensator plates). In general, the power leaking out of the dark port due to a mode mismatch between the arms cannot be addressed by this method of thermal compensation.

LIGO II Sapphire (Compensated)




Figure 4-10: Melody model of a thermally compensated idealized Advanced LIGO with sapphire test masses. See figure 3-11 for the uncompensated performance.

LIGO II Sapphire (Compensated, Asymmetric Absorption)




Figure 4-11: Melody model of a thermally compensated Advanced LIGO with sapphire test masses and differential absorption. See figure 3-12 for the uncompensated performance.

4.3.2 Silica Advanced LIGO

Figure 4-12 shows the results for thermal compensation by a shielded ring actuating on a fused silica compensation plate in the power recycling cavity of an ideal Advanced LIGO with fused silica test masses, as diagrammed in figure 4-7. As in section 3.2.3 of Chapter 3, where we considered the uncompensated performance of the identical interferometer, all of the optics' curvatures are perfect and optimized for distortion-free operation, the beamsplitter is perfect, and the arms are identical. The power recycling cavity gain for the sidebands begins to slowly fall off at about 1 kW in the recycling cavity, due to the fact that such a large amount of compensation is required that the hardly noticeable "ripple" in the corrected transmissive optical path distortion similar to that seen for sapphire in figure 4-8 begins to become large enough to cause an appreciable loss for the sidebands. The loss is small enough, however, that increasing stored sideband power is still seen with increasing carrier power past 2 kW stored in the power recycling cavity. As in the case for the compensated sapphire interferometer, a slightly larger amount of compensation is introduced into the X arm to compensate the beamsplitter thermal lens (about 1% in this case, which is smaller than the sapphire case due to the fact that we're now compensating a $\sim 5\times$ larger transmissive thermal distortion in the silica test masses), and the carrier leakage out the dark port is again reduced by a factor of about 100. The phase noise and arm cavity gains are perfectly maintained in the interferometer, and we see a slight decrease in the power recycling gain for the carrier again due to the enlarging of the arm cavity mode by thermal expansion of the arm cavity faces (this drop-off is less significant than the sapphire case because the silica test masses are $\sim 6\text{cm}$ larger in diameter).

Figure 4-13 shows the results for a thermally compensated Advanced LIGO with fused silica test masses as before, except now $1.5\times$ more power is absorbed in the Y input test mass and $0.5\times$ less power is absorbed in the X input test mass (i.e., the common mode thermal distortion is the same, while the differential mode distortion is equal to half the common mode distortion). As in the sapphire case discussed above,

the compensated performance is essentially identical to the symmetric case, except for the relatively unaffected amount of carrier leakage out of the dark port due to the mode mismatch between the arms which this method of thermal compensation cannot address.

LIGO II Silica (Compensated)




Figure 4-12: Melody model of a thermally compensated idealized Advanced LIGO with fused silica test masses. See figure 3-13 for the uncompensated performance.

LIGO II Silica (Compensated, Asymmetric Absorption)




Figure 4-13: Melody model of a thermally compensated Advanced LIGO with fused silica test masses and differential absorption. See figure 3-14 for the uncompensated performance.

4.4 Noise Considerations

As we discussed in section 2.3, radiative thermal compensation acts to alter the optical path through the interferometer, so we must consider how random intensity fluctuations in our compensator induce random path length fluctuations in the instrument and make certain that they will not dominate over the path length fluctuations we wish to detect. The root mean square fluctuation in photon number Δn_{ω_0} for a single field mode in an ideal blackbody source of temperature T is given by [24]:

$$\Delta n_{\omega_0} = \sqrt{\bar{n}_{\omega_0}^2 + \bar{n}_{\omega_0}}$$

where n is the mean photon number in the mode, given by the Planck thermal excitation function $\bar{n}_{\omega_0} = 1/(e^{\frac{\hbar\omega}{k_B T}} - 1)$. For $\bar{n}_{\omega_0} > 1$, the fluctuations are approximately equal to the mean number of photons in the mode, while for $\bar{n}_{\omega_0} < 1$ the statistics become Poissonian. The wavelength where $\bar{n}_{\omega_0} = 1$ is calculated to obey the relation:

$$\lambda_c T = 0.02 \text{ m}^\circ\text{K}.$$

Utilizing Wien's displacement law, $\lambda_{\max} T = 0.003 \text{ m}^\circ\text{K}$, we arrive at the general relation:

$$\frac{\lambda_c}{\lambda_{\max}} = 6.67.$$

Using this result, and integrating the blackbody energy density, it can be shown that 1.6% of the total radiated energy resides in the non-Poissonian tail of the Planck distribution.

In a realistic blackbody source, the effect of spatial extent of the source must be taken into account to properly calculate the total statistical fluctuation. The coherence length of the emitted light, l_c , (approximately equal to the wavelength of the mode for blackbody light) defines the scale of coherence between spatially separated radiators. Light emitted from separated "coherence patches" (patches of area l_c^2) will thus be statistically independent. An identical argument applies to the "detector" (the face of the optic under actuation), since radiation from a single coherence patch

arriving at spatially separated positions of the optic will be uncorrelated if the difference in travel time is longer than the coherence time $t_c = l_c/c$. For N_c coherence patches (detector plus radiator), the mean photon number in a given field mode becomes $\bar{n}_\omega = N_c \bar{n}_{\omega 0}$. Similarly, the root mean square is $(\Delta n_\omega)^2 = N_c (\Delta n_{\omega 0})^2$, thus yielding:

$$\frac{\Delta n_\omega}{n_\omega} = \frac{\sqrt{N_c(\bar{n}_{\omega 0}^2 + \bar{n}_{\omega 0})}}{N_c \bar{n}_{\omega 0}} \quad (4.7)$$

$$= \sqrt{\frac{1}{N_c} + \frac{1}{\bar{n}_\omega}}. \quad (4.8)$$

Note that for $\bar{n}_\omega < N_c$, the fractional photon fluctuation is closely Poissonian, while it is a constant $1/N_c$ otherwise. The new λ_c (where $\bar{n}_\omega = N_c$) is calculated to be:

$$\frac{\lambda_c}{\lambda_{\max}} = 6.67 \times \log(N_c + 1). \quad (4.9)$$

Since $N_c = A/l_c^2$ depends on the coherence length $l_c \sim \lambda_c$, this equation becomes transcendental in λ_c . We are forced to numerically solve for λ_c given area emitter/detector area A and emitter temperature T .

For our radiator/detector system, the typical peak wavelength is $\lambda_{\max} \simeq 7 \mu\text{m}$, with a total detector/emitter area of $A \simeq 0.05 \text{ m}^2$. Solving equation (4.9) for λ_c , we find the number of disjoint (statistically independent) coherence patches to be $N_c \sim 1 \times 10^5$, which results in 0.0013% of the total emitted power in the non-Poissonian tail. Thus, for the geometry and wavelengths we are concerned with, the counting statistics on radiated photons may be considered purely Poissonian.

At each field mode ω in the blackbody emission spectrum of our ring compensator, the root mean square fluctuation in the “detected” power $p_d(\omega)$ (i.e., the power absorbed on the surface of the optic per unit bandwidth) is then:

$$\Delta p_d(\omega) = \sqrt{2\hbar\omega p_d(\omega)}$$

and the total detected power fluctuation is simply the sum in quadrature of the

single-mode fluctuations over all modes, i.e.,:

$$\Delta P_d = \sqrt{\int (\Delta p_d(\omega))^2 d\omega} = \sqrt{2\hbar \int \omega p_d(\omega) d\omega}.$$

Utilizing the fact that $p_d(\omega)$ follows the Planck distribution and the total detected power is P_d , i.e.,:

$$p_d(\omega) \propto \frac{\omega^3}{e^{\frac{\hbar\omega}{k_B T}} - 1} \quad \text{and} \quad P_d = \int p_d(\omega) d\omega$$

the total absorbed power fluctuation from the ring compensator is found to be:

$$\Delta P_d \approx \sqrt{2P_d(4k_B T)} \tag{4.10}$$

The broadband ripple \mathcal{R}_{comp} on the detected thermal actuator power P_d , as defined in section 2.3, is now:

$$\mathcal{R}_{comp} \equiv \frac{\Delta P_d}{P_d} = \sqrt{\frac{8k_B T}{P_d}}$$

where T is the temperature of the thermal actuator.

Nominally, we would now utilize equation 2.24 and integrate the compensator's intensity pattern against the probe beam profile to recover the detected displacement noise. This must be done numerically, however, so instead we'll note that the shielded ring's intensity profile at the optimum behaves in a very similar manner to the "perfect" correction of power P_a absorbed from a Gaussian beam of waist radius w :

$$I_{ring}(r) \approx \frac{2P_a}{\pi w^2} \left(1 - e^{-2\frac{r^2}{w^2}}\right)$$

and integrate the probe beam profile against this. This was done back in section 2.3 where we found that, if we are actuating on the face of the arm cavity (the most sensitive place in the interferometer), the displacement noise will not be appreciably affected if the ripple on the compensator's profile is less than the ripple on the


absorbed probe power, i.e.:

$$\mathcal{R}_{comp} = \sqrt{\frac{8k_B T_r}{P_d}} \leq \mathcal{R}_{probe} = \sqrt{\frac{2\epsilon}{P_a}}$$

where $\epsilon \approx 2 \times 10^{-19}$ J is the energy of a single probe photon. In general, a necessary condition for the compensation of a Gaussian beam whose waist is smaller than the mirror radius is that the absorbed compensator power is greater than the absorbed probe power, i.e., $P_d > P_a$ (which is obvious by inspection of the approximate intensity profile I_{ring} listed above). Also, we must have $k_B T \ll \epsilon$, for otherwise the compensator is emitting in the transmissive band of the optical material, which is a condition which we have already designed our compensator to avoid. Thus, $\mathcal{R}_{comp} < \mathcal{R}_{probe}$ is always true for shielded ring thermal compensation, and we may utilize the shielded ring *anywhere* in the interferometer without increasing the detected displacement noise above the ever-present photothermal noise induced by probe beam absorption.

Chapter 5

Scanning Laser Thermal Compensation



Up until now we've discussed *nominal* thermal distortions in high power gravitational wave interferometers, where the shape of the resulting wavefront distortions are predictable and essentially cylindrically symmetric. As such, we've seen that these anticipated thermal distortions can be largely repaired by constructing radiative thermal corrections through the careful placement of a radiative heating rings and shields. There is the strong possibility, however, that optical absorption in core optics will be inhomogeneous to some unknown degree, whether it be due to a small flake of absorptive material bonding to the face of an optic subjected to high optical power, or the inhomogeneous substrate absorption that has been observed in sapphire substrates [2]. This chapter discusses thermally correcting arbitrary wavefront distortions by tailoring the amount of power deposited by a radiative compensator over the face of an affected optic (or, equivalently, on a separate compensation plate as discussed in the previous section). For the radiative corrector we'll use a laser beam at a wavelength chosen for strong surface absorption in the material under actuation, which we'll rapidly scan in a fixed pattern over a single face of the actuated optic and change the amplitude of the laser at each point in the scan pattern according to what is necessary to generate the desired wavefront correction. The two key problems that must be solved are thus:

1. **How do we determine the “Actuation Basis”?** In other words, given a wavefront correction we wish to generate, how do we determine how much laser power must be deposited at each point in the scan pattern?
2. **What is the “optimal” scan pattern?** What should the scan resolution be, and what size of beam is required for such a resolution?

5.1 Finding the Actuation Basis

Given a wavefront correction we wish to generate, we need to find the power that our laser must deposit at each point in the scan pattern to generate it. We can easily go the other way: given the powers and pattern of the laser actuator, we can

compute the induced wavefront distortion with the finite element models discussed in Chapter 3. Now we must work backwards, which is perhaps most easily done in terms of a mapping (i.e., a matrix) linking the M dimensional vector space representing the actuator power at each of its M scan points with the N -dimensional vector space representing wavefront distortions decomposed in terms of N linearly independent functions $\{Z_n(x, y)\}$ (e.g., Zernike polynomials). Provided the choice of scan points and representation functions is done properly, the mapping between these vector spaces (originally, scan powers \rightarrow wavefronts) will be invertible (yielding wavefronts \rightarrow scan powers).

More precisely, with the three-dimensional finite element model discussed in Chapter 3 we can readily compute the steady-state wavefront distortion resulting from shining a Gaussian actuator beam at unit power and constant waist radius w_a :

$$I_m(x, y) = \frac{2P_m}{\pi w_a^2} e^{-2\frac{(x-x_m)^2+(y-y_m)^2}{w_a^2}}$$

on each separate point (x_m, y_m) in the scan pattern. Let $\mathcal{I}_m(x, y)$ denote the wavefront distortion resulting from unit power deposited in the m th point in the scan pattern (termed hereafter the “influence functions”). The total induced wavefront distortion $\phi(x, y)$ given a list of powers $\{P_m\}$ is then simply:

$$\phi(x, y) = \sum_{m=1}^M P_m \mathcal{I}_m(x, y) \quad (5.1)$$

Utilizing some basis $\{Z_n(x, y)\}$ of smooth, linearly independent functions chosen to accurately represent the wavefronts we wish to generate over the optical aperture, we can write:

$$\phi(x, y) = \sum_{n=1}^N c_n Z_n(x, y) \quad \text{and} \quad \mathcal{I}_m(x, y) = \sum_{n=1}^N a_{nm} Z_n(x, y). \quad (5.2)$$

Inserting into equation 5.1:

$$\sum_{n=1}^N c_n Z_n(x, y) = \sum_{m=1}^M P_m \sum_{n=1}^N a_{nm} Z_n(x, y)$$

which is equivalent to the matrix equation:

$$\vec{c} = \vec{P} \cdot \underline{a} \quad (5.3)$$

where \vec{c} is a vector of length N consisting of amplitudes of the linearly independent functions $\{Z_n(x, y)\}$, \vec{P} is a vector of length M representing the power deposited at each scan point in the pattern, and \underline{a} is an $M \times N$ matrix whose m th column represents the amplitudes of the decomposition of the m th influence function \mathcal{I}_m in terms of the basis functions $\{Z_n(x, y)\}$. We are able to fully calculate the elements of \underline{a} with our existing finite element models.

Equation 5.3 represents a compact method of computing the resulting wavefront distortion given a set of powers in the scan pattern. To go “backwards”, i.e., retrieving the required powers to generate a given wavefront distortion, we must find an $N \times M$ matrix \underline{A} such that $\underline{a} \cdot \underline{A}$ is equal to the $M \times M$ identity matrix, so that:

$$\vec{c} \cdot \underline{A} = \vec{P} \cdot \underline{a} \cdot \underline{A} = \vec{P} \quad (5.4)$$

In general, such a matrix \underline{A} exists if $N \leq M$ and the full rank of \underline{a} is N . In plain English, this means that the number of functions in the decomposition $\{Z_n\}$ must be less than or equal to the number of influence functions $\mathcal{I}_m(x, y)$, and the actuation functions have to look sufficiently different from each other (they all must be linearly independent of each other, meaning that any one cannot be accurately expressed in terms of a linear combination of all the others). Provided these conditions are met, then the inverse matrix \underline{A} can be found from the known matrix \underline{a} using any number of different numerical methods (see, for example, Chapter 2 of [34] and Chapter 4 of [26]).

We now know how to calculate $\{P_m\}$ given $\phi(x, y)$, for arbitrary scan patterns

and arbitrary basis functions. There is no mathematical constraint, however, that the P_m we find will be greater than zero, which is a constraint that is most certainly enforced in reality (we can't actuate on an optic with negative power). In fact, we can guarantee that we can find neither a set of basis functions $\{Z_n\}$ nor a scan pattern $\{(x_m, y_m)\}$ that will make $\{P_m\}$ exclusively positive, since, by the formulae above, if we generate a wavefront $\phi(x, y)$ such that $\{P_m\} > 0$, then the conjugate wavefront $-\phi(x, y)$ necessarily requires $\{P_m\} < 0$.

Making \vec{P} Positive

The most direct way to bypass this is to utilize the fact that a deformation which is constant over the aperture (typically termed "piston") is generated by simply displacing the optic we wish to compensate. Thus, we may add (or subtract) an arbitrary amount of piston to the requested wavefront correction in order to force the powers $\{P_m\}$ to be purely positive. More precisely, if $\vec{c}_0 = \{c_n^{(0)}\}$ is the representation of unit piston over the basis functions $\{Z_n\}$, i.e.:

$$1 = \sum_{n=1}^N c_n^{(0)} Z_n(x, y)$$

then the corresponding actuator powers necessary to generate piston (termed the "piston power vector") are simply:

$$\vec{P}_0 = \vec{c}_0 \cdot \underline{A}$$

where the actuation matrix \underline{A} is determined as outlined before. The behavior of \vec{P}_0 is thus exclusively dependent on the choice of basis functions and scan pattern. We then add some constant amount C of the piston power vector \vec{P}_0 to the requested power vector $\vec{P} = \vec{c} \cdot \underline{A}$ so that the resulting powers are all greater than zero. In other words, we choose C such that:

$$\min_m (P_m + CP_m^{(0)}) = 0$$

and the power vector we use in reality is then $\vec{P} + C\vec{P}_0$. Note that this means that some point k in the scan pattern will have $P_k = 0$.

A sufficient condition that ensures this “fix” will always work is to have all of the elements of the piston power vector nonzero, identical in sign, and comparable in magnitude. The behavior of the piston power vector \vec{P}_0 depends heavily on the choice of the actuator beam waist and scan pattern. For a uniformly spaced scan pattern covering the entire face of the actuated optic and a carefully chosen beam size, the elements of \vec{P}_0 will all be comparable in magnitude and positive, and the shape of the induced piston wavefront looking flat. If the beam size is too small, then the quality of the DC offset in our correction will be poor, as we’ll generate a “pincushion” phase distortion instead of flat piston, which will generally reduce the overall quality of compensation. If the beam is chosen too large, however, then the influence functions at neighboring points look nearly identical. This near degeneracy forces the influence function decomposition matrix \underline{a} close to singular (i.e., uninvertible) and the elements of the corresponding actuation operator \underline{A} thus begin to vary widely over large ranges of values. For a given scan pattern, we’ll vary the waist of the actuator (or vice-versa) until \vec{P}_0 becomes well-behaved. If the scan pattern is chosen poorly, then this task could be impossible.

5.1.1 Choosing a Scan Pattern

Until now we’ve made no assumptions about the pattern of points $\{(x_m, y_m)\}$ at which the compensator beam is directed. As discussed in the previous paragraph, we must be able to generate a DC offset in phase (piston), in order to avoid negative powers required to correct certain wavefront distortions. As such, scan pattern should be reasonably uniform and the actuator beam waist should be sufficiently small to generate clean piston while maintaining sufficient distinction between neighboring actuation functions. Consideration must be given to the optical aperture which we are concerned with (i.e., within the probe beam waist), which will be considerably smaller than the radius of the optic under actuation. Actuating at points outside of the circular aperture lying near the same radial ray will all have nearly identical




Figure 5-1: Different types of scan patterns.

effects within the circular aperture: tilt with a little astigmatism along that radial ray. To avoid degeneracy, we must be careful to include only a few points outside of the aperture of concern at well-separated radial angles in the scan pattern. Finally, we must consider the load on the mechanical mirrors which will be used to direct the compensator beam from point to point. The scan path through the entire pattern should be minimized in order to minimize the work done by the mirrors over a single repetition of the pattern. With these considerations in mind, we'll now discuss the three different types of scan patterns diagrammed in figure 5-1.

Raster Scan: This pattern is a grid uniformly spaced in rectangular coordinates x and y . The scan path is left-to-right, top-to-bottom, similar to how a picture is generated on a television screen. To even the load on the scanners, one can instead scan in concentric boxes about the center of the pattern without increasing the total length of the scan path; however, the square geometry is, in general, not well suited to compensate on a circular aperture, and it is found to be difficult to find an actuator beam waist such that the piston power vector \vec{P}_0 is well-behaved.

Polar Scan: This pattern is a uniform grid in polar coordinates r and θ . The scan path is in concentric circles, thus evenly distributing the load between the scanners. Since the pattern is unevenly spaced and most dense in the innermost circle, we must choose an actuator beam size such that each of the influence functions on the innermost circle look sufficiently distinct. This means that a high quality correction can be provided near the innermost circle, but nowhere else (even in the very center, except for the special case where there are 6 angular arms in the pattern: see the next pattern). The degree to which a correction can be provided for distortions away from the center of the aperture falls off very rapidly with increasing radius. Also, unlike the rectangular pattern, there are two degrees of freedom in specifying the polar pattern (number of radial points and number of angular points) which makes finding “good” polar patterns (i.e., patterns with $\vec{P}_0 > 0$) considerably more difficult to locate.

Hexagonal Scan: This pattern, a natural conclusion after considering the previous two cases, is a grid of hexagonal close-packed points. The scan path is in concentric hexagons, thus evenly distributing the load between the scanners. The hexagonal geometry is better suited for circular apertures than the square raster scan, and, unlike the unevenly spaced polar scan, it can provide a high quality correction over the entire optical aperture. Like the raster pattern, there exists only one degree of freedom in the pattern choice: the number of hexagonal rings in the pattern.

5.1.2 Using Zernike Polynomials as Basis Functions

The first attempt at implementing scanning laser thermal compensation using the methodology discussed above was done in MIT Bachelor’s theses by Robert Bennett [4], who performed the numerical calculations, and Phil Marfuta [25], who performed the experimental test. A polar scan pattern was used on a fused silica optic, and Zernike functions were used as the wavefront decomposition functions. Zernike functions are a complete set of orthogonal functions defined on the unit circle, and are often used to describe wavefront distortions on finite circular apertures (the lowest order ones are piston, tilt, focus, astigmatism, coma, etc.). Since the orthogonal Zernike functions span the space of smooth functions on the unit circle, if one can




Figure 5-2: A few Zernike polynomials (astigmatism and coma, as well as an unnamed one of higher order).

generate them with an adaptive optics system (or generate quite a few of the lowest order ones, at least), then one can correct any arbitrary wavefront distortion. In general, these functions are expressed in polar coordinates over two positive indices n and m with $n \geq m$:

$$Z_{nm}(r, \theta) \equiv \sum_{j=0}^{\frac{n-m}{2}} \frac{(-1)^j (n-j)!}{j! \left(\frac{n+m}{2} - j\right)! \left(\frac{n-m}{2} - j\right)!} r^{n-2j} \cdot \begin{cases} \sin(n-2m)\theta & \text{for } n > 2m \\ \cos(n-2m)\theta & \text{for } n \leq 2m \end{cases}$$

and represents a surface with $2(n-2m)$ zero crossings in θ and $n - \|n-2m\|$ zero crossings in r . A few of these functions are graphed in figure 5-2.

Zernike functions are orthogonal in the sense of the L^2 metric [36] on the unit circle (i.e., the inner product on this functional space is simply an area integral over the unit circle):

$$\delta_{n'n} \delta_{m'm} C_{nm} = \int_0^{2\pi} \int_0^1 Z_{nm}(r, \theta) Z_{n'm'} r dr d\theta$$

where C_{nm} is some constant of normalization. So, given a wavefront $\phi(x, y)$, the corresponding representation in the space of Zernikes in terms of amplitudes c_k as

written in equation 5.2 is:

$$c_k = \frac{1}{C_k} \int_0^{2\pi} \int_0^1 \phi(r, \theta) Z_k(r, \theta) r dr d\theta$$

where we've normalized the original radial coordinate r to the size of the aperture, and the Zernike double index (n, m) has been written in terms of a single index $k \equiv n(n+1)/2 + m + 1$. Using Zernike functions with this metric (to find \vec{c} from a given wavefront ϕ), a polar scan pattern, and influence functions computed from a finite element model (to give us \underline{a}), the actuation operator \underline{A} and piston power vector \vec{P}_0 can now be calculated as discussed in the first section of this chapter.

That the actuation operator can, in this case, be calculated was demonstrated by Bennett [4], who also showed that low order Zernike functions can be generated via scanning laser thermal compensation. On further analysis in the experimental demonstration by Marfuta [25], it was noticed that piston was very difficult to generate with a polar scan pattern, since heat exiting the aperture must be compensated by inserting a large amount of power just outside the aperture of concern (to get a flat wavefront, there can be no net radial heat loss out of the aperture of concern), where the density of scan points is the least. This low scan resolution at the periphery of the aperture also hindered the reproduction of higher order Zernike functions, where the most radical behavior occurs. In the end, it was found that Zernikes up to $(n, m) = (3, 3)$ could be experimentally generated [25] with the apparatus further described in Chapter 6.

The conclusions drawn from the work of Bennett and Marfuta are thus:

1. **Zernikes are bad idea for thermal compensation.** While they yield a complete representation of wavefront distortions over a finite aperture, their radical behavior at the edge of the aperture means that they are difficult to reproduce with relatively ponderous thermal actuation. Equivalently, it takes a tremendous number of Zernike functions to accurately represent any of our influence functions, and similarly for thermal distortions we might eventually want to compensate (e.g., a point absorber).

2. **A polar scan pattern is also not a good idea.** While we don't necessarily need high resolution at the edge of the aperture (remember that we care most about the center, where most of the probe beam power traverses), we do need to fulfill the condition that, to generate piston, there cannot be any net radial heat flow from the aperture of concern. The amount of power required to generate piston thus increases rapidly with radius when the aperture of concern is considerably smaller than the optic's total radius, as is the case for LIGO core optics. For the polar scan pattern, the number of scan points lying near a circle of radius r does not increase with r (as it does with the raster scan and hexagonal scan), meaning that considerably larger amounts of power must be deposited at scan points at larger radii when compared to the other scan patterns.

5.1.3 Using Influence Functions as Basis Functions

Instead of working in a functional basis of orthogonal polynomials, it is perhaps more natural to work in the basis of influence functions themselves. Recall that the m th influence function $\mathcal{I}_m(x, y)$ is the net distortion generated per unit power by the laser actuating on the m th scan point (x_m, y_m) . The space of compensatable wavefronts is automatically spanned, but the influence functions are not generally orthogonal in the normal L^2 sense. As long as each function $\mathcal{I}_m(x, y)$ is linearly independent of the previous $m - 1$ functions, though, it is possible to construct an orthogonal basis from the \mathcal{I}_m using the standard L^2 metric (see, for example, [36]).

In the terminology of §5.1 (equation 5.3, in particular), the influence function decomposition matrix \underline{a} is now simply the $M \times M$ identity matrix, as is the corresponding actuation operator \underline{A} . The problem is now how to compute the vector representation $\vec{c} = \vec{P}$ for a given wavefront $\phi(x, y)$:

$$\phi(x, y) = \sum_{m=1}^M P_m \mathcal{I}_m(x, y).$$

Blindly applying the L^2 metric obviously doesn't work since, if ϕ is the k th influence

function \mathcal{I}_k , the resulting decomposition is:

$$P_m = \iint \mathcal{I}_k(x, y) \mathcal{I}_m(x, y) dx dy$$

which is *not* identically zero for all $m \neq k$.

To determine how to properly calculate the decomposition $\vec{c} = \vec{P}$ for a given wavefront ϕ , suppose we are given wavefront data $\{\phi_k\}$ over a set of K discrete points $\{(x_k, y_k)\}$ (*not* the scan basis $\{(x_m, y_m)\}$). We construct a merit function comparing the requested wavefront to the wavefront we are able to generate with the influence functions:

$$\chi^2 = \sum_{k=1}^K W_k (\phi_k + C - \phi_{\text{fit}}(x_k, y_k))^2$$

where C is an arbitrary constant with units of length, W_k is a dimensionless weight dependent on the spatial coordinate (x_k, y_k) , and:

$$\phi_{\text{fit}}(x_k, y_k) \equiv \sum_{m=1}^M P_m \mathcal{I}_m(x_k, y_k)$$

is the wavefront we are able to generate. The weight W_k is included so that some part of the aperture can be considered “more important” than the rest (if we’re trying to fix the wavefront distortion experienced by a Gaussian probe beam, then these weights will follow the intensity profile of the probe beam). Minimizing the merit function yields the best possible reconstruction of the requested wavefront.

Differentiating the merit function with respect to each P_m and equating to zero yields a set of M equations, the n th of which is:

$$\begin{aligned} 0 &= \sum_{m=1}^M P_m \sum_{k=1}^K W_k \mathcal{I}_m(x_k, y_k) \mathcal{I}_n(x_k, y_k) - \sum_{k=1}^K (\mathcal{I}_n(x_k, y_k) + C) W_k \phi_k \\ &= \sum_{m=1}^M P_m \mathcal{I}_{n,m} - d_n \end{aligned}$$

where we have defined:

$$d_k \equiv \sum_{k=1}^K (\mathcal{I}_n(x_k, y_k) + C) W_k \phi_k \quad \mathcal{I}_{n,m} \equiv \sum_{k=1}^K W_k \mathcal{I}_n(x_k, y_k) \mathcal{I}_m(x_k, y_k). \quad (5.5)$$

which is equivalent to the matrix equation:

$$\vec{d} = \vec{P} \cdot \underline{\mathcal{I}}$$

In the limit where the wavefront data are continuous, all sums over k become double integrals over the aperture, the “poke matrix” $\underline{\mathcal{I}}$ simply represents the overlap of the influence functions among each other, and \vec{d} is simply the direct overlap of the requested wavefront $\phi(x, y)$ against the actuation functions in the \mathbf{L}^2 metric. The M dimensional power vector is found by inverting the $M \times M$ poke matrix $\underline{\mathcal{I}}$:

$$\vec{P} = \vec{d} \cdot \underline{\mathcal{I}}^{-1} \quad (5.6)$$

Using wavefront slope data

If the phase measurement is performed with a Shack-Hartmann sensor (see Chapter 6) then wavefront slope data $\left(\frac{\partial\phi_k}{\partial x}, \frac{\partial\phi_k}{\partial y}\right)$ are obtained over a set of K discrete points (x_k, y_k) . The merit function here is then slightly modified:

$$\chi^2 = \sum_{k=1}^K W_k \left[\left(\frac{\partial\phi_k}{\partial x} + C_x - \frac{\partial}{\partial x} \phi_{\text{fit}}(x_k, y_k) \right)^2 + \left(\frac{\partial\phi_k}{\partial y} + C_y - \frac{\partial}{\partial y} \phi_{\text{fit}}(x_k, y_k) \right)^2 \right]$$

where C_x and C_y are arbitrary dimensionless constants, the spatially dependent weight W_k is defined as before, and:

$$\frac{\partial}{\partial x} \phi_{\text{fit}}(x_k, y_k) \equiv \sum_{m=1}^M P_m \frac{\partial}{\partial x} \mathcal{I}_m(x_k, y_k)$$

with a similar result for the y -dimension terms. Organizing to form the matrix equation 5.6, we find the matrix elements to be:

$$d_n \equiv \sum_{k=1}^K W_k \left[\left(\frac{\partial}{\partial x} \mathcal{I}_n(x_k, y_k) + C_x \right) \frac{\partial \phi_k}{\partial x} + \left(\frac{\partial}{\partial y} \mathcal{I}_n(x_k, y_k) + C_y \right) \frac{\partial \phi_k}{\partial y} \right]$$

$$\mathcal{I}_{n,m} \equiv \sum_{k=1}^K W_k \left[\frac{\partial}{\partial x} \mathcal{I}_n(x_k, y_k) \frac{\partial}{\partial x} \mathcal{I}_m(x_k, y_k) + \frac{\partial}{\partial y} \mathcal{I}_n(x_k, y_k) \frac{\partial}{\partial y} \mathcal{I}_m(x_k, y_k) \right].$$

and \vec{P} is again found via equation 5.6.

5.1.4 Example: Compensating a Point-absorber

This section discusses the practical implementation of scanning laser thermal compensation, which until now has only been discussed in very general terms. The case considered is a scaled down version of a LIGO test mass used in the experiment discussed in Chapter 6: a cylindrical fused silica optic 10 cm in diameter over which is phase is measured in transmission over a 3 cm diameter coaxial aperture. In the experiment, thermal distortions are induced in the optic by a tailored external heating beam, the resulting phase distortion is measured, the necessary scanning laser powers are calculated from these phase data, and the scanning laser system then actuates to remove the initial distortion. There isn't a "probe beam", in the LIGO sense per se, but we'll assume for weighting and effective scattering calculations that the radius of the aperture is the waist of a pretend LIGO probe beam.

Step 1: Choose a scan pattern

The first step in implementing scanning laser thermal compensation is to choose a scan pattern. By the previous discussion, the hexagonal scan appears to be the best, so it's now a question of how many hexagonal rings to have in the pattern. There must be one (and only one) ring outside the measured aperture. The spacing of the points in the pattern, hence the number of rings within the aperture, is nominally determined by the spatial scale of absorbers we wish to compensate. Here the goal




Figure 5-3: The scan pattern used in this example. The measured aperture is 1.5 cm, denoted by the dashed circle, and the actuator beam waist is 0.19 cm. The pattern itself is denoted by “+” markers, and the solid line represents the path over which the beam is scanned from point to point. All of the influence functions in the outermost ring, denoted with the bold line, are merged into one enormous influence function in order to bypass problems with degeneracy. The diameters of the circles represent the piston power vector \vec{P}_0 for this pattern (the amount of power that must be deposited at each point to produce piston).

is to address the worst possible case: an absorber with a spatial scale considerably smaller than any pattern spacing we can reasonably generate.

To compensate localized absorption of power P_a at a spatial scale less than or equal to that of the scan pattern, each point in the scan pattern (excluding the one nearest the point absorber) will be required to deposit an equal or larger power. For M points in the scan pattern, the total power input to the optic by the laser is then at least MP_a . More scan points results in more necessary input power, for the reason that smaller spatial scale components of the point absorber are being addressed. Recall from §2.2 that, for perfect thermal compensation, the total intensity absorbed across the face of the optic is equal to the maximum intensity of the absorber we wish to compensate. In the case of a strong absorber in the Advanced LIGO Power Recycling Cavity, the maximum intensity absorbed is equal to the intensity of the Power Recycling cavity beam, thus the total power input by the perfect compensator would need to be on the order of the power of the Power Recycling Cavity beam itself (about 2 kW). By using a coarser scan pattern, we exercise a tradeoff in the quality of correction for total necessary input power. The total power deliverable by our laser and/or the total power we are comfortable depositing into the optic are the criteria that will drive the pattern spacing selection here.

For a perfect absorber of diameter $d_a = 300 \mu\text{m}$ at powers similar to that seen in the Advanced LIGO power recycling cavity, the total amount of power absorbed is:

$$P_a = \frac{2P_{\text{PRC}}}{\pi w^2} \frac{\pi d_a^2}{4} \approx 25 \text{ mW}$$

The laser used in the experiment is limited to 10 Watts of output, and the duty cycle of the scan pattern itself will be 50% (the laser is switched off when traveling from point to point). This limits the total number of scan points to $10 \text{ W} \times 50\% / 25 \text{ mW} = 200$ points, which corresponds to 7 hexagonal rings (169 points). However, we'll reduce this by a factor of 2 further, since we expect significantly larger amounts of power to be necessary just outside of the aperture in order to negate radial heat flow from within the aperture. In the end, we settle on 4 hexagonal rings (61 points) with spacing such




Figure 5-4: A few of the influence functions calculated in this example.

that the outermost ring lies exclusively outside of the measured aperture, as shown in figure 5-3. Also, anticipating some degree of degeneracy between points outside the aperture, we'll "link" them together into one influence function: i.e., constrain that the same amount of power must always be deposited at each point in this outermost ring.

Step 2: Calculate the influence functions and find a "good" actuator beam waist

With the scan pattern in-hand, we now choose an actuator beam waist and proceed with calculating the influence functions. We expect the necessary waist to be smaller than the pattern spacing, otherwise there will be significant overlap of neighboring actuation beams as well as their corresponding influence functions. So, we pick a waist which is about half the pattern spacing, and then calculate the actuation functions using the three dimensional finite element model discussed in Chapter 3. A few influence functions are shown in figure 5-4. The poke matrix is then numerically calculated (equation 5.5):

$$\mathcal{I}_{n,m} = \frac{\iint e^{-2\frac{x^2+y^2}{R^2}} \mathcal{I}_n(x,y) \mathcal{I}_m(x,y) dx dy}{\iint e^{-2\frac{x^2+y^2}{R^2}} dx dy}$$

and is inverted to get the modified actuation matrix (equation 5.6). Piston is then decomposed in this basis:

$$d_n = \frac{\iint e^{-2\frac{x^2+y^2}{R^2}} \mathcal{I}_n(x, y) dx dy}{\iint e^{-2\frac{x^2+y^2}{R^2}} dx dy}$$

and the behavior of the resulting piston power vector is examined by applying the modified actuation matrix (equation 5.6). If the behavior is erratic, we pick a smaller actuator beam waist at start over. With the scan pattern shown in figure 5-3, we find an actuator beam waist of 0.19 cm (about $\sim 2.5\times$ smaller than the pattern spacing) yields a well-behaved piston power vector which is plotted as circles in the same figure. We now have the basis of influence functions $\{\mathcal{I}_m(x, y)\}$ and the corresponding poke matrix $\underline{\mathcal{I}}$ and are now ready to compensate measured wavefront distortions.

Step 3: Compensate wavefront distortions

Given any measured measured wavefront distortion $\phi(x, y)$, we wish to compensate it by thermally actuating to induce the conjugate $-\phi(x, y)$. In this example, we obtain $\phi(x, y)$ by modeling 25 mW of absorption over a $300\ \mu\text{m}$ spot on the face of the test optic. The conjugate wavefront $-\phi$ is then decomposed in terms of the influence functions (\vec{d} in equation 5.5) and equation 5.6 is applied to get the raw power vector. The piston power vector is then added in sufficient quantity to make the resulting total power vector purely positive (and identically zero at one point). The original distortion, the calculated powers, and the resulting corrected wavefront are all detailed in figure 5-5.

5.1.5 Addressing the time dependence of scanning laser actuation

The previous analysis was done assuming a steady-state heating pattern, which is certainly untrue for a scanning laser which is being switched on and off while moving point to point across the pattern. Here we'll consider the effect of a scan pattern

Scanning Laser Compensation of 20mW absorbed on 300 μ m spot

0.19cm compensator beam waist, 5nm contour interval

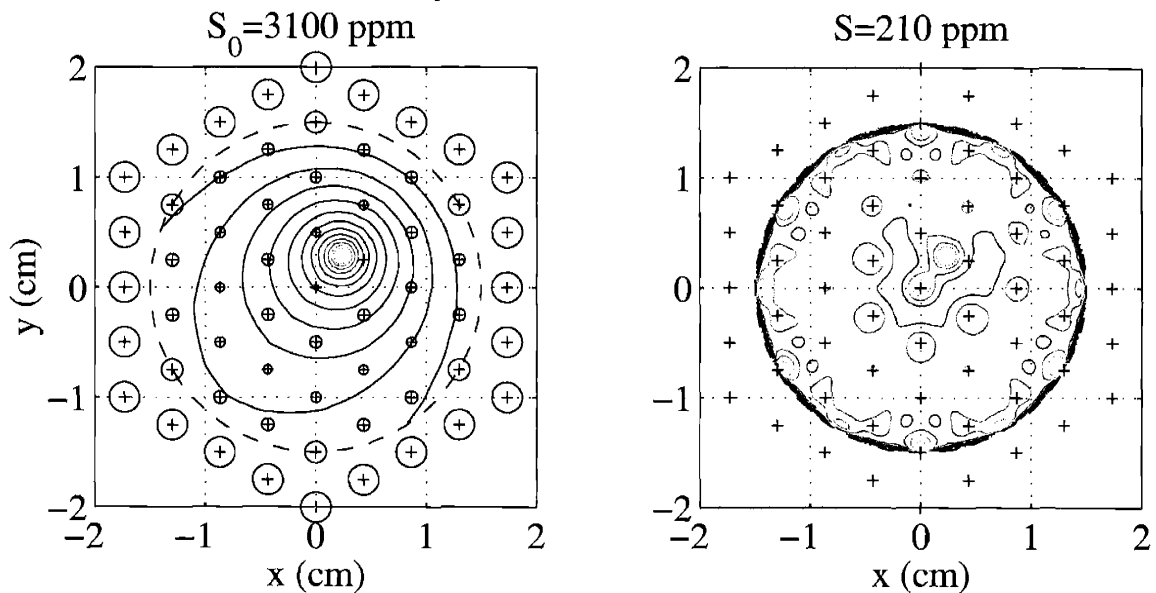


Figure 5-5: The result of thermally compensating a small region of large absorption. As in figure 5-3, the measured aperture is denoted by the dashed circle, and the points of the scan pattern are denoted by “+” markers. The diameters of the circles in the left diagram represent the powers calculated to correct the wavefront distortion shown. The right diagram shows the resulting compensated wavefront. The TEM_{00} scatter for a Gaussian beam whose waist is the size of the measured aperture is listed at the top of each diagram.

identical to that analyzed in the previous section, scaled up to LIGO sizes (same number of scan points and the same power laser, but the aperture is now 5 cm in radius and the actuator beam waist will be approximately 1 cm).

Consider the m th point of the M -pointed actuation pattern where an average power P_m is deposited, and suppose that the actuator beam returns to this point after a time τ_p (the pattern period). The actuator beam shines on this point for a time τ_p/M with peak power MP_m , so the net energy input per scan cycle is simply $\tau_p P_m$. In order for the steady-state optical path distortion to dominate over the time dependent component, the total scan period τ_p must be small compared to the local thermal time constant (the time it takes for radial heat flow to be established over the actuator beam area):

$$\tau_a \sim \pi w_a^2 \frac{\rho c}{k} \approx 120 \text{ s} \left(\frac{w_a}{1 \text{ cm}} \right)^2 \quad (5.7)$$

where ρ is the density, c is the heat capacity, k the thermal conductivity of the optic (fused silica), and w_a the actuator beam waist radius. As such, the approximation of negligible radial heat flow applies over the time-scale τ_p/M (recall the discussion of section §2.3), and the net *local* optical path change induced by introducing energy $\tau_p P_m$ over the actuation beam waist w_a in the short amount of time $\tau_p \ll M\tau_a$ is then bounded above by:

$$\begin{aligned} \max \|\delta x\|_l &\lesssim \frac{P_m \tau_p}{w_a^2} \frac{\beta}{\rho c} \\ &\approx 60 \text{ nm} \left(\frac{P_m}{0.1 \text{ W}} \right) \left(\frac{\tau_p}{10 \text{ s}} \right) \left(\frac{1 \text{ cm}}{w_a} \right)^2 \left(\frac{\beta}{10 \text{ ppm}/^\circ\text{K}} \right) \end{aligned} \quad (5.8)$$

which is *extremely* large in terms of the length scales that LIGO hopes to measure.

At any given time, though, this large and relatively rapid increase in optical path is occurring on only a small portion of the aperture (around the actuator beam waist, which is much smaller than the LIGO beam waist), so that the measured length

fluctuation on time scales of τ_P/M is actually quite smaller:

$$\begin{aligned} \|\delta x\|_{\tau_P/M} &\lesssim \frac{w_a^2}{w^2} \|\delta x\|_l \lesssim \frac{P_m \tau_P}{w^2} \frac{\beta}{\rho c} \\ &\approx 2 \text{ nm} \left(\frac{P_m}{0.1 \text{ W}} \right) \left(\frac{\tau_P}{10 \text{ s}} \right) \left(\frac{5 \text{ cm}}{w} \right)^2 \left(\frac{\beta}{10 \text{ ppm}/^\circ\text{K}} \right). \end{aligned} \quad (5.9)$$

This is still large compared to what LIGO hopes to measure, but occurs as a result of a well-known input at a within a well-defined frequency band $f \approx M/\tau_P = 6 \text{ Hz}$ (60 points on a 10 second pattern period).

After the actuator leaves the pattern point, the local optical path will decay by the same amount that it abruptly increased (equation 5.8), but now over the longer time scale τ_P . However, we are not saved by the small actuator beam size here, since every scan point in the aperture, excluding the one on which the actuator is shining, is experiencing this slow decay. Thus, on time scales of τ_P the total fluctuating optical path is equal to the large quantity in equation 5.8:

$$\begin{aligned} \|\delta x\|_{\tau_P} &\lesssim \|\delta x\|_l \lesssim \frac{P_m \tau_P}{w_a^2} \frac{\beta}{\rho c} \\ &\approx 60 \text{ nm} \left(\frac{P_m}{0.1 \text{ W}} \right) \left(\frac{\tau_P}{10 \text{ s}} \right) \left(\frac{1 \text{ cm}}{w_a} \right)^2 \left(\frac{\beta}{10 \text{ ppm}/^\circ\text{K}} \right). \end{aligned} \quad (5.10)$$

Again, this is large compared to what LIGO hopes to measure, but occurs as a result of a well-known input at a within a well-defined (and now much lower) frequency band $f \approx 1/\tau_P = 0.1 \text{ Hz}$.

Both of these numbers, 60 nm fluctuations at 0.1 Hz and 2 nm fluctuations at 6 Hz, are large but are the result of a well-known input and lie in narrow frequency bands below the low-frequency edge of LIGO's detection band. Increasing the number of points, hence reducing the actuator beam size and forcing a shorter pattern period, necessarily moves the frequencies of these fluctuations into the LIGO detection band.

5.2 Noise Considerations

The relatively large optical path fluctuation induced by shining the actuator on-and-off occurs within well-defined frequency bands at the pattern repetition frequency (~ 0.1 Hz) and the laser amplitude modulation frequency (~ 6 Hz), hence any effect seen at the output of the interferometer will lie within this pair of very narrow frequency bands. However, broadband intensity noise on the actuator beam will induce broadband optical path fluctuations through the compensated optics which will be seen at the output of the instrument. We'll now calculate this effect using the formulae derived in §2.3.

At a given instant, the actuator beam will be depositing peak power $P = MP_m$ at the m th pattern point $\vec{r}_m = (x_m, y_m)$, and we may write the intensity noise of this beam as:

$$I(x, y, \omega) = \frac{2P}{\pi w_a^2} \frac{\Delta P(\omega)}{P} e^{-2 \frac{(x-x_m)^2 + (y-y_m)^2}{w_a^2}}$$

where w_a is the waist of the beam and ΔP is the spectral density of the beam power at frequency ω . In the terminology of §2.3, the broadband ripple is simply:

$$\mathcal{R}(\omega) = \frac{\Delta P(\omega)}{P}$$

Without loss of generality, we may set $y_m = 0$, $x_m = \|r_m\|$ and now use equation 2.24 to calculate the necessary upper limit required on $\Delta P(\omega)$ as a function of the actuation radius r_m . Assuming the actuator beam is much smaller than the LIGO beam, we can remove the LIGO beam term from within the integral in equation 2.24:

$$\begin{aligned} \delta x(\omega t) &\approx \frac{\beta \mathcal{R}(\omega) i}{\rho c \omega} e^{i\omega t} \frac{2}{\pi w^2} e^{-2 \frac{r_m^2}{w^2}} \iint \frac{2P}{\pi w_a^2} e^{-2 \frac{(x-x_m)^2 + y^2}{w_a^2}} dx dy \\ &\approx \frac{\beta i}{\rho c \omega} \frac{2 \Delta P(\omega)}{\pi w^2} e^{-2 \frac{r_m^2}{w^2}} e^{i\omega t} \end{aligned} \quad (5.11)$$

Assuming that actuation is performed on a fused silica compensation plate within the power recycling cavity, the optical path fluctuation caused by scanning laser intensity

noise must be less than:

$$\|\delta x\| < 1 \times 10^{-24} / \sqrt{\text{Hz}} \times 4 \text{ km} \times 800 \approx 3 \times 10^{-18} \text{ m} / \sqrt{\text{Hz}}$$

at 100 Hz, in order to keep from eclipsing the instrument's strain sensitivity goal of $1 \times 10^{-24} / \sqrt{\text{Hz}}$ at 100 Hz. The factor of 800 in this expression is simply the arm cavity gain. Thus, the total intensity fluctuation of the actuator beam at the scan radius r_m must be less than:

$$\Delta P(\omega) < 2 \times 10^{-7} \frac{W}{\sqrt{\text{Hz}}} \left(\frac{\omega}{100 \text{ Hz}} \right) \left(\frac{w}{5 \text{ cm}} \right)^2 \left(\frac{10 \text{ ppm}}{\beta} \right) e^{2 \frac{r_m^2}{w^2}}. \quad (5.12)$$


Note that this limit relaxes exponentially with radius and is about 3 orders of magnitude above the shot noise limit for a 10 Watt CO2 laser beam ($\lambda = 10.6 \mu\text{m}$):

$$\Delta P_{\text{shot}}(\omega) \approx 6 \times 10^{-10} \frac{W}{\sqrt{\text{Hz}}} \left(\frac{P}{10 \text{ W}} \right)^{\frac{1}{2}}.$$

Thus, in principle, it is possible to adequately intensity stabilize the CO2 laser such that 20 Watts or more of peak compensator power can be delivered to the optic without compromising the interferometer's sensitivity goal.

Chapter 6

The Proof-of-Principle Experiment



The previous chapters have identified and analyzed the problem of thermal distortions in high power laser interferometric gravitational wave antennae, and thoroughly analyzed the potential solution of thermally adaptive optics. The work of the previous chapters is exclusively analytic and/or numeric, however, and the cliché ‘talk is cheap’ may have crossed the experimentalist reader’s mind. Perhaps we overlooked something subtle but crucial in the analysis, or perhaps the solution is so sensitive that it is not practically realizable. Such is the motivation for building a testing a prototype of both heating ring and scanning laser thermal compensation. Using LIGO itself is not an option since, at the time of writing this thesis, it does not operate anywhere near the powers necessary to induce thermal distortions which affect the performance of the instrument (perhaps more significantly, it is trying to detect gravity waves, and using it as a test-bed will interfere with this primary mission). The prototype is thus a scaled-down test designed to demonstrate thermal wavefront distortions and their correction, via shielded heating ring or scanning laser compensation, on an optic rigidly suspended in high vacuum.

6.1 General Setup

The ultimate goal of the experiment is to examine absorption-induced wavefront distortions in a cylindrical optic, as well as the correction produced by a prototype compensator (shielded ring or scanning laser). There are thus three key components to the experiment: (1) the probe, by which we measure wavefront distortions, (2) the pump, by which we generate LIGO-like wavefront distortions, and (3) the compensator(s), by which we thermally correct the LIGO-like wavefront distortions. Figure 6-1 diagrams the experimental setup, which is now discussed in detail. Excluding the materials measurement, the test optic is a flat-flat cylindrical fused silica test mass, 10 cm in diameter and 8 cm thick with a 1° wedge, inherited from the 40 meter LIGO prototype at Caltech which is mounted on an optical breadboard in a cylindrical vacuum enclosure 2 meters in diameter and 3 meters tall. The scaled-down aperture that we will probe is 3 cm in diameter (recall that Initial LIGO has a 7.5 cm diameter

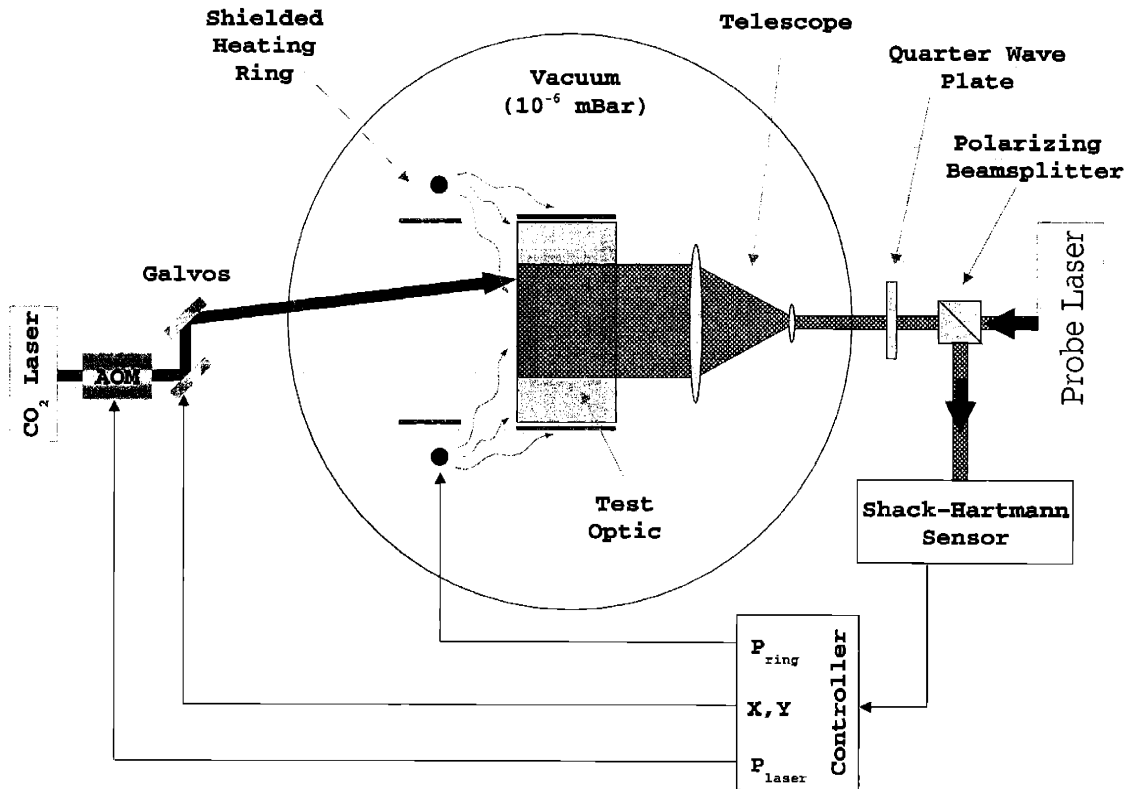


Figure 6-1: General setup of the experiment.

probe beam over 25 cm diameter optics).

The Probe

To measure wavefront distortions induced in a test optic mounted in high vacuum, we illuminate the optic with a well collimated probe beam of very low optical power and examine the reflection of this beam on a commercially available Shack-Hartmann sensor. The right side of figure 6-1 diagrams the probe in this experiment.

A red ($\lambda = 633$ nm) diode laser is first coupled through a single mode fiber to remove any time dependent wavefront fluctuations from the input probe light (these fluctuations are converted into intensity fluctuations at the output of the fiber, to which our Shack-Hartmann sensor is insensitive). It passes through a polarizing beamsplitter which selects a linear polarization, and then through a $\lambda/4$ plate where

the polarization becomes circular. The probe passes through a simple telescope to collimate the beam as well as expand it to appreciably cover the desired 3 cm aperture (the beam waist lies beyond the effective aperture which we will image on the wavefront sensor) and is retroreflected off of either test optic surface we chose: “front” meaning the surface closest to the probe, and “back” meaning the far surface. When the retroreflected probe beam returns to the $\lambda/4$ plate, it is of the opposite circular polarization as it was at the input (physically speaking, angular momentum of the beam is conserved) so, after passage, the returning polarization is orthogonal to that of the input. The polarizing beamsplitter fully reflects this orthogonal polarization, which is directed through collimating optics and onto the Shack-Hartmann wavefront sensor.

A Shack-Hartmann wavefront sensor is a nominally simple device consisting of an array of tiny lenses (typically called “lenslets”) etched on a piece of glass with a CCD camera mounted at the focal plane of the array. For a beam illuminating the lenslet array, each lenslet produces a focal spot on the CCD behind it at a location determined by the mean local slope of the portion of the beam illuminating the lenslet. If the local slope of the input light illuminating a lenslet changes, then the focal spot on the CCD displaces a distance proportional to the change in wavefront slope. By tracking the focal spots on the CCD camera after the pump is turned on, we are able to reconstruct the pump-induced wavefront distortion from the slope data. The wavefront sensor utilized is commercially available from Wavefront Sciences, Inc. [44] and consists of a 31×24 lenslet array with each lenslet $198 \mu\text{m}$ in diameter and a focal length of 15.5 mm hard-mounted on a CCD camera, a COHU type 6612, which has a $4.8 \times 6.3 \text{ mm}$ active area consisting of 640×480 pixels.

The key to the phase measurement is to relay image the measured optical surface onto the lenslet array. In other words, if M is the magnification of the optical system between the measured optical surface and the wavefront sensor surface, then the wavefront slope $(\frac{\partial\phi}{\partial x}, \frac{\partial\phi}{\partial y})$ at coordinate (x, y) on the optic will result in a measurement of slope $(M\frac{\partial\phi}{\partial x}, M\frac{\partial\phi}{\partial y})$ at sensor coordinates $(\frac{1}{M}x, \frac{1}{M}y)$. If this condition is not met, then wavefront slopes induced in the test optic will cause the image on the wavefront

sensor to change: i.e., the mapping of coordinates from the measured plane to the sensor plane depends on local wavefront slope. To properly set up the optical train so that the image is properly relayed, we first set up and align the optical system as best we can, making the beam as collimated as we can easily get and of the proper size at the measured plane as well as on the wavefront sensor. The beam at the wavefront sensor plane is then examined on a test CCD via a pickoff placed slightly before the wavefront sensor, the probe laser is switched off, and a back-illuminated USAF 1951 resolution test pattern is placed at the measured plane (back-illumination is provided by a simple flashlight). The relay imaging condition is satisfied if there is a clear image of the test pattern produced on the test CCD. The large optic of the in-vacuum telescope (the lens closest to the test optic) is axially displaced until the image of the test pattern on the test CCD becomes the sharpest possible. If we were careful about the initial alignment, then the necessary axial displacement of the large lens is very small, but is nonetheless *extremely* critical to fulfilling the imaging condition (as was found in the process of performing this experiment).

Finally, we must calibrate the system so that we know what to make of a phase measurement $\phi(x_s, y_s)$ at sensor coordinates (x_s, y_s) . The absolute measure of wavefront slope depends on the focal length of the lenslet array at the probe wavelength as well as the pixel size of the CCD camera. Since all of these quantities are known extremely well in absolute terms (one part in 1000, according to the manufacturer [44]), we'll assume that the absolute wavefront slope measurement is also better than one part in 1000. Thus trusting the wavefront sensor, the magnification of the optical train between the sensor plane and detected plane will be the greatest potential source of error. Although the magnification has no effect on the phasemap reconstructed from the wavefront slopes (since the wavefront slopes are amplified by the magnification of the optical train while the spatial coordinates are attenuated by the inverse magnification), it determines the mapping of coordinates from the measured plane to the sensor plane. In other words, if phase $\phi(x_s, y_s)$ is measured at sensor coordinates (x_s, y_s) , then this corresponds to the identical phase $\phi(x_d, y_d)$ at coordinates $(x_d, y_d) = (Mx_s, My_s)$ in the measured plane.

To precisely determine the magnification of the system, we execute two measurements: (1) block the aperture at the measured plane with a horizontal rod attached to a vertical micrometer, and correlate the image coordinates on the test CCD with the readings on the micrometer as the rod is swept over the entire aperture, and (2) tilt the test optic very slightly (a few micro-radians) and examine how the slopes measured on each lenslet of the wavefront sensor vary. The first measurement gives us the mean magnification of the system, and the second yields a map of how the magnification varies across the measured aperture.

The Pump

To generate LIGO-like wavefront distortions in the test optic we use a carbon dioxide laser, whose emission is at a wavelength which is strongly absorbed in the immediate surfaces of both sapphire and fused silica ($\lambda = 10.6 \mu\text{m}$). The setup of the CO₂ laser is shown on the left in figure 6-1. The full output of the laser, a Synrad J48-1SW, is nominally 10 Watts in a 3.5 mm diameter Gaussian beam at the output (quoted by the manufacturer at 95% TEM₀₀) and is directed through an acousto-optic modulator (AOM), an Intra-Action AGM-406B1, which allows us to control beam power without disturbing the drive power of the laser. This means of controlling the laser power was found to be extremely important, as directly changing the drive power of the CO₂ laser itself was observed to slightly change the pointing of the output mode and, more significantly, change the output wavelength of the laser (there are multiple CO₂ lines between 10.57 and 10.63 microns which can be active in this particular model of laser, and are selected by the laser cavity length which thermally expands or contracts with the drive power provided). The beam is then sent through a pair of galvanometer scanners, GSI Lumonics model VM-2000, to control the x and y position of the beam across the face of the test optic and enters the tank through a 3 inch Zinc Selenide viewport.

To shape the beam to what is desired in a particular test, a trio of Zinc Selenide optics (two aspheric lenses with $f=1.5''$ and one lens with $f=10''$) are used. Alignment of the beam through these optics proves to be quite a challenge due to the very long

wavelength of the laser, but this task is greatly simplified with the aid of small paddles coated in a phosphor whose UV fluorescence depends heavily on temperature (commercially available from Macken Instruments). The CO₂ beam shining on one of these UV illuminated paddles shows up as a black spot on a bright neon yellow background, and allows for precise alignment through the optics and mirrors in the setup. The paddles only give an inkling as to what the precise beam size is, however, so we resort to other methods to obtain these data.

To precisely measure the size of the beam at the optic face, the optic is removed and a pair of measurement techniques is used, depending on the intended beam size: (1) for small beams (diameter less than 7 mm) a Photon Inc. model XYFIR pyroelectric beam profiler is placed in the plane of the optic face, and (2) for larger beams, a 1 mm diameter thin-film thermopile is placed in the plane of the optic face and the galvanometers raster scan the the beam across the small detector, yielding a slow “television picture” of the large pump beam.

To monitor the power entering the tank, a ZnSe pickoff is introduced between the AOM and galvanometers which directs a small fraction of the main beam onto a high power thermopile (Newport model 818T-02) whose output voltage is continuously monitored. To determine the absolute amount of pump power entering the vacuum tank versus the monitor voltage, a NIST-traceable thermopile power meter (Ophir model Orion calibrated thermopile laser power meter, with a quoted absolute accuracy $\pm 1\%$ for CO₂ beams [29]) is placed in the tank and the absolute reading is correlated to the voltage output of the monitor.

Finally, to determine the absolute amount of power that is absorbed in the test optic versus the monitor voltage, we measure the $\lambda = 10.6 \mu\text{m}$ reflectivity of the pumped optical surface. The pump beam in the experiment will, nominally, not be normal to the optical surface due to space constraints of the input optical table and vacuum tank. As such, the reflectivity is easily measured by locating the reflected beam in the tank and measuring its absolute power versus on-table monitor voltage. Comparing this to the calibration of the total absolute power in the tank versus monitor voltage, we recover the reflectivity of the surface under test and, since no

power is transmitted through the substrate, the absolute power absorbed is known versus monitor voltage. To illustrate the fact that there is really *no* appreciable transmission into, much less through, the substrate material, we note the results of a simple transmission measurement which was attempted by directing the CO₂ beam at a 300 μm thick fused silica cover slide: absolutely no transmission could be measured (the measurement resolution was 1 mW) until the laser was elevated to such a power where it quickly bored a hole through the slide (a few Watts).

The Compensators

More thorough descriptions of the two compensators are described in their own dedicated sections of this chapter (§6.3 for the shielded ring and §6.4 for the scanning laser), although a cursory overview is provided here.

In the case of scanning laser thermal compensation, the compensator is simply the CO₂ laser setup described above with the galvanometers scanning the CO₂ beam in a fixed pattern over the face of the optic while the power of the laser is adjusted as necessary from point to point via the acousto-optic modulator. Wavefront data are processed from the Shack-Hartmann sensor and decomposed in (i.e., numerically integrated against) the modeled actuation basis (recall the discussion in §5.1.3). The modeled orthogonalization matrix is numerically applied to this decomposition to yield the powers required at each scan point, which is then converted into voltages fed to the AOM driver.

The compensator in the shielded ring test is a ring of resistive wire rigidly mounted to (but thermally insulated from) an aluminum frame. Also attached to the frame is a hollow aluminum cylinder (the shield) mounted coaxially to the ring.

6.2 Measuring Material Parameters

The first measurement performed is a validation that we do, in fact, spatially and temporally understand Gaussian beam heating in sapphire and fused silica optics. Figure 6-2 shows a diagram of the setup for this measurement. The CO₂ beam is




Figure 6-2: Setup of the materials measurement.

shaped and directed to a relatively small spot (about 5 mm in diameter) at the center of the front face of a slightly wedged test optic, which is optionally probed in either transmission or reflection. A pickoff is introduced in the beam to a thermopile which allows us to precisely monitor and control the absolute power absorbed in the optic under test as described by the pump thruput calibration procedure in the previous section. An electronically actuated fast mechanical shutter is placed in the beam path before the tank which allows us to abruptly switch the pump beam on and off. By abruptly switching on the intensity-stabilized pump beam, we are able to examine the evolution of the induced optical path distortion in either transmission or reflection versus the time elapsed since the pump was switched on. From this measurement we'll recover the necessary materials parameters to accurately explain what is observed in terms of the two dimensional numerical model thoroughly discussed in Chapter 3. The quantities which we will assume as unknown are: (1) the thermal expansion coefficient α , (2) the thermooptic coefficient $\frac{dn}{dT}$, and (3) the thermal conductivity k .

We'll first examine the time evolution of the surface distortion induced by abruptly illuminating the test optic with the small circular CO2 beam of constant intensity, in order to recover (in the case of a generic uniaxial crystal with the axis of symmetry aligned with the symmetry axis of the cylindrical optic) the radial and axial thermal expansion coefficients α_r and α_z as well as the radial and axial thermal conductivities k_r and k_z . The actual wavefront measurement is performed via a Shack-Hartmann sensor as described in the previous section, thus yielding local wavefront slope data $(\frac{\partial\phi}{\partial x}(x_n, y_m, t), \frac{\partial\phi}{\partial y}(x_n, y_m, t))$ over the rectangular grid of Shack-Hartmann lenslets (x_n, y_m) ; however, the test geometry is cylindrically symmetric (the circular pump beam illuminates the center of the cylindrical test optic) so the measured quantity of interest is simply:

$$\frac{\partial\phi}{\partial r}(r, t) = \frac{1}{r} \left(x \frac{\partial\phi}{\partial x}(x, y, t) + y \frac{\partial\phi}{\partial y}(x, y, t) \right).$$

Now, the behavior of $\frac{\partial\phi}{\partial r}$ in the part of the aperture where the pump beam deposits power (i.e., for radii r smaller than the radius of the pump beam r_p) is heavily depen-

dent on profile and size of the small pump beam, as well as exhibiting the heaviest variation of wavefront slope over each Shack-Hartmann lenslet. We'll therefore ignore the slope data taken here, and instead concentrate on the smoothly varying data taken on the periphery of the aperture $r > r_p$, which, by the analogy drawn to Gauss' law discussed at the end of §3.1.2, should be almost completely independent of the intensity profile of the interior pump beam. Furthermore, the fact that the thermal wavefront distortion away from the pump beam is smooth and monotonic means that it is well approximated by a quadratic function of r , thus the slope will look approximately linear over r here. So, over a “donut” aperture with inner diameter r_1 greater than the pump beam radius r_p and outer diameter r_2 smaller than the total aperture r_a (i.e., $r_p < r_1 \leq r \leq r_2 < r_a$) we fit a line to the measured wavefront slope data:

$$\frac{d\phi}{dr}(r, t) \approx a(t) + b(t)r.$$

We then fit the equivalent results of the numerical model to the data $a(t)$ and $b(t)$ to recover the desired material parameters. Physically, the term $a(t)$ represents the raw magnitude of the surface distortion, while $b(t)$ represents the local curvature. While the axial thermal expansion coefficient governs the total magnitude of both a and b , an excessive radial thermal expansion coefficient will cause the entire test optic to gently “bow” more than expected, thus resulting in a measured value of b which is smaller than anticipated. The timing of these two events (raw axial expansion and the gentle “bowing” induced by radial expansion) are determined by the radial and axial thermal conductivities (respectively).

To determine how to efficiently fit the model to these data a and b , recall the discussion of the functional dependencies of the surface distortion on thermal expansion and thermal conductivity in §3.1.3. In particular, the net axial surface distortion can be expressed in terms of the thermal conductivities and stress temperature moduli (which are simply related to the thermal expansion coefficients by the known elastic

modulus matrix elements via $\vec{\gamma} = \vec{\alpha}\underline{c}_2$):

$$u_z(r, t) \approx \frac{\gamma_r}{k_z} g_r(r, k_z t) + \frac{\gamma_z}{k_r} g_z(r, k_r t)$$

where g_z is the component of axial expansion due to unit axial stress temperature modulus (i.e., $\gamma_r = 0$ and $\gamma_z = 1$), g_r is the component of axial expansion due to unit radial stress temperature modulus (i.e., $\gamma_r = 1$ and $\gamma_z = 0$), and both functions g_r and g_z are dependent only on the geometry and the elastic modulus matrix elements. The functions $g_r(r, t)$ and $g_z(r, t)$ are obtained from the numerical model, the radial derivatives $\frac{\partial}{\partial r}$ of each are calculated, and the results are then interpolated onto the wavefront sensor coordinates. The four model “fitting functions” $\mathbf{a}_r(t)$, $\mathbf{a}_z(t)$, $\mathbf{b}_r(t)$, and $\mathbf{b}_z(t)$ are then calculated:

$$\begin{aligned} \frac{\partial g_r}{\partial r}(r, t) &= \mathbf{a}_r(t) + \mathbf{b}_r(t)r \\ \frac{\partial g_z}{\partial r}(r, t) &= \mathbf{a}_z(t) + \mathbf{b}_z(t)r \end{aligned}$$

which we will fit to the data to recover the thermal expansion coefficients and thermal conductivities:

$$\begin{aligned} \mathbf{a}(t) &= \frac{\gamma_r}{k_z} \mathbf{a}_r(r, tk_z) + \frac{\gamma_z}{k_r} \mathbf{a}_z(r, tk_r) \\ \mathbf{b}(t) &= \frac{\gamma_r}{k_z} \mathbf{b}_r(r, tk_z) + \frac{\gamma_z}{k_r} \mathbf{b}_z(r, tk_r). \end{aligned}$$

As was the case for g_r and g_z , the fitting functions \mathbf{a} and \mathbf{b} are dependent only on the geometry and the elastic modulus matrix elements. The fitting functions for the C-axis sapphire test optic (discussed further below) are shown in figure 6-3.

Once the thermal conductivities and thermal expansion coefficients are obtained, the optic is tilted so that it is probed in transmission and the steady-state wavefront distortion on the periphery of the aperture is examined to obtain a measurement of the thermooptic coefficient $\frac{dn}{dT}$.

C-axis Sapphire Fitting Functions

2.6 mm pump beam, probed interval $r=(6\text{ mm},13\text{ mm})$




Figure 6-3: Numerically determined fitting functions derived from g_z (solid line), and g_r (dashed line) for the C-axis sapphire sample.

6.2.1 Systematic Parameters

The parameters necessary to completely construct the fitting functions \mathbf{a} , and \mathbf{b} are the elastic moduli matrix elements (see tables A.1 and A.2), the boundaries of the partial differential equations (the test optic's dimensions), and the conditions on those boundaries (pump beam size, r_p , and the material's emissivity, ϵ). The pump beam size is much smaller than the physical dimensions of the test optic, so the fitting functions approach those for an infinite half space and depend only very weakly on the physical dimensions of the optic and the emissivity of the radial surface. Since the measured aperture excludes the central region around the pump beam, the thermal distortion that we sense only weakly depends on the pump beam size or shape. The optic's dimensions are precisely obtained with a set of calipers, and the outer edge is thermally isolated from its aluminum mount by restricting contact to three small Teflon spacers on the radial edge of the optic.

The emissivity of the face under test at room temperature roughly corresponds

to the absorptivity of wavelengths near the peak of the room temperature blackbody distribution (about $10\ \mu\text{m}$). When calibrating the power thru-put of the pump, we measure the reflectivity of the optic at the pump wavelength ($10.6\ \mu\text{m}$) which closely corresponds to the peak wavelength of room temperature blackbody radiation. Since no power is transmitted (recall the cover slide transmission test mentioned in the previous section), one minus the measured reflectivity is the absorptivity at $10.6\ \mu\text{m}$, hence the emissivity, which is approximately equal to the broadband room temperature emissivity.

The most significant potential sources of systematic error arise from the pump thru-put calibration, which determines how well we know the absolute amount of pump power absorbed in the test sample, and the probe magnification, which determines the conversion between the test optic and sensor wavefront slope as well as the conversion between test optic and sensor coordinates. The measurement of these parameters was discussed in the previous section. The error in the pump thru-put calibration will couple directly into the measurement of thermal expansion coefficients and thermo-optic coefficient, and is dominated by the 1% error in absolute calibration of the power meter that is used to measure absolute thru-put. The error in magnification will have little effect on the measured parameter $a(t)$, as this parameter indicates the gross magnitude of the wavefront distortion over the aperture (recall also that it has units of $1/\text{absorbed power}$). The magnification error will have a more significant effect on the measured parameter b , as the magnitude of local wavefront curvature depends on the radial coordinates which it is defined (recall that the units of b are $1/\text{absorbed power}/\text{length}$). Thus, the magnification error will couple most significantly into the measurement of radial thermal expansion.

Table 6.1 lists the relevant systematic parameters that have just been discussed, as well as their corresponding errors.

6.2.2 Fused Silica Measurement

The sample in this measurement is a fused silica (Corning 7940) cylinder, 3" in diameter and 1" thick, salvaged from the Phase Noise Interferometer [21]. The surfaces

Systematic Parameters and Errors			
Parameter	Value	Error	Units
Probe Magnification	7.25	0.10	-
Pump Beam Radius	2.6	0.2	mm
Pump Beam Thru-put	12.77	0.13	W/V
Fused Silica Emissivity	0.89	0.01	-
Sapphire Emissivity	0.93	0.01	-

Table 6.1: Systematic parameters and their errors.

of the sample are polished optically flat with a wedge of $\sim 1^\circ$, and neither surface is coated. The first test consists of abruptly shuttering the pump on (about 1.5 Watts of pump power) and observing the time evolution of the surface distortion by probing the optic in reflection. Data are taken by examining the output of the wavefront sensor twice per second starting 0.5 seconds after the shutter is triggered open and ending after a total of 300 seconds have elapsed. Twelve runs are taken, and the results are averaged at each measured time t_n after normalizing to the measured absorbed beam power for each run. The resulting averaged data are shown in figure 6-4.

Amorphous fused silica is isotropic, which we will assume when constructing the fitting functions (so, instead of four in the general case, there are now simply two: $\mathbf{a}(t)$ and $\mathbf{b}(t)$). To recover the required parameters, a χ^2 merit function is constructed and numerically minimized over α and k :

$$\chi^2(\alpha, k) = \sum_{n=1}^{600} \left(\frac{a_n - \frac{\alpha}{k} \mathbf{a}(t_n k)}{\sigma_n^a} \right)^2 + \left(\frac{b_n - \frac{\alpha}{k} \mathbf{b}(t_n k)}{\sigma_n^b} \right)^2$$

where the errors σ_n^a are the standard deviation of the data from all twelve runs at the measured times t_n . The model with the best fit parameters are shown as solid lines in figure 6-4. The statistical errors on the measured quantities α and k are obtained by computing the 2×2 covariance matrix for this χ^2 function. The best-fit parameters as well as their total errors (the root of the diagonal elements of the covariance matrix) are listed in table 6.2.

To measure the thermo-optic coefficient, the test optic is tilted so that it is probed




Figure 6-4: Fused Silica reflected data.

in transmission, and the test is repeated. Now, however, we examine the wavefront distortion induced by the pump on the periphery of the aperture (*not* the slope) in the steady-state. Also, since we're looking at silica in transmission, where the thermo-optical optical path distortion is $\sim 30\times$ larger than the surface distortion, the measurement is performed at a much lower pump power (about 200 mW). The measurement is performed four times, the region around the pump is thrown out, the wavefront data are averaged, and the local statistical error is computed for each Shack-Hartmann pixel of the wavefront sensor. The data are then least-squares fit to the model to obtain the thermo-optic coefficient $\frac{dn}{dT}$. The data and the resulting best-fit model are shown in figure 6-5, and the best-fit value of $\frac{dn}{dT}$ as well as the statistical error are listed in table 6.2.

Fused Silica Results			
Parameter	Value	Error	Units
k	1.44	0.02	W/m/ $^{\circ}$ K
α	0.55	0.03	ppm/ $^{\circ}$ K
$\frac{dn}{dT}$	9.0	0.3	ppm/ $^{\circ}$ K

Table 6.2: Measured materials parameters for fused silica.




Figure 6-5: Fused Silica steady-state transmitted phase map. Data are marked by “+” markers and the fit is the solid gray line. Data taken around the pump beam ($r < 2.6$ mm) are thrown out.




Figure 6-6: Sapphine crystal structure, adapted from [13].

6.2.3 Sapphire Measurement

The measurement for crystalline sapphire is the same as that for silica, although we can no longer assume isotropy in the model. Sapphire is a crystal with 3-fold rotational symmetry (trigonal symmetry) about the “C” axis of symmetry. A diagram of the sapphire unit cell is shown in figure 6-6. Note that the sapphire structure is very close to hexagonal (6-fold rotational symmetry): close enough that the elastic moduli and elasto-optic matrix elements that manifest the trigonal symmetry are often ignored [23, 32], and the crystal structure is assumed hexagonal. The materials parameters of a hexagonal crystal are cylindrically symmetric about the C-axis of the crystal, thus leading to two thermal conductivities, thermal expansion coefficients, and thermo-optic coefficients (one along the symmetry axis and the other perpendicular to it). Nominally, we should be able to estimate all of these parameters with the identical measurement procedure as before on a sapphire sample whose C axis lies along the optical axis, but we’ll find that the error in estimating so many parameters is rather large. To help remove these extra degrees of freedom and improve the qual-

ity of the results, we'll perform additional steady-state measurements in reflection and transmission on an identically-sized sample whose C axis is perpendicular to the optical axis which, in comparison to the C-axis data, yields a measurement of the relative thermal expansion coefficients and thermal conductivities, respectively.

C-axis Data

The sample in this measurement is a sapphire cylinder, 3" in diameter and 1" thick, with the C-axis oriented perpendicular to the optical surfaces (i.e., parallel to the optical axis). The surfaces of the sample are polished optically flat with a wedge of $\sim 1^\circ$, and neither surface is coated. The test proceeds as described in the fused silica case, although only 100 seconds of data are taken in each run as the steady-state is quickly reached due to sapphire's high thermal conductivity. All four anisotropic fitting functions $\mathbf{a}_r(t)$, $\mathbf{a}_z(t)$, $\mathbf{b}_r(t)$, and $\mathbf{b}_z(t)$ are numerically calculated (shown previously in figure 6-3) and fit to the data by minimizing the χ^2 merit function:

$$\chi^2(\gamma_z, \gamma_r, k_z, k_r) = \sum_{n=1}^{200} \left(\frac{a_n - \frac{\gamma_r}{k_z} \mathbf{a}_r(t_n k_z) - \frac{\gamma_z}{k_r} \mathbf{a}_z(t_n k_r)}{\sigma_n^a} \right)^2 + \left(\frac{b_n - \frac{\gamma_r}{k_z} \mathbf{b}_r(t_n k_z) - \frac{\gamma_z}{k_r} \mathbf{b}_z(t_n k_r)}{\sigma_n^b} \right)^2.$$

The data as well as the model with the best-fit parameters are shown as solid lines in figure 6-7. The statistical errors on the four measured quantities γ_z , γ_r , k_z , and k_r are obtained by computing the 4×4 covariance matrix for this χ^2 function. The best-fit parameters as well as their total errors (the root of the diagonal elements of the covariance matrix) are listed in table 6.3. The uncertainties in the fitted parameters here are considerably larger than in the Fused Silica case, mainly due to the fact that we are fitting more parameters to the data, as well as the data are considerably noisier (sapphire dissipates the heat deposited by the pump beam much more effectively, thus reducing our signal to noise, and the thermal distortion also takes hold much quicker, thus reducing the total number of data points recorded in the interesting region between "cold" and "hot"). To refine these measurements,

C-axis Sapphire Surface Data

2.6 mm pump beam, probed interval $r=(6\text{ mm}, 13\text{ mm})$




Figure 6-7: C-axis Sapphire reflected data.

we'll examine the steady-state wavefront distortions in reflection in transmission for an identically sized M-axis sapphire sample (i.e., the M-axis of the crystal now lies along the optical axis).

Initial Sapphire Results			
Parameter	Value	Error	Units
k_m	36.8	2.6	W/m/°K
k_c	35.4	4.5	W/m/°K
α_m	6.15	1.19	ppm/°K
α_c	5.68	0.89	ppm/°K

Table 6.3: Initial measured materials parameters for sapphire.

M-axis Data

The sample in this measurement is again a sapphire cylinder, 3" in diameter and 1" thick, with the M-axis now oriented parallel to the optical axis (i.e., the C-axis is perpendicular to the optical axis). The surfaces of the sample are polished optically flat with a wedge of $\sim 1^\circ$, and neither surface is coated. Since the crystal symmetry

axis is now perpendicular to the cylindrical symmetry axis of the optic, we must utilize the three dimensional finite element model.

Comparing the size of the M-axis steady state thermoelastic deformation with that seen in the C-axis will give us a measure of the relative size of α_c compared to α_m . Furthermore, examining the ellipticity of the steady-state phase map of transmitted beams now allows us to determine the relative size of k_c as compared to k_m . These measurements will further constrain the results found in the previous section. Finally, using the best fit parameters consistent with all the data we have so far taken, we may recover the thermo-optic coefficient $\frac{dn}{dT}$ by looking at the amplitude of the steady-state phase maps in M-axis sapphire.

Figure 6-8 shows the steady-state reflected phase maps for the C-axis sample compared to the M-axis sample. Both distortions are seen to be very circular, although the M-axis sample's distortion is about 10% smaller than that seen in the C-axis sample. Fitting the relative magnitude to that which is predicted in the model, yields the result:

$$\frac{\alpha_c}{\alpha_m} = 1.11(\pm 0.05) \quad (6.1)$$

Figure 6-9 shows the steady-state transmitted phase maps for probe polarizations along the C-axis as well as perpendicular to the C-axis. The elasto-optic effect serves to ellipticize both phase maps, although this ellipticity is much more prominent for polarizations perpendicular to the C-axis. If the thermal conductivities in each direction were identical, then the semi-minor axis of the elliptical phase map would track the polarization and the net ellipticity would not appreciably change (recall figure 3-8 and the corresponding discussion back in Chapter 3 showing the elasto-optic effect in sapphire). However, a larger C-axis thermal conductivity would tend to “wash out” the optical path distortion along the C-axis regardless of the probe polarization, thus amplifying the ellipticity for polarizations perpendicular to the C-axis, while nullifying the ellipticity for the orthogonal polarization. Fitting the measured ellipticities




Figure 6-8: C and M-axis steady-state reflected phase map.

to that predicted by the model yields the result:

$$\frac{k_c}{k_m} = 1.08(\pm 0.03). \quad (6.2)$$

Sapphire Results

Using the constraints found from the M-axis data in equations 6.2 and 6.1, we re-fit the C-axis data to over k_c , k_m , α_c and α_m . With these parameters now known, we then fit to the transmitted M-axis data to find $\frac{dn}{dT}$. Assuming the elastic moduli and elasto-optic matrix elements listed in table A.2, we find the following parameters for sapphire and the corresponding uncertainties measured around the mean temperature $T_\infty \approx 310^\circ\text{K}$ listed in table 6.4.

M-axis Sapphire Thermal Lens (Measured)

Contour Interval 10 nm/W, Pump Beam Waist=2.6mm




Figure 6-9: Steady-state phase map for transmission through M-axis sapphire. Probe polarization is perpendicular to the C-axis on the left, and parallel to the C-axis on the right.

M-axis Sapphire Transmitted Data




Figure 6-10: M-axis steady-state transmitted phase map (blue crosses), and the principal axes of the best fit model to the data (red lines).




Figure 6-11: Setup of the heating ring experiment. The probe examines the test optic in transmission, and the pump shines on the far optical surface.

Sapphire Results			
Parameter	Value	Error	Units
k_m	36.0	0.5	W/m/°K
k_c	38.9	0.5	W/m/°K
α_m	5.1	0.2	ppm/°K
α_c	5.6	0.2	ppm/°K
$\frac{dn}{dT}$	7.2	0.3	ppm/°K

Table 6.4: Measured materials parameters for sapphire.

6.3 Heating Ring Thermal Compensation

This experiment tests shielded heating ring thermal compensation by demonstrating the compensation of a LIGO-like thermal distortion induced by the pump laser. A diagram of the setup is in figure 6-11. The test optic is a flat-flat cylindrical fused silica test mass, 10 cm in diameter and 8 cm thick with a 1° wedge, and the probed aperture is the same as before (3 cm in diameter). The optic is probed in transmission for this experiment.

The shielded ring compensator consists of a 5 mm diameter solenoid of thin nichrome wire snaked around a core of ceramic fish-spine beads shaped by a frame of music wire into a circle of diameter 13.4 cm. The ring itself is mounted to a square aluminum frame at 4 points via alumina ceramic grommets held in 10 mil shim steel mounts. Ceramic is used at the mounting point because it is an electrical insulator with a melting point higher than nichrome, and shim steel is used to carry the grommets because of its very low effective thermal conductivity compared to that of ceramic and aluminum. The ring is powered by an adjustable DC power supply (100 Watts maximum output) whose output current and voltage are monitored to give the total power input to the ring. For the coaxial shield, a hollow aluminum cylinder with inner diameter 10.1 cm and outer diameter 11.4 cm is turned down at a 10° angle to a sharp edge located at the inner diameter nearest the optic (so the cross-section of the shield looks like a right triangle). The shield is carefully covered in high quality aluminum foil to minimize the absorption of ring radiation. The thermal time constant of the ring itself is measured out of vacuum by examining the resistance of a




Figure 6-12: Measured intensity profile of the pump beam cut along the x and y axes. The best fit Gaussian beam, shown by the grey line, has a waist radius of 1.45 cm.

thermistor placed inside the solenoid winding after the ring power supply is abruptly switched on, and is found to be about 300 seconds.

The pump beam is shaped by a single positive lens to the size of the probed aperture (3 cm diameter). The beam travels through the galvanometer scanners near its focus, thus allowing us to control the position of the beam on the optic. The probe beam's size is measured as described for large beams in §5.1 (by scanning it over a small thermopile), the results of which are shown in figure 6-12.

The ring center is located 11.9 cm from the optic's face, and the sharp edge of the coaxial shield is placed 3.4 cm closer (i.e., 8.5 cm away from the optic's face). Because the ring is not infinitesimal (it's 5 mm thick), we'll model it as an equal sum of two rings at coordinates $R_r = 6.7$ cm, $H_r = 11.65$ cm and $R_r = 11.9$ cm, $H_r = 6.45$ cm. We'll assume that all of the power provided to the ring is radiated away, as the only other means of dissipating the power is through the ceramic contact mounts, up the thin steel leads crimped to the narrow wires of the nichrome solenoid (the alloys of

steel and nichrome both have low thermal conductivities), and in the low resistance steel cable running from the DC power supply to the vacuum tank (the resistance of these leads is measured to be over 100 times smaller than the resistance of the ring itself).

6.3.1 Heating Ring Results

After all the relevant calibrations are performed, the pump is switched on and the steady-state phase map of the pump induced distortion is measured after 15 minutes (the calculated time constant for the aperture probed). The ring is switched on, and the phase map is examined after 20 minutes (time constant of the aperture plus the ring time constant). The servo loop here is closed by the author, where pump power is adjusted to flatten the wavefront generated by the ring compensator (the relevant thermal time constant for changing ring power is longer compared to that for changing the laser power). After the steady-state is again reached, the powers on the ring and pump are noted and the test is repeated four times.

During the course of the measurements, it was noticed that the ring compensator would slightly heat the back side of the rigid test optic mount, thus inducing over 30 micro-radians of tilt across the measured aperture (about 1 μm of optical path distortion across the 3 cm aperture). A separate aluminum plate, with a 10.2 cm diameter hole cut in it to leave the face of the optic exposed, was placed between the optic mount and compensator and the resulting tilt was seen to be greatly reduced. It was not completely eliminated, however, due to the fact that the optic itself is wedged, thus a global increase in the optic's temperature results in some tilt, but was at a level where it could be cleanly removed from the wavefront data.

The results of the experiment for 50 mW of pump power absorbed are shown in figures 6-13 and 6-14. The calculated TEM_{00} scatter for the pump beam is seen to be reduced by nearly 2 orders of magnitude, and the residual optical path distortion is well explained by both imperfections in the pump beam as well as a slight ellipticity of the ring (shown as dashed grey lines in figure 6-14).

The one major discrepancy between the compensation results and the model is




Figure 6-13: Results of the shielded heating ring test. The left diagram is the pump-induced distortion and the shielded ring induced correction is on the right. The calculated TEM₀₀ scatter for a Gaussian beam whose waist is the size of the pump beam is listed at the top of each diagram.

in the ring power necessary to generate such a correction. The full power of the DC power supply is used (100 Watts), but the power necessary to generate such a correction is only 70 Watts. It was noticed during the course of the experiments that the ring glows a dull orange at full power, thus some amount of power is being emitted in the transmissive band of the glass (wavelengths less than about 4 microns). If the emissivity of the nichrome ring is lower than expected (about 1), then, for equivalent radiated powers, the resulting blackbody radiation occurs at shorter wavelengths which are not absorbed in the optic's surface (from the Stefan-Boltzmann law $P_{rad} \propto \epsilon\sigma(T^4 - T_{\infty}^4)$, if emissivity ϵ is smaller then $(T^4 - T_{\infty}^4)$ must be larger to radiate equivalent power). To test this hypothesis, the steady-state optical path distortion induced by the ring alone is examined at intervals of 10 Watts of electrical input power, and the results are fit to that which is expected from the model (with the optic's emissivity fixed at 0.9, as we previously measured at 10.6 microns). The results are shown in figure 6-15. At low ring power, the optical path distortion looks almost




Figure 6-14: Shielded heating ring results as compared to the finite element model's prediction. The "+" marker represents the measured distortion due to the pump beam alone, and the "x" marker represents the result under shielded ring compensation. The solid lines are what is expected according to the finite element model. The dashed lines represent the predicted envelope of wavefront data if the ring is slightly elliptical (ellipticity of ~ 0.1).




Figure 6-15: Measured heating ring coupling efficiency versus electrical power input to the ring per unit length.

exactly as expected (i.e., the coupling efficiency is nearly 1), but quickly deteriorates as ring power is increased. To roughly model these results, we assume a constant but unknown ring emissivity ϵ_r , and assume that the optic has an emissivity of 0.9 at wavelengths greater than $\lambda_c = 4 \mu\text{m}$ and is otherwise purely transmissive. Utilizing the Planck distribution, the coupling efficiency is thus:

$$\frac{P_{eff}}{P_{in}}(T) \approx \frac{\int_{\lambda_c}^{\infty} \frac{\lambda^5 d\lambda}{e^{\frac{hc}{\lambda kT}} - 1}}{\int_0^{\infty} \frac{\lambda^5 d\lambda}{e^{\frac{hc}{\lambda kT}} - 1}}$$

and the temperature is related to the input power as

$$\frac{P_{in}}{2\pi R_r} = \epsilon_r \sigma (T^4 - T_{\infty}^4) \pi d_r.$$

where $d_r = 5 \text{ mm}$ is the ring thickness. The results for $\epsilon_r = 1$ and the best-fit ring emissivity $\epsilon_r = 0.65$ are shown as the grey lines in figure 6-15.




Figure 6-16: Setup of the scanning laser test. The optic is probed in transmission and the pump, a heavily attenuated portion of the immediate CO₂ laser output, is focused to a $300\ \mu\text{m}$ spot on the probe-side surface of the test optic. Actuation is performed on the opposite surface of the test optic on the pattern shown in figure 5-3 with powers determined by the measured wavefront distortion induced by the pump.

The upshot of this is that ring power parameters greater than $2.5\ \text{W/cm}$ are difficult to achieve with a 5 mm nichrome ring, in stark contrast to the rough limit of $11\ \text{W/cm}$ that we calculated in Chapter 4. This is by no means a hard limit, however, as simply increasing the diameter of the ring (hence its effective surface area) can increase this limit, although there will be a trade in the resulting quality of correction (recall the ring placement tolerance examined in figure 4-9).

6.4 Scanning Laser Thermal Compensation

This experiment tests scanning laser thermal compensation by demonstrating the compensation of a point absorber off the optical axis. A diagram of the setup is in figure 6-16. As in the shielded ring experiment, the test optic is the same flat-flat cylindrical fused silica test mass which is probed in transmission over the same (3 cm diameter) aperture.

The scanning laser compensator consists of a CO₂ laser beam of fixed shape (0.19 cm waist radius) scanned with a pair of mirror galvanometers over a hexagonal pattern with 0.5 cm spacing. The galvanometers are located approximately 50 cm from the face under actuation, and the entire 3 cm probed aperture is covered by the pattern. The choice of this pattern was previously discussed in Chapter 5 (in particular, see figure 5-3). The beam passes through an acousto-optic modulator (AOM), which facilitates the control of the laser power passing through the mirror galvanometers. There are three signal outputs provided by a desktop computer: the X galvanometer input voltage, the Y galvanometer input voltage, and the AOM drive voltage. The mirror galvanometer positions are proportional to drive voltage, which are calibrated by shining a red diode laser beam through the scanners and recording the beam's position on a piece of paper placed at the location of the optic's face versus various input voltages. The laser power transmitted through the AOM goes as \sin^2 of the AOM drive voltage, the period of which is calibrated beforehand by examining the AOM thru-put versus drive voltage.

The galvanometers continuously execute the fixed pattern shown in figure 5-3, with a pattern repetition frequency of about 4 Hz (by equation 5.7 the local time constant for this beam size on fused silica is about 5 seconds). The laser beam dwells at each scan point for 2 ms, then switches off for the following 2 ms where it travels to the next point and repeats (the elapsed time for point-to-point movement is ~ 0.5 ms, measured via capacitive position sensors built into the galvanometers).

The AOM voltages (i.e., laser powers) are determined by the wavefront distortion registered on the Shack-Hartmann wavefront sensor, taken at 15 minute intervals

(approximately the thermal time constant of the 3 cm aperture). This distortion is numerically integrated against the 38 modeled influence functions (which are evaluated over the coarse wavefront sensor coordinates) and the modeled inverted poke matrix is applied yielding the required powers. These powers are added to the powers found in the previous iteration, and then piston is applied (or subtracted) until all powers are positive and one is identically zero. These positive powers are translated into AOM voltages, and the new set of AOM voltages is implemented without interruption. 15 minutes are then allowed to elapse and the process is repeated.

The wavefront distortion is generated by focusing a steady CO₂ beam to a small spot somewhere on the optical aperture. The beam is picked off from the full 10 Watt CO₂ output, and is attenuated using a pair of uncoated ZnSe brewster polarizers to about 20 mW. The low power beam is directed into the vacuum tank, and encounters a 10" ZnSe lens which focuses the beam to a $\sim 300 \mu\text{m}$ spot on one face of the optic within the probed aperture.

6.4.1 Scanning Laser Results

After all the relevant calibrations are performed (laser thru-put, galvo calibration, etc.) the laser is stabilized near its maximum power (about 7 Watts), and the pump beam is switched on. After 15 minutes, when the distortion reaches the steady-state, wavefront data are taken by the Shack-Hartmann sensor and the scanning laser servo loop is switched on. The initial wavefront data are shown on the left in figure 6-17. The additional lobes in the phase map are caused by "ghost" spots in the CO₂ beam: i.e., weak reflections off the uncoated surfaces of the two brewster polarizers used to attenuate the beam. The initial powers are calculated from the model, and are indicated by the diameters of the circles in the left diagram. The servo loop is then abruptly switched on.

For a few minutes after the servo loop is abruptly switched on, the resulting wavefront develops a tremendous amount of curvature (over 200 nm across the measured aperture) as heat rushes into the aperture from the large amounts of power delivered on the periphery. After about 15 minutes, the temperature within the aperture in-

creases relative to the rest of the optic and this curvature disappears (the net radial heat flow entering the aperture caused by the large actuator powers on the aperture's periphery equals the radial heat flow exiting the aperture due to its elevated temperature relative to the rest of the optic). This problem can be easily circumvented by slowly ramping up the actuator powers over the thermal time constant of the aperture; however, this was not attempted in this experiment.

At 15 minute intervals thereafter, the wavefront is automatically reexamined, new powers are calculated, and the AOM voltages are updated without interruption. The quality of the wavefront correction, in terms of effective TEM₀₀ scatter for a beam the size of the probed aperture, was seen to slowly improve, finally reaching a minimum after about 2 hours (approximately the time constant of the entire optic). The resulting corrected wavefront is shown on the right of figure 6-17, and figure 6-18 shows the initially calculated (at t=0) and the finally executed (at t=6300s) powers to achieve this wavefront correction. The final powers are typically within a few percent of the initially calculated powers; however, there are a few points (primarily those with initially very low powers) that differ by a factor of 2 or more. This could be due to imperfections in the local magnification at those points in the aperture, imperfections in the scan pattern (i.e., the actuator beam doesn't shine exactly on those points), or it could be due to lower amounts of 10.6 μm absorption in those areas of the surface (although we expect this to only be a few percent).

In the end, the model is seen to agree well enough with the experiment to implement a stable servo loop for thermal wavefront correction. However, it might be simpler and more robust to use the model to examine the *suitability* of a particular actuation basis only, and instead measure each influence function and compute the resulting measured poke matrix. Although it would take twice the full time constant of the aperture (15 minutes here) to measure each influence function, symmetry could be used to reduce the measurement to only a fraction of the total number of influence functions. Most significantly, measuring the influence functions would negate the need for most of the precise calibrations that were necessary to perform this experiment (mainly the CO₂ system thru-put and reflectivity), and would account for

Scanning Laser Compensation of ~20 mW absorbed on 300 μ m spot

0.19cm compensator beam waist, 1nm contour interval

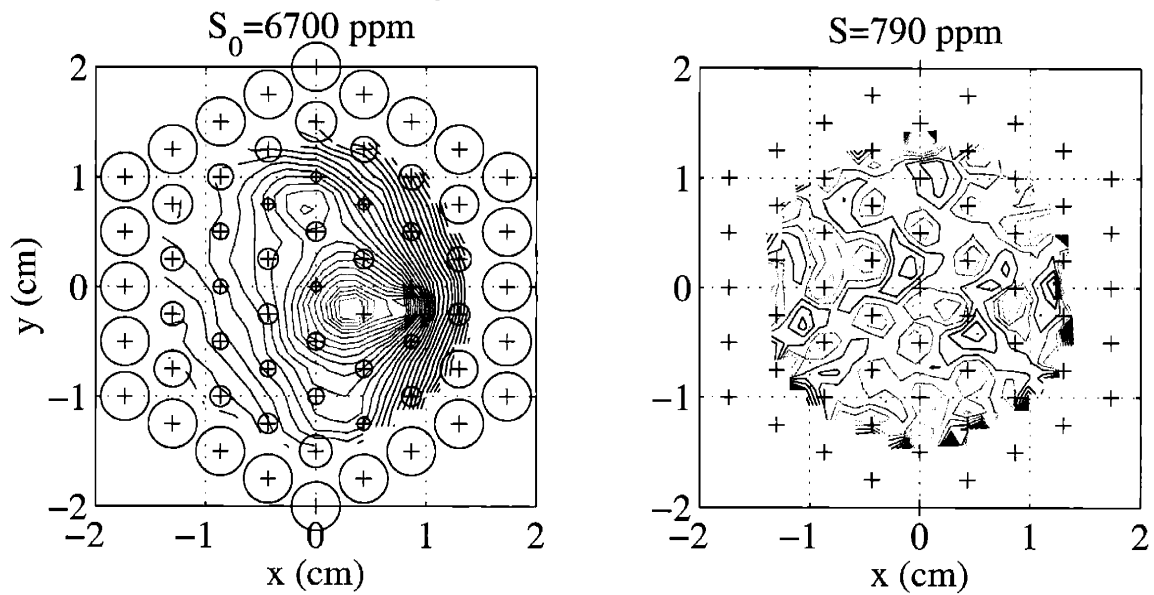


Figure 6-17: Results of the scanning laser test. The points of the scan pattern are denoted by “+” marker in each diagram. The measured wavefront distortion induced by the pump beam is shown at left, and the diameters of the circles represent the powers calculated to correct the wavefront distortion shown. The right diagram shows the resulting compensated wavefront, taken after 5 complete iterations of the servo loop (about 2 hours). The TEM_{00} scatter for a Gaussian beam whose waist is the size of the measured aperture is listed at the top of each diagram.




Figure 6-18: Initial and final powers in the scanning laser test, for the two wavefronts shown in figure 6-17. The initial actuator powers are indicated by the diameters of the black circles, and the final powers are shown by the grey circles. When no grey circle is seen, it is simply overlapped by the black circles (i.e., initial and final powers are identical).

imperfections in the probe optical system (i.e., fluctuations in effective magnification across the probed aperture due to spherical aberration or astigmatism).

Conclusion

In order for Advanced LIGO to reach its design sensitivity goal, where anticipated event rates are such that a major scientific payoff is guaranteed (even if nothing is seen), the necessary optical power must be increased to such a level where thermal distortions induced by realistic levels of optical absorption begin to significantly detract from the sensitivity of the instrument. For a Fused Silica Advanced LIGO this is a very serious issue, precluding the attainable sensitivity to a level not much better than that of Initial LIGO. For a Sapphire Advanced LIGO the problem does not appear as serious by this analysis (for optimistic assumptions on absorption and ignoring effects on signal recycling, see below), but is still enough that a factor of ~ 2 is lost in sensitivity. This thesis has presented a way of correcting these distortions through the use of thermally adaptive optics, two modes of which are discussed: (1) a simple heating ring with some carefully placed shielding to compensate distortions induced by uniform absorption, and (2) a scanning heating beam to tailor the correction in the case of nonuniform optical absorption. Careful modeling and analysis has demonstrated that, for realistic amounts of uniform absorption, heating ring thermal compensation can adequately repair these distortions and restore the instrument's full sensitivity. Some degree of nonuniform absorption can also be compensated via scanning laser thermal compensation, although one must be extremely careful of the displacement noise that actuator laser intensity noise will introduce into the instrument. A successful experiment was conducted to verify the validity of the single-optic model as well as the basic technical feasibility of both modes of thermally adaptive optics.

This thesis did not consider signal recycling, a means by which the integrating

effects of a large arm cavity storage time can be bypassed for a relatively narrow band of gravitational wave signal frequencies. This scheme is critical for Advanced LIGO, and relies on the ability to store gravity wave signal sidebands (which are nonresonant in the high finesse arm cavities) in a signal recycling cavity formed by a partially reflective mirror at the dark port and the high reflective surfaces of the input test masses (and thus shares the path with the power recycling cavity between the beamsplitter and input test masses). Since the effective power recycling cavity gain for the control sidebands is so seriously reduced by thermal lensing in the input test masses, it seems likely that the gravity wave sidebands will experience a similar falloff in signal recycling cavity gain, thus possibly diminishing any signal recycling benefit. If this is the case, then thermal compensation of input test mass distortions is absolutely necessary to maintain the full gain of both signal and control sidebands. Simulations of this effect are forthcoming in the next release of Melody[35].

This work also has not discussed a means by which thermally adaptive optics may be implemented in Advanced LIGO in a closed loop manner, which is a necessary condition for this work to truly be considered an incarnation of “adaptive optics”. There is an obvious brute force solution to the problem, where each optic of concern (i.e., an input test mass and/or its compensator plate) is resolved slightly off-axis by a dedicated wavefront sensor (e.g., a Shack-Hartmann sensor) whose signal is used to drive the thermal actuator. A more elegant solution, perhaps, is to use the existing pickoffs in the instrument to spatially resolve the phase of sideband light. Such a scheme exists in Initial LIGO to sense and control tilt (with quad RF photodiodes)[14] and focus (with “bullseye” RF photodiodes)[28], so a finer spatial sensor may make it possible to sense higher order distortions which can be used to drive the thermal actuator¹.

¹Work on constructing such a sensor, termed a “phase camera” or “sideband camera”, which utilizes a pair of mirror galvanometers to rapidly scan the beam one wishes to measure over a simple RF photodiode, is underway at the MIT LIGO Laboratory by David Ottaway, Rana Adhikari, and Keisuke Goda [30].

Appendix A

Tables of Parameters

A.1 Materials Parameters

Fused Silica (Amorphous SiO ₂)				
Parameter	Symbol	Value	Units	Reference
Thermo-Optic Coefficient (633nm)	$\frac{dn}{dT}$	8.7	ppm °K ⁻¹	This Work
Thermal Expansion	α	0.55	ppm °K ⁻¹	This Work
Thermal Conductivity	k	1.38	W/m/°K	[23]
Young's Modulus	E	72.80	GPa	[23]
Poisson Ratio	ν	0.170	-	[23]
Elasto-Optic Coefficient	p_{11}	0.121	-	[23]
Elasto-Optic Coefficient	p_{12}	0.270	-	[23]
Emissivity (10.6 μ m)	ϵ	0.90	-	This Work
Density	ρ	2196	Kg/m ³	[23]
Specific Heat	c	740	J/kg/°K	[23]
Refractive Index (633nm)	n	1.4497	-	[23]

Table A.1: Fused silica (SiO₂) material parameters.

Sapphire (Crystalline Al₂O₃)				
Parameter	Symbol	Value	Units	Reference
Thermo-Optic Coefficient (633nm)	$\frac{dn}{dT}$	7.2	ppm °K ⁻¹	This Work
Thermal Expansion	α_o	5.6	ppm °K ⁻¹	This Work
Thermal Expansion	α_e	5.1	ppm °K ⁻¹	This Work
Thermal Conductivity	k_o	39	W/m/°K	This Work
Thermal Conductivity	k_e	36	W/m/°K	This Work
Elastic Moduli	c_{11}	499.1	GPa	[23]
Elastic Moduli	c_{33}	496.8	GPa	[23]
Elastic Moduli	c_{13}	112.2	GPa	[23]
Elastic Moduli	c_{12}	164.0	GPa	[23]
Elastic Moduli	c_{44}	147.4	GPa	[23]
Elasto-Optic Coefficient	p_{11}	-0.23	-	[23]
Elasto-Optic Coefficient	p_{33}	-0.20	-	[23]
Elasto-Optic Coefficient	p_{12}	-0.03	-	[23]
Elasto-Optic Coefficient	p_{13}	0.02	-	[23]
Elasto-Optic Coefficient	p_{31}	-0.04	-	[23]
Emissivity (10.6μm)	ϵ	0.89	-	This Work
Density	ρ	3970	Kg/m ³	[23]
Specific Heat	c	775	J/kg/°K	[23]
Refractive Index (633nm)	n	1.765	-	[23]

Table A.2: Sapphire (Al₂O₃) material parameters. Axis 3 represents the ordinary (C) axis.

A.2 Instrument Parameters

Input Laser and Length Parameters					
Parameter		Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
Input Laser					
Laser Wavelength	λ	1064	1064	1064	nm
Laser Input Power	P_i	6	125	80	W
Lengths					
Arm Length	L	3995	3995	3995	m
PRC Common Length	l_c	9.4	9.4	9.4	m
PRC Differential Length	l_d	0.2	0.2	0.2	m

Table A.3: LIGO input laser and length parameters.

Mirror Parameters				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
End Test Mass				
ETM Material	SiO2	Al2O3 (e)	SiO2	-
ETM Radius	0.125	0.157	0.190	m
ETM Thickness	0.10	0.13	0.154	m
ETM Transmission	5.5	1	1	ppm
ETM Curvature	8.5	60	60	km
ETM Surface Absorption	0.5	0.5	0.5	ppm
ETM HR Loss	30	15	15	ppm
Input Test Mass				
ITM Material	SiO2	Al2O3 (e)	SiO2	-
ITM Radius	0.125	0.157	0.190	m
ITM Thickness	0.10	0.13	0.154	m
ITM Transmission	0.03	0.005	0.005	-
ITM Curvature	14.6	60	60	km
ITM Substrate Absorption	4	30	0.5	ppm/cm
ITM Surface Absorption	0.5	0.5	0.5	ppm
ITM HR Loss	30	15	15	ppm
ITM AR Loss	800	600	600	ppm
Beamsplitter				
BS Material	SiO2	SiO2	SiO2	-
BS Radius	0.125	0.175	0.175	m
BS Thickness	0.06	0.06	0.06	m
BS Substrate Absorption	2.5	0.5	0.5	ppm/cm
BS Surface Absorption	0.5	0.5	0.5	ppm
BS AR Loss	100	100	100	ppm
Recycling Mirror				
RM Material	SiO2	SiO2	SiO2	-
RM Radius	0.125	0.133	0.133	m
RM Thickness	0.10	0.10	0.10	m
RM Transmission	0.03	0.06	0.06	-
RM Curvature	15.8	34	41	km
RM Substrate Absorption	13	20	20	ppm/cm
RM Surface Absorption	0.5	0.5	0.5	ppm
RM HR Loss	30	15	15	ppm

Table A.4: LIGO mirror parameters. The term “loss” is used to describe the sum of aperture diffraction loss, diffuse scatter loss, and reflection loss (for anti-reflection coated surfaces). The entry “Al2O3 (e)” means that the extraordinary axis of the crystal lies along the optical axis (i.e., the C-axis (the ordinary axis) is parallel to the optical surface).

Derived Cavity Parameters (Cold)				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
Power Recycling Cavity				
\mathcal{G}_{PRC}	62	23	28	-
P_{PRC}	0.4	2.1	1.3	kW
g_{ITM}	1.00094	1.00027	1.00023	-
g_{RM}	0.99941	0.99972	0.99977	-
w_0	Undefined	0.043	0.045	m
z_0	Undefined	1.0	1.0	km
Arm Cavity				
\mathcal{G}_{ARM}	130	815	810	-
P_{ARM}	26	850	530	kW
g_{ITM}	0.72637	0.93342	0.93342	-
g_{ETM}	0.53000	0.93342	0.93342	-
w_0	0.036	0.059	0.059	m
z_0	1.2	2.0	2.0	km

Table A.5: Derived LIGO cavity parameters (Cold). The quantity z_0 is defined as the location of w_0 relative to the ITM, with $+z$ towards the ETM.

Heated Arm Cavities				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
Input Test Mass				
P_a (coating)	12	400	260	mW
P_a (substrate)	16	800	10	mW
R_0	14.6	60	60	km
R_t	-410	-63	-120	km
End Test Mass				
P_a (coating)	12	400	260	mW
P_a (substrate)	0.0	0.3	0.0	mW
R_0	8.5	60	60	km
R_t	-950	-190	-120	km
Hot Arm Cavity				
g_{ITM}	0.73611	0.99683	0.96671	-
g_{ETM}	0.53421	0.95444	0.96671	-
w_0	0.036	0.075	0.070	m
z_0	1.2	0.25	2.0	km

Table A.6: Parameters for heated arm cavities using the approximations from Chapter 2. Materials and instrument parameters are taken from tables A.2, A.1, A.4 and A.5.

Heated Power Recycling Cavities				
Parameter	Initial LIGO Typical	Advanced LIGO Sapphire	Advanced LIGO Silica	Units
Input Test Mass				
$P_a(\text{coating})$	12	400	260	mW
$P_a(\text{substrate})$	16	800	10	mW
$R_0(\text{effective})$	-10	-34	-41	km
R_t	23	27	6.5	km
S	8000	43000	430000	ppm
Beamsplitter				
$P_a(\text{coating})$	0.3	1.5	1.0	mW
$P_a(\text{substrate})$	3.0	3.0	1.9	mW
R_t^X	190	400	610	km
S^X	110	210	90	ppm
R_t^Y	-3500	-7000	-11000	km
S^Y	0.4	0.7	0.3	ppm
Recycling Mirror				
$P_a(\text{coating})$	0.2	1.0	0.7	mW
$P_a(\text{substrate})$	0.8	24	16	mW
R_0	15.8	34	41	km
R_t	-11000	-1300	-1900	km
S	0.0	20	10	ppm
Hot Power Recycling Cavity				
g_{ITM}	1.00052	0.99993	0.99878	-
g_{RM}	0.99941	0.99973	0.99978	-
w_0	0.019	0.013	0.009	m
z_0	70	-2.0	-7.9	m

Table A.7: Parameters for hot power recycling cavities, using the approximations of Chapter 2. Materials and instrument parameters are taken from tables A.2, A.1, A.4 and A.5.

Appendix B

Calculation of the Heating Ring Intensity Profile

Consider a small piece of heating wire dl oriented in some unit direction \hat{l} (so that $d\vec{l} = dl \hat{l}$) at position \vec{r}' which is radiating a total power P_{dl} . The thermal (incoherent) radiation pattern of $d\vec{l}$ is:

$$\vec{S} = \frac{P_{dl} (\vec{r} - \vec{r}')}{\pi^2 |\vec{r}' - \vec{r}|^3} |\cos \theta| = \frac{P_{dl} (\vec{r} - \vec{r}')}{\pi^2 |\vec{r}' - \vec{r}|^3} \left| 1 - \left(\frac{(\vec{r}' - \vec{r})}{|\vec{r}' - \vec{r}|} \cdot \hat{l} \right)^2 \right|^{\frac{1}{2}}$$

where θ is the angle between $(\vec{r} - \vec{r}')$ and $d\vec{l}$. If there is a small area $d\vec{A}$ with unit normal \hat{n} at a position \vec{r} which absorbs all the energy radiated from $d\vec{l}$ that is incident on $d\vec{A}$, the general form for power incident (absorbed) on dA due to $d\vec{l}$ is:

$$\begin{aligned} P_{dA} &= \left| \vec{S} \cdot d\vec{A} \right| \\ &= \frac{P_{dl} dA}{\pi^2 |\vec{r}' - \vec{r}|^2} \left| 1 - \left(\frac{(\vec{r}' - \vec{r})}{|\vec{r}' - \vec{r}|} \cdot \hat{l} \right)^2 \right|^{\frac{1}{2}} \left| \frac{(\vec{r}' - \vec{r})}{|\vec{r}' - \vec{r}|} \cdot \hat{n} \right|. \end{aligned} \quad (\text{B.1})$$

For a coaxial heating ring as described in Chapter 4, we have $\vec{r}' = R_r \hat{r}' + H_r \hat{z}$ and $\vec{r}' \cdot d\vec{l} = 0$.

B.1 Intensity Distribution on the Optic's Face




Figure B-1: Schematic of the coordinates used in the calculation of the intensity distribution of the heating ring on the face of the optic.

A schematic of the geometry is shown in figure (B-1). For a differential area element dA on the optic's face, we have $\hat{n} = \hat{z}$ and $\vec{r} = r\hat{r}$. For brevity, let $\vec{\rho} \equiv \vec{r}' - \vec{r}$. We now have:

$$|\vec{\rho}|^2 \equiv |\vec{r}' - \vec{r}|^2 = (R_r^2 + r^2 + H_r^2) - 2R_r r \cos \phi$$

where ϕ is the angle between \vec{r}' and \vec{r} . Equation B.1 now becomes:

$$P_{dA} = \frac{P_{dl} H_r dA}{\pi^2 \rho^3} \left(1 - \frac{(\vec{\rho} \cdot \hat{l})^2}{\rho^2} \right)^{\frac{1}{2}} = \frac{P_r d\phi h dA}{2\pi^3 \rho^4} ((R_r - r \cos \phi)^2 + H_r^2)^{\frac{1}{2}}$$

where we have utilized the fact that the total ring power P_r is related to the line element power P_{dl} by $P_{dl} = P_r \frac{d\phi}{2\pi}$. The power incident on dA due to the entire ring is then:

$$P_{total} = \frac{P_r H_r dA}{2\pi^3} \int_{-\pi}^{\pi} \frac{((R_r - r \cos \phi)^2 + H_r^2)^{\frac{1}{2}}}{\rho^4} d\phi$$

Thus, the power incident per unit area on the face of the optic is:

$$I_{face}(r) = \frac{P_r H_r}{2\pi^3} \int_{-\pi}^{\pi} \frac{((R_r - r \cos \phi)^2 + H_r^2)^{\frac{1}{2}}}{((R_r^2 + r^2 + H_r^2) - 2R_r r \cos \phi)^2} d\phi \quad (\text{B.2})$$

B.2 Intensity Distribution on the Optic's Edge




Figure B-2: Schematic of the coordinates used in the calculation of the intensity distribution of the heating ring on the outer edge of the optic.

A schematic of the geometry is shown in figure B-2. For a differential area element dA on the optic's outer edge, we have $\hat{n} = \hat{r}$ and $\vec{r} = R_o \hat{r} + z \hat{z}$. Once again, let $\vec{\rho} \equiv \vec{r}' - \vec{r}$. We have:

$$\begin{aligned} |\vec{\rho}|^2 &\equiv |\vec{r}' - \vec{r}|^2 = |R_r \hat{r}' - R_o \hat{r} + (H_r - z) \hat{z}|^2 \\ &= (R_r^2 + R_o^2 + (H_r - z)^2) - 2R_r R_o \cos \phi \end{aligned}$$

where ϕ is the angle between \vec{r}' and \vec{r} . Equation B.1 now becomes:

$$P_{dA} = \frac{P_{dl} dA}{\pi^2 \rho^3} \left(1 - \frac{(\vec{\rho} \cdot \vec{l})^2}{\rho^2} \right)^{\frac{1}{2}} (R_r \cos \phi - R_o)$$

$$= \frac{P_r d\phi (R_r \cos \phi - R_o) dA}{2\pi^3 \rho^4} \left((R_r - R_o \cos \phi)^2 + (H_r - z)^2 \right)^{\frac{1}{2}}$$

Glancing at figure B-2, it is easy to see that the only section of the ring which is not shielded from dA is the region where $|\phi| \leq \arccos\left(\frac{R_o}{R_r}\right)$. The power incident on dA due to the entire ring is then:

$$P_{total} = \frac{P_r dA}{2\pi^3} \int_{-\arccos\left(\frac{R_o}{R_r}\right)}^{\arccos\left(\frac{R_o}{R_r}\right)} \frac{(R_r \cos \phi - R_o) \left((R_r - R_o \cos \phi)^2 + (H_r - z)^2 \right)^{\frac{1}{2}}}{\rho^4} d\phi$$

Thus, the power incident per unit area along the edge of the optic is:

$$I_{edge}(z) = \frac{P_r}{2\pi^3} \int_{-\arccos\left(\frac{R_o}{R_r}\right)}^{\arccos\left(\frac{R_o}{R_r}\right)} \frac{(R_r \cos \phi - R_o) \left((R_r - R_o \cos \phi)^2 + (H_r - z)^2 \right)^{\frac{1}{2}}}{\left((R_r^2 + R_o^2 + (H_r - z)^2) - 2R_r R_o \cos \phi \right)^2} d\phi$$

B.3 Intensity Distribution for the Shielded Ring

Figure B-3 shows the geometry and coordinates for for the following calculation. The slope of the line connecting the surface element dA at radial coordinate r and the ring element dl at angular coordinate ϕ is simply the height of the ring divided by the distance in the plane of the optic's face from dA to dl . Similarly, the slope of the line connecting dA with the bottom edge of the metal shield in the direction of the ring element dl is the height of the shield divided by the distance in the plane of the optic's face from dA to the edge of the optic in the direction of dl . The element dA is not shielded from dl if the slope from dA to dl is smaller than the slope from dA to the shield, i.e.:

$$\frac{R_o^2 + r^2 - 2R_o r \cos \phi'}{R_r^2 + r^2 - 2R_r r \cos \phi} \leq \frac{H_s^2}{H_r^2}$$




Figure B-3: Schematic of the coordinates used in the calculation of the intensity distribution of the shielded heating ring.

where we have utilized the law of cosines to find the distances in the plane of the optic's face, and ϕ' is the angular coordinate of the bottom of the metal shield along the line of sight from dA to dl . To find $\cos \phi'$, we utilize similar triangles to find:

$$\frac{R_o^2 + r^2 - 2R_o r \cos \phi'}{R_r^2 + r^2 - 2R_r r \cos \phi} = \frac{R_o \sin \phi'}{R_r \sin \phi}$$

and $\cos \phi'$ is found to be:

$$\cos \phi' = \frac{r R_r^2}{R_o \rho^2} \sin^2 \phi \pm \sqrt{\left(\frac{r R_r^2}{R_o \rho^2} \sin^2 \phi \right)^2 - \left(\frac{r}{R_o} + \frac{R_o}{r} \right) \frac{r R_r^2}{R_o \rho^2} \sin^2 \phi + 1.}$$

where ρ is defined here as the distance in the plane of the optic's face from dA to the line element dl :

$$\rho \equiv R_r^2 + r^2 - 2R_r r \cos \phi.$$

Bibliography

- [1] A. Abramovici, W. Althouse, R. Drever, Y. Gürsel, S. Kawamura, F. Raab, D. Shoemaker, L. Sievers, R. Spiro, K. Thorne, R. Vogt, R. Weiss, S. Whitcomb, and M. Zucker. Ligo-the laser interferometer gravitational-wave observatory. *Science*, 256:325–333, 1992. LIGO-G920010-00-R.
- [2] A. Alexandrovski, M. M. Fejer, V. Kondilenko, and R. Route. Absorption studies in optical coatings and sapphire crystals. LIGO Scientific Collaboration Meeting, March 2002. LIGO-G020122-00-Z.
- [3] R. Beausolei and D. Sigg. Spatiotemporal model of the LIGO interferometer. *J. Opt. Soc. Am.*, 16:2990–3002, 1999. LIGO-P990004-00-D.
- [4] Robert Bennett. *Thermally Adaptive Optics in LIGO II*. Bachelor’s thesis, Massachusetts Institute of Technology, Department of Physics, February 2001.
- [5] L. Bildsten. Gravitational radiation and rotation of accreting neutron stars. *Astrophys. J. Lett.*, 501:L89–L93, 1998.
- [6] GariLynn Billingsley. LIGO core optics components website. <http://www.ligo.caltech.edu/gari/index.htm>, 2002. A compilation of all specifications and metrology data for the LIGO core optics.
- [7] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A*, 264:1–10, 1999.

- [8] H. Dimmelmeier, J. A. Font, and E. Müller. Gravitational waves from relativistic rotational core collapse in axisymmetry. *Class. Quantum Grav.*, 19:1291–1296, 2002.
- [9] D. Djambazov and D. Coyne. Transient thermal lensing and thermoelastic aberrations of mirrors. LIGO-T960185-00-R, October 1996. Mathematica Notebook, Version 2.2.
- [10] A. Freise, G. Heinzel, K. A. Strain, J. Mizuno, K. D. Skeldon, H. Lück, B. Willke, R. Schilling, A. Rüdiger, W. Winkler, and K. Danzmann. Demonstration of detuned dual recycling at the garching 30m laser interferometer. *Phys. Lett. A*, 277:135–142, 2000.
- [11] P. Fritschel, G. Gonzalez, B. Lantz, P. Saha, and M. Zucker. High power interferometric phase measurement limited by quantum noise and the application to the detection of gravitational waves. *Phys. Rev. Lett.*, 80:3181–3184, 1998. LIGO-P970009-00-R.
- [12] E. Gustafson, D. Shoemaker, K. Strain, and R. Weiss. Ligo scientific collaboration (lsc) white paper on detector research and development. Technical report, LIGO Scientific Collaboration, September 1999. LIGO-T990080-00-D.
- [13] Daniel C. Harris, Frederick Schmid, David R. Black, Ender Savrun, and Herbert E. Bates. Factors that influence mechanical failure of sapphire at high temperature. In Randal W. Tustison, editor, *Window and Dome Technologies and Materials V*, volume 3060, page 226. SPIE, SPIE-The International Society for Optical Engineering, 1997.
- [14] Y. Hefetz, N. Mavalvala, and D. Sigg. Principles of calculating alignment signals in complex resonant optical interferometers. *J. Opt. Soc. Am. B*, 14(7):1597+, 1997. LIGO-P960024-A-D.
- [15] P. Hello and J. Vinet. Analytical models of thermal aberrations in massive mirrors heated by high power laser beams. *J. Phys. France*, 51:1267–1282, 1990.

- [16] P. Hello and J. Vinet. Analytical models of transient thermoelastic deformations of mirrors heated by high power cw laser beams. *J. Phys. France*, 51:2243–2261, 1990.
- [17] Kerson Huang. *Statistical Mechanics*. John Wiley & Sons, New York, 2nd edition edition, 1987.
- [18] Meers B. J. Recycling in laser interferometric gravitational wave detectors. *Phys. Rev. D*, 38:2317–2326, 1988.
- [19] D. I. Jones. Gravitational waves from rotating strained neutron stars. *Class. Quantum Grav.*, 19:1255–1265, 2002.
- [20] V. Kalogera. Compact binary mergers and accretion-induced collapse: Event rates. In Sydney Meshkov, editor, *Gravitational Waves, Proceedings of the Third Edoardo Amaldi Conference*, volume 523 of *AIP Conference Proceedings*, pages 41–50. American Institute of Physics, 2000. astro-ph/9911532.
- [21] Brian T. Lantz. *Quantum Limited Optical Phase Detection in a High Power Suspended Interferometer*. PhD dissertation, Massachusetts Institute of Technology, Department of Physics, February 1999. LIGO-P990003-00-R.
- [22] Thomas J. Lardner, editor. *An Introduction to the Mechanics of Solids*. McGraw-Hill, New York, 2nd edition, 1972.
- [23] David R. Lide, editor. *CRC Handbook of Chemistry and Physics*. CRC Press, Boca Raton, FL, 3rd electronic edition, 2000. Available online at <http://www.knovel.com>.
- [24] Rodney Loudon. *The Quantum Theory of Light*. Oxford University Press, New York, 2nd edition, 1983.
- [25] Phil Marfuta. *Testing Dynamic Thermal Compensation of Optics for Use in LIGO II*. Bachelor’s thesis, Massachusetts Institute of Technology, Department of Physics, June 2001. LIGO-P010011-00-R.

- [26] The Mathworks, Inc., Natick, MA. *Using MATLAB*, version 5 edition, December 1996.
- [27] Charles W. Misner, Kip S. Thorne, and John A. Wheeler. *Gravitation*. W. H. Freeman and Company, New York, 1973.
- [28] G. Mueller, Q. Shu, R. Adhikari, D. Tanner, D. Reitze, D. Sigg, N. Mavalvala, and J. Camp. Determination and optimization of mode matching into optical cavities by heterodyne detection. *Opt. Lett.*, 25(4):266–268, 2000. LIGO-P000001-00-R.
- [29] Ophir Optronics Ltd., Jerusalem, Israel. *Orion Laser Power/Energy Monitor User Manual*, version 1.3 edition, November 1999.
- [30] David Ottaway. Optics and laser research at MIT. Technical report, LIGO Laboratory, June 2002. LIGO-G020270-00-R.
- [31] B. J. Owen and L. Lindblom. Gravitational radiation from the r-mode instability. *Class. Quantum Grav.*, 19:1247–1253, 2002.
- [32] Douglas A. Pinnow. Elastooptical materials. In Robert J. Pressley, editor, *CRC Handbook of Lasers with Selected Data on Optical Technology*, pages 478–488. CRC Press, Cleveland, Ohio, 1971.
- [33] S. F. Portegies-Zwart and S. L. W. McMillan. Black hole mergers in the universe. *Astrophys. J. Lett.*, 528:L17–L20, 2000.
- [34] William Press, Saul Teukolsky, William Vetterling, and Brian Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, 2nd edition edition, 1992.
- [35] Raymond Beausoleil, personal communication.
- [36] Frigyes Riesz and Bela Sz.-Nagy. *Functional Analysis*. Dover, Mineola, NY, 1990.

- [37] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, Singapore, 1994.
- [38] Bernard F. Schutz. *A First Course In General Relativity*. Cambridge University Press, Cambridge, UK, 1985.
- [39] Anthony E. Siegman. *Lasers*. University Science Books, Sausalito, California, 1986.
- [40] D. Sigg, N. Mavalvala, J. Giame, P. Fritschel, and D. Shoemaker. Signal extraction in a power-recycled michelson interferometer with fabry-perot arm cavities using a multiple-carrier frontal modulation scheme. *Appl. Optics*, 37:5687–5693, 1998. LIGO-P970016-01-R.
- [41] Walter A. Strauss. *Partial Differential Equations: An Introduction*. John Wiley & Sons, New York, 1992.
- [42] Kip Thorne. Science section of ligo-ii proposal - draft 4. Technical report, LIGO Laboratory, November 2000. LIGO-P000024-00-R.
- [43] Pin Tong and John N. Rossettos. *Finite-element Method: Basic Technique and Implementation*. MIT Press, Cambridge, Massachusetts, 1977.
- [44] Wavefront Sciences, Inc., Albuquerque, NM. *Complete Light Analysis System 2D*, 1.71.01 edition, August 2001.
- [45] Steven Weinberg. *Gravitation and Cosmology*. John Wiley & Sons, New York, 1972.
- [46] R. Weiss and B. Barish. LIGO and the detection of gravitational waves. *Physics Today*, 52:44–50, 1999. LIGO-P990039-00-R.
- [47] W. Winkler, K. Danzmann, A. Rüdiger, and R. Schilling. Heating by optical absorption and the performance of interferometric gravitational wave detectors. *Phys. Rev. A*, 44:7022+, 1991.