

TECHNIQUES FOR RESONANT OPTICAL INTERFEROMETRY WITH APPLICATIONS
TO THE ADVANCED LIGO GRAVITATIONAL WAVE DETECTORS

By

CHRISTOPHER L. MUELLER

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2014

© 2014 Christopher L. Mueller

This dissertation is dedicated to my parents
Walter Mueller and Janine Mueller
whose never-ending love and support put me on the path to success.

ACKNOWLEDGMENTS

The work in this dissertation was supported in countless ways by countless individuals. It is impossible to explicitly thank all of them, but I will make an attempt.

I owe my deepest appreciation to my advisor, Guido Mueller, without whom this dissertation would not have been possible. He is the reason that I stayed in graduate school and the reason that I ended up at the LIGO Livingston Observatory. He allowed me the freedom to explore my own interests, with intermittent pushes in the right direction to keep me on course.

I also owe a dept of gratitude to my surrogate advisor at LLO, Valery Frolov. He has shaped this dissertation in innumerable ways during discussions at our many ‘outside meetings.’ His honesty and candor are refreshing in a world where people are often too polite to tell the truth.

Many thanks go to the other people who played an advisory role in my studies. To Rana Adhikari; discussions with whom shaped many of the ideas in this dissertation, particularly in Part [IV](#). To Koji Arai and Paul Fulda; who helped me to understand the results of Chapter [8](#). To Muzammil Arain; who taught me the value of a quick and dirty experiment. To Peter Fritschel; who gave me the initial idea for the IMC absorption measurements. To David Tanner; whose wisdom in the field of optics is seemingly endless. And to Hiro Yamamoto; who taught me about optics and the power of modeling.

The input optics, discussed at length in Part [III](#), are the brainchild of my colleagues at the University of Florida: Muzammil Arain, Giacomo Ciani, David Feldbaum, Joe Gleason, Matt Heintze, Rodica Martin, Guido Mueller, David Reitze, David Tanner, and Luke Williams. The ease of installation, testing, and commissioning, and the high level of performance achieved is a testament to their brilliance.

The ideas of this dissertation would also not have come together without the support of my fellow young scientists, and I owe them all a sincere thank you. To Katie Leonard; her ambition and enthusiasm were a constant source of inspiration, and her

willingness to listen to me babble about my experiments still amazes me. To Anamaria Effler and Keiko Kokeyama; for driving me to the airport and feeding me delicious dishes from around the world. To Zach Korth, Adam Mullavey and Jan Poeld; whose friendship reminded me that there is more to life than work. To David Feldbaum and Eleanor King for help with the absorption and scattering measurements. To Ryan Derosa for teaching me about seismic isolation and bourbon. And to Denis Martynov for demonstrating what it is possible to accomplish through hard work.

I would never have made it out of the first years of graduate school without the help of my student colleagues at the University of Florida. I owe particular appreciation to Bobby Bond, Yitzi Calm, and Aaron Spector for countless hours of time spent on collaborative homework. A big thanks to the rest of my UF colleagues, of whom there are too many to enumerate, but whose friendship made the decision to move to Louisiana very difficult.

I also need to acknowledge all of the hard working scientists, managers, and staff of the LIGO Livingston Observatory. Their kindness and willingness to help made my stay there a pleasant one, and it is impossible to imagine how that place would run without them. The people I owe particular thanks to are: Carl Adams, Joseph Betzwieser, Suresh Doravari, Joe Giaime, Chris Guido, Ed Merilh, Richard Oram, Brian O'Reilly, Janeen Romie, Michael Thomas, and Keith Thorne.

I also should say thank you to Alexander Franzen whose optical components library (available from <http://www.gwoptics.org/ComponentLibrary/>) is used for all of the optical layout illustrations in this dissertation. Finally, I need to say thanks to the agency responsible for funding this work. The National Science Foundation supports both the LIGO project as a whole and my role in it. The field of gravitational wave detection would be nearly non-existent in the United States without their support.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGMENTS	4
LIST OF TABLES	10
LIST OF FIGURES	11
ABSTRACT	14
PART I INTRODUCTION	16
CHAPTER	
1 GRAVITATIONAL RADIATION AND DETECTION WITH LASER INTERFEROMETRY	17
1.1 Gravitational Radiation	18
1.1.1 The Weak Field Einstein Equations	19
1.1.2 An Observable Effect of Gravitational Radiation	20
1.2 Expected Sources of Gravitational Radiation	22
1.2.1 Inspiring Neutron Stars and Low Mass Black Hole Binaries	23
1.2.2 Spinning Neutron Stars	25
1.2.3 Supernovae and Other Collapse Phenomena	25
1.3 Topology of the Advanced LIGO Interferometers	26
1.4 Noise Sources in the aLIGO Detectors	31
1.4.1 Terrestrial Noise	32
1.4.2 Thermal Noise	34
1.4.3 Quantum Noise	36
PART II MATHEMATICAL TECHNIQUES FOR RESONANT OPTICAL INTERFEROMETRY	38
CHAPTER	
2 LASER RADIATION AND LASER BEAM PROPAGATION	39
2.1 Electromagnetic Radiation in Free Space	39
2.1.1 Propagation of the Fundamental Gaussian Mode	41
2.1.2 Comments on Propagation of Higher Order Modes	43
2.2 The ABCD Matrices of Geometric Optics	44
2.2.1 Geometric Optics and Ray Matrices	44
2.2.2 The ABCD Law for Gaussian Beams	45
2.3 Matrix Methods for Treating Misaligned Optical Systems	50
2.3.1 Derivation of the 4x4 Matrix Formalism	50
2.3.2 Definition of the Optical Axis	52

3	RESONANT OPTICAL INTERFEROMETERS	54
3.1	Quasi-static Interactions Between Laser Radiation and Resonant Optical Interferometers	54
3.1.1	Plane Wave Calculation of Resonant Optical Cavities	54
3.1.2	Finesse, Buildup, and Storage Time of a Two Mirror Optical Cavity	57
3.1.3	Spatial Modes, Stability, and Gouy Phase in a Two Mirror Optical Cavity	61
3.1.4	Spatial Modes, Stability, and Gouy Phase in Complex Resonant Optical Interferometers	64
3.1.5	Alignment Modes of Resonant Optical Cavities	66
3.2	Dynamics of Two Mirror Optical Cavities	68
3.2.1	Laser Frequency Noise and a Two Mirror Cavity	69
3.2.2	Length Noise in a Two Mirror Cavity	72
3.3	Symmetry in Optical Systems	74
3.3.1	Energy Conservation in Lossless Systems with no Optical Activity	74
3.3.2	Optical Symmetry Breaking with Optical Diodes	77
3.4	Analytical Models of Thermal Deformations	79
3.4.1	The Heat Equation	80
3.4.2	The Deformation Caused by a Point Absorber on the Mirror Surface	81
3.4.3	Realistic Heat Sources and Time Dependence with the Hello-Vinet Bessel Series Solution	84

PART III THE ADVANCED LIGO INPUT OPTICS 88

CHAPTER

4	THE ADVANCED LIGO ELECTRO-OPTIC MODULATOR	89
4.1	Requirements	90
4.2	Design	92
4.3	Measured Phase Modulation Indices	96
4.4	Residual Amplitude Modulation	98
4.4.1	Causes and Effects of RAM in Advanced LIGO	99
4.4.2	Measurements of RAM at the LIGO Livingston Observatory	100
4.5	Conclusion	103
5	THE ADVANCED LIGO FARADAY ISOLATOR	105
5.1	Requirements	105
5.2	Design	106
5.3	Measured Optical Losses	110
5.4	Measured Thermal Lensing	112
5.5	Measured Isolation Ratio	115

5.6	Conclusion	118
6	THE ADVANCED LIGO INPUT MODE CLEANER	119
6.1	Requirements	119
6.1.1	Frequency Noise at the Interferometer Input	120
6.1.2	Beam Alignment Noise at the Interferometer Input	122
6.2	Design	123
6.3	Control Loops	128
6.3.1	Length Control	128
6.3.2	Angular Control	130
6.4	Losses and Thermal Effects	132
6.4.1	Cavity Pole	133
6.4.2	Scattering	135
6.4.3	Absorption	138
6.5	Length and Frequency Noise	144
6.5.1	Noise Budget	145
6.5.2	Noise Deficit	148
6.6	Conclusion	151

PART IV EIGENSPECTRA MEASUREMENTS IN RESONANT OPTICAL INTERFEROMETERS 152

CHAPTER

7	AN INTRODUCTION TO EIGENSPECTRA MEASUREMENTS IN RESONANT OPTICAL INTERFEROMETERS	153
7.1	Can One Hear the Shape of a Drum?	153
7.2	Creating and Detecting Higher Order Modes	157
7.3	Fitting the Measurements	160
8	EIGENSPECTRA MEASUREMENTS OF THE INPUT MODE CLEANER	164
8.1	Experimental Setup	164
8.1.1	Calculation of the Signal	166
8.1.2	Example Scan	169
8.2	Frequency/Length Offsets and Measurements of the Length	169
8.3	Input Polarization Tuning	172
8.4	Absorption Measurements	172
9	EIGENSPECTRA MEASUREMENTS OF THE POWER RECYCLED MICHELSON INTERFEROMETER	178
9.1	Experimental Setup	178
9.1.1	The Phase-Locked Loop	181
9.1.2	Feedback Electronics	182
9.1.3	Calculation of the Signal	184
9.2	Non-Ideal Clipping Effects	186

9.3	PRC Cavity Parameters	187
9.4	A Thermal Error Signal	190
9.4.1	TCS Sensors and Actuators	190
9.4.2	The Error Signal	191
9.4.3	Continuous Tracking of the Thermal State	193
9.4.4	Conclusion	195
10	Conclusion and Outlook	198
10.1	The Input Optics	198
10.2	Eigenspectra Measurements	200
10.2.1	Arm Cavity Absorption	201
10.2.2	Probing the Full Interferometer in the Gravitational Wave Band	202
10.2.3	Thermal State Sensing	203
10.2.4	Parametric Instability Cooling	203
	APPENDIX: RAY MATRICES	205
	REFERENCES	207
	BIOGRAPHICAL SKETCH	216

LIST OF TABLES

<u>Table</u>	<u>page</u>
4-1 The key properties of a number of different materials considered for the aLIGO phase modulators.	93
4-2 The phase and amplitude modulation indices measured on the aLIGO EOM at the Livingston Observatory.	101
5-1 The requirements of the aLIGO Faraday isolator.	106
5-2 The possible optimizations of the LLO aLIGO FI at 20 W.	111
5-3 The measured focal length of the LLO FI with the three different thicknesses of DKDP that were tried.	114
6-1 The measured sensing matrix of the angular control loops of the input mode cleaner.	132
6-2 The scattering results for the MC1 and MC2 mirrors.	137
6-3 The data and inferred absorption from numerous repetitions of the Gouy phase absorption measurements in the IMC.	143
8-1 Two measurements of the IMC length with the eigenspectra probe.	171
8-2 The results of tracking the eigenfrequencies of the IMC HOMs up to order 4 as the power is stepped up and down.	175
8-3 The motion of the beam spot location on the three IMC mirrors during the measurements summarized in Table 8-2.	177
9-1 The cavity parameters measured with the PRC probe.	189
A-1 A table of 4x4 ray matrices.	206

LIST OF FIGURES

<u>Figure</u>	<u>page</u>
1-1 An numerically calculated inspiral waveform showing the inspiral, merger, and ringdown phases.	24
1-2 The optical configuration of the Advanced LIGO interferometers with the vacuum chambers and major optics labeled.	27
1-3 The theoretical noise budget for the Advanced LIGO interferometers.	32
2-1 A Gaussian beam propagates through a cascaded, paraxial optical system described in the language of geometric optics by an ABCD matrix.	46
2-2 A decentered ABCD system.	50
2-3 A system with a single center of curvature is shown as well as three different possible definitions of the misalignment terms.	53
3-1 A two mirror Fabry-Perot cavity.	55
3-2 The effects of differing finesse on the transmission and reflection coefficients of a two mirror Fabry-Perot resonator.	59
3-3 The reflection and transmission coefficients of a two mirror Fabry-Perot cavity together with their low frequency approximations.	71
3-4 Three different examples of lossless optical systems; a beamsplitter, a mirror, and a black-box system which could contain e.g. an optical cavity.	75
3-5 Two common types of optical diodes.	78
3-6 The green function solution to the heat equation together with the result of the Hello-Vinet calculation with a delta function absorber.	83
3-7 The thermal cross sections of an optic of infinite extent heated by a small square.	84
3-8 The temperature profiles of two of the aLIGO optics calculated from the Hello-Vinet calculation.	86
3-9 The development of the thermal bump on two different optics in the Advanced LIGO interferometers with only surface heating assumed.	87
4-1 The amplitude and phase noise requirements of the aLIGO electro-optic modulator.	91
4-2 A schematic of the EOM crystal and one of the resonant circuits.	94
4-3 The response of the EOM resonant circuit.	95
4-4 The measured phase modulation index of the LLO EOM as a function of frequency.	98
4-5 The spatial variation of the amplitude modulation index of the aLIGO EOM.	102
4-6 A long term study of the RAM directly after the EOM.	103
5-1 A schematic of the aLIGO Faraday isolator.	109

5-2	The measured magnetic field strength of the LLO and LHO Faraday rotators.	110
5-3	The theoretical self induced thermal lens in one of the TGG crystals in the FI at the full 165 W of input power.	113
5-4	The measured focal powers for the FI without the DKDP and the different thicknesses of DKDP.	115
5-5	A set of example beamscreens together with the fits to the data.	116
5-6	The measured isolation and transmission rejection ratios as a function of power.	118
6-1	A schematic layout of the input optics.	120
6-2	The Advanced LIGO design sensitivity is shown together with the measured sensitivity from the LIGO Hanford Observatory during the sixth LIGO science run.	121
6-3	The frequency noise requirements at the input to the interferometer both with and without the common arm loop suppression.	122
6-4	The mode spacing of the input mode cleaner.	124
6-5	A front and side view of the LIGO small triple suspensions.	127
6-6	A model of the open loop transfer functions of the various actuators in the IMC length/frequency control servo.	130
6-7	The measured IMC cavity pole is shown together with a simple single pole fit.	134
6-8	The shape of the first eight resonances of the IMC mirrors together with the eigenfrequency.	139
6-9	The frequency shift of each of the three input mode cleaner optics together with the input power.	140
6-10	The noise budget of the Advanced LIGO input mode cleaner.	146
6-11	The effect of driving white noise into both channels of the piezo-actuated mirror on the measured IMC length noise.	149
7-1	The losses from the various higher order modes at a mirror with radius equal to 3ω	155
7-2	The separation between the lowest order zeros of the Hermite polynomial term in a Hermite-Gauss beam.	156
7-3	The scattering matrices for pie shaped clipping with a wedge angle of 0.2π radians.	158
7-4	The power of the RF beat note is shown for different levels of cylindrical clipping with a carrier and sideband of unit power.	160
7-5	The imaginary part (I phase) and real part (Q phase) of the PDH signal used to fit the data from the eigenspectra measurements.	161
8-1	The experimental setup of the IMC eigenspectra measurement.	165

8-2	An example HOM sweep of the IMC across a full FSR together with a Finesse model.	170
8-3	The effect of tuning the control loop offset on the eigenspectra measurement of the free spectral range.	171
8-4	The p polarization peak of the IMC eigenspectra measurement and the effect of adjusting the input polarization.	173
8-5	The height of various deformations on the surface of MC2.	176
9-1	The layout of the PRC probe experiment.	179
9-2	The feedback controller used for the phase-locked loop of the PRC probe.	183
9-3	The measured signals at the zeroth and first order peaks for a number of different types of clipping.	187
9-4	The evolution of the I phase component of the first and second order modes as central heating is applied to ITMY.	193
9-5	The frequency and loss of the lowest order modes as 400 mW of CO ₂ power is applied to ITMY.	194
9-6	The results of tracking the thermal state of the PRC by locking the auxiliary laser to the resonance of the first order mode.	196
10-1	Comparison between the current noise of the aLIGO interferometer at LLO and the best eLIGO sensitivity.	199

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

TECHNIQUES FOR RESONANT OPTICAL INTERFEROMETRY WITH APPLICATIONS
TO THE ADVANCED LIGO GRAVITATIONAL WAVE DETECTORS

By

Christopher L. Mueller

August 2014

Chair: Guido Mueller

Major: Physics

A worldwide effort to detect gravitational radiation directly with large scale laser interferometers has been underway for the past several decades. In the United States the Laser Interferometer Gravitational-Wave Observatories (LIGO) have been operating since the early 2000s. These detectors use multiply resonant optical interferometers operating at a wavelength of 1064 nm to read out the strain induced by passing gravitational waves. This dissertation discusses a number of techniques and applications related to the operation of these interferometers at there designed sensitivity.

Part I introduces the reader to the concept of gravitational wave detection. We start by showing how the Einstein equations lead to the concept of propagating gravitational radiation and discuss some astrophysical events which are expected to emit enough radiation to be seen here on Earth. The reader is then introduced to the topology of modern terrestrial gravitational wave detectors and the limiting noise sources for such devices.

In Part II mathematical techniques for understanding and modeling these complex instruments are discussed; with a focus on the interaction of monochromatic radiation with resonant optical interferometers. Part II also contains a discussion of the effects of heating in high power resonant optical interferometers such as the Advanced LIGO (aLIGO) gravitational wave detectors.

Part III focuses on the design and testing of a particular subsystem of the aLIGO interferometers; the input optics. In particular it begins by examining the Electro-optic modulator whose task it is to add radio frequency sidebands to the main laser beam for sensing and detection of the various degrees of freedom of the interferometer. It discusses the in-vacuum, high power Faraday isolator which is responsible for deflecting the reflected light from the interferometer for sensing and control and for isolating the reflected light from returning to the rest of the input chain. Finally, it describes the input mode cleaner which is used to stabilize and prepare the laser beam before injection into the main interferometer.

Finally, Part IV of this dissertation describes a technique for characterizing these resonant optical interferometers through measurement of their eigenspectra. This technique and some of the nuanced details are explained before moving on to a description of the ways in which it has been used to characterize the Advanced LIGO interferometers. The key results obtained with this technique are described in which the eigenspectra measurement is used to monitor the thermal state of the input mode cleaner and of the power recycling cavity.

Part I
INTRODUCTION

CHAPTER 1 GRAVITATIONAL RADIATION AND DETECTION WITH LASER INTERFEROMETRY

A worldwide effort to directly detect gravitational radiation with large scale laser interferometers has been underway for the past several decades. In the United States the Laser Interferometer Gravitational-Wave Observatories (LIGO) have been operating since the early 2000s. During this time of operation a significant amount of effort was invested by the LIGO Scientific Collaboration (LSC) to research, design, and build upgrades to the initial LIGO interferometers. As of 2011 the initial LIGO detectors were decommissioned and installation of these upgrades began. As installation has progressed, the integration and commissioning phase has begun for many of the upgraded subsystems at the LIGO observatories. As of this writing the Livingston interferometer is almost fully installed and the full Advanced LIGO detector configuration has been brought to and held at its operating point for time periods up to a few hours.

This dissertation is broken into 4 distinct parts. Part I (this chapter) is a general introduction to gravitational wave detection with laser interferometers and will look at the topic from a bird's eye view. In particular, it will discuss how gravitational radiation falls out of the Einstein equations of general relativity, how it propagates, and what kind of sources can generate it. From there it will discuss the topology of the Advanced LIGO gravitational wave detectors and describe the noise sources which set the fundamental limits to their sensitivity.

Part II of this dissertation (Chapters 2 - 3) will discuss techniques for mathematically understanding resonant optical interferometers.

Part III (Chapters 4 - 6) discusses the Advanced LIGO input optics (aLIGO IO) which is one of the first major subsystems which was installed and commissioned at the two observatories. In particular, it will describe the three major components of the IO subsystem; the electro-optic modulator, the Faraday isolator, and the input mode cleaner. In each of these three cases I will describe how the requirements for

each component are derived from the interferometer sensitivity, we will show how these systems were designed to meet these requirements, and finally, we will examine measurements to see how well they are performing.

The final part of this dissertation, Part IV (Chapters 7 - 9), will discuss a very useful technique for characterization of resonant optical interferometers. This technique employs an auxiliary beam, which could take the form of phase sidebands, amplitude sidebands, or a true auxiliary beam, to measure the eigenspectra of a resonant optical interferometer. We will see that a lot of information can be extracted from measurements of the eigenspectra of such devices.

1.1 Gravitational Radiation

In the times after Sir Isaac Newton, before Albert Einstein, space and time were taken for granted as the static background in which one works on the problems of physics. A modern observer with knowledge of the vastness of space is lead by this view to wonder what reference frame is actually at rest. Some other strange characteristics of Newtonian gravity are that: bodies with mass know about each other instantaneously (action at a distance), massless objects like photons are undeflected by massive objects, and that space and time are disparate entities.

In addition to these problems, the Newtonian understanding of wave propagation led scientists naturally to believe that light needed a medium for propagation. In 1885 Michelson and Morley[57] built one of the most complex optical experiments of their time and showed that whatever the medium of propagation was, we could not see any motion relative to it.

Einstein's special and general theories of relativity solved all of these problems in one fell swoop. The fundamental principle behind relativity is that the laws of physics should be the same for any observer in any reference frame, i.e. two people moving or accelerating with respect to each other should still be able to agree on the laws of physics. Since then test after test of relativity has confirmed the validity of these

theories. A few of the most notable tests are the predication of the precession of the perihelion of Mercury by Einstein himself, the bending of light by stars[26], and the loss of energy in an orbiting system by gravitational radiation[38].

However; no one has yet observed directly the quivering of space itself due to the presence of moving mass distributions. *This* is the challenge of gravitational wave detection; to build an instrument which is capable of directly observing the small deformations in spacetime from the presence of massive moving bodies.

1.1.1 The Weak Field Einstein Equations

In this section we will see how the theory of general relativity leads to the idea of propagating gravitational radiation. We will also see why an interferometer is uniquely suited to detecting these perturbations in spacetime.

Recall the deceptively simple form of Einstein's equations with no cosmological constant,

$$G^{\alpha\beta} = 8\pi T^{\alpha\beta}. \quad (1-1)$$

The Einstein tensor $G^{\alpha\beta}$ is given in terms of the Ricci curvature tensor and the Ricci scalar by

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}. \quad (1-2)$$

In the weak field limit we may work in a nearly Lorentz coordinate system in which the metric has the form

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta}, \quad (1-3)$$

where $\eta_{\alpha\beta}$ is the Minkowski metric and the perturbations satisfy $|h_{\alpha\beta}| \ll 1$. This coordinate system has two very convenient properties which often simplify calculations. First, when transforming coordinate systems with the simple Lorentz transformations of special relativity, it can be shown that the perturbation tensor transforms as if it were a special relativity tensor unto itself (rather than a piece of the metric). Second, under

‘small’ gauge transformations it can be shown that $h_{\alpha\beta}$ is still small and that we therefore remain in an acceptable coordinate system even under gauge transformations.

The Einstein tensor expressed in this weak field limit is given by[80]

$$G_{\alpha\beta} = -\frac{1}{2} [\bar{h}_{\alpha\beta,\mu}{}^{,\mu} + \eta_{\alpha\beta} \bar{h}_{\mu\nu}{}^{,\mu\nu} - \bar{h}_{\alpha\mu,\beta}{}^{,\mu} - \bar{h}_{\beta\mu,\alpha}{}^{,\mu} + O(h_{\alpha\beta}^2)], \quad (1-4)$$

where $\bar{h}^{\alpha\beta}$ is the trace reverse of $h^{\alpha\beta}$ given by

$$\bar{h}^{\alpha\beta} = h^{\alpha\beta} - \frac{1}{2} \eta^{\alpha\beta} \bar{h}. \quad (1-5)$$

If we use the gauge freedom discussed above to demand that $\bar{h}^{\mu\nu}{}_{,\nu} = 0$, then the Einstein tensor can be written in a familiar form

$$G^{\alpha\beta} = -\frac{1}{2} \square \bar{h}^{\alpha\beta}. \quad (1-6)$$

Here the *D'Alembertian* or wave operator is defined as

$$\square f = f^{,\mu}{}_{,\mu} = \eta^{\mu\nu} f_{,\mu\nu} = \left(-\frac{\partial^2}{\partial t^2} + \nabla^2 \right) f. \quad (1-7)$$

Finally, we arrive at the weak field Einstein equations

$$\square \bar{h}^{\mu\nu} = -16\pi T^{\mu\nu}, \quad (1-8)$$

which have the form of a wave equation sourced by the stress energy tensor. These are also known as the linearized Einstein equations since they result from keeping only terms linear in $h_{\alpha\beta}$.

1.1.2 An Observable Effect of Gravitational Radiation

The linearized Einstein equations in vacuum ($T^{\mu\nu} = 0$) are given by

$$\left(-\frac{\partial^2}{\partial t^2} + \nabla^2 \right) \bar{h}^{\alpha\beta} = 0. \quad (1-9)$$

This equation clearly admits the standard wave solutions

$$\bar{h}^{\alpha\beta} = A^{\alpha\beta} e^{ik_\alpha x^\alpha}, \quad (1-10)$$

where $\{k_\alpha\}$ are the constant components of some one-form and $\{A^{\alpha\beta}\}$ are the constant components of some tensor.

We imposed one gauge condition while deriving the linearized Einstein equations whose effects we must consider. Applying $\bar{h}^{\mu\nu}{}_{,\nu} = 0$ to (1-10) gives

$$A^{\alpha\beta}k_\beta = 0. \quad (1-11)$$

I.E. our initial restriction of gauge was to a plane wave coordinate system in which $A^{\alpha\beta}$ is orthogonal to \vec{k} . Let us use our remaining gauge freedom to restrict ourselves to the transverse-traceless gauge which is defined by

$$A^\alpha{}_\alpha = 0 \quad \& \quad A_{\alpha\beta}U^\beta = 0, \quad (1-12)$$

where \vec{U} is some fixed four-velocity of our choosing. This uses up all of our gauge freedom which means that any remaining independent components must have physical significance. Notice that the traceless condition $A^\alpha{}_\alpha = 0$ implies that the trace reversed perturbation tensor is equal to the original, i.e.

$$\bar{h}^{\alpha\beta} = h^{\alpha\beta}, \quad (1-13)$$

where TT denotes the fact that we are now working in the transverse traceless gauge. If we choose our four-velocity in our gauge condition to be the time basis vector $U^\beta = \delta^\beta_0$, then the time amplitudes must be zero $A_{\alpha 0} = 0$. If we further choose to orient ourselves so that the wave is traveling in the z direction, $\vec{k} = (\omega, 0, 0, \omega)$, then the amplitude tensor reduces to

$$A^{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1-14)$$

We would now like to ask the question ‘what happens to the separation between two free particles in the presence of one of these waves?’ If we choose a Lorentz frame in which the particles are at rest and enforce our TT gauge conditions with respect to this frame, then it is fairly straightforward to show from the geodesic equation that the particles will remain at rest for all time. A naive interpretation might therefore lead to the conclusion that the separation between the particles remains the same for all time. The position of the particles is a coordinate dependent number however, but the separation between the particles is given by the coordinate invariant proper distance;

$$\Delta l = \int |ds^2|^{\frac{1}{2}}. \quad (1-15)$$

If we let one particle be at the origin and the other at $x = L, y = z = 0$, then this becomes

$$\Delta l_x \approx \left[1 + \frac{1}{2} h_{xx} \right] = [1 + A_{xx} e^{i\omega t}] L. \quad (1-16)$$

Similarly, if we consider the separation between a particle at the origin and a particle at $x = 0, y = L, z = 0$, then we find

$$\Delta l_y \approx \left[1 + \frac{1}{2} h_{yy} \right] = [1 - A_{xx} e^{i\omega t}] L. \quad (1-17)$$

If we were to rotate the detector by 45° (i.e. $x = \frac{L}{\sqrt{2}}, y = \pm \frac{L}{\sqrt{2}}$, and $z = 0$), then we would find that the detector in this orientation is sensitive to the other polarization of the gravitational wave A_{xy} .

This is why Michelson interferometers are ideally suited to measure gravitational waves. A passing gravitational wave squeezes one dimension while simultaneously expanding the orthogonal dimension. Michelson interferometers are extraordinarily sensitive devices for making precisely this measurement.

1.2 Expected Sources of Gravitational Radiation

Beginning with the initial LIGO proposal in 1989, LIGO was envisioned to come on line in two major steps. Initial LIGO was the first of those steps and included the building

of facilities, the operation of a reasonably complex large scale interferometer with strain sensitivities on the order of 10^{-21} , and the time to research new technologies for the second step, Advanced LIGO. It was hoped that these initial interferometers would provide the first detection of gravitational waves with reasonably high SNR, but alas this goal was unattained.

Advanced LIGO is the step which will transform the LIGO interferometers into gravitational wave *observatories*. The advanced interferometers will lower the strain noise by a factor of ~ 15 in the region of highest sensitivity, will widen the band of high sensitivity down to ~ 20 Hz and up to ~ 1 kHz, and will be capable of reshaping the noise in the sensitive region to target specific types of sources.

Lowering the noise at optimal frequencies will increase the event rates for distant sources by a factor of $\sim 15^3 = 3000$. The spreading of the sensitive frequency band will bring new sources into the LIGO window such as colliding black holes at low frequencies and fast pulsars at high frequencies. Noise curve reshaping will allow this generation to target specific periodic sources such as low mass X-ray binaries.[92]

Let us look more closely at some specific sources which Advanced LIGO will be able to see. Most of the information in this section will come from Kip Thorne's 2001 science case proposal for the Advanced LIGO upgrades[92]. It should be noted that estimates of rates rely on making population estimates within our own galaxy and then extrapolating to the nearby universe based on the density of massive stars.

1.2.1 Inspiring Neutron Stars and Low Mass Black Hole Binaries

Inspirals are considered the standard candle for gravitational wave detectors. The gravitational radiation generated by the inspiraling of massive systems can be broken up into three distinct phases: the inspiral, the merger, and the ringdown. The three distinct phases are visible in the cartoon of a merger waveform in Figure 1-1.

Using the technique of matched filtering it should be possible to finely measure these type of signals in Advanced LIGO data. Using these measurements we will be

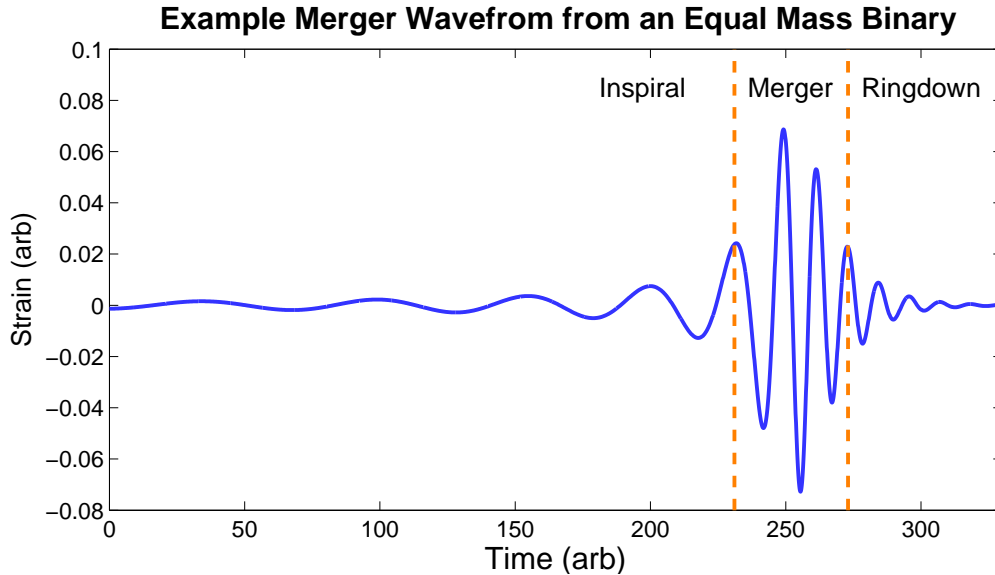


Figure 1-1: An numerically calculated inspiral waveform showing the inspiral, merger, and ringdown phases. Reproduced with permission from Scheel et al. “High-accuracy waveforms for binary black hole inspiral, merger, and ringdown.” *Phys. Rev. D* 79 (2009).

able to measure well the chirp mass from which it is possible to infer (with reduced accuracy and dependent upon models) the reduced mass and total mass of the system. The luminosity distance to these binaries can also be inferred to a precision of $\sim 10\%$. Some interesting general relativistic effects which should be visible in the inspiral waveforms include backscattering off of the spacetime curvature and frame dragging.

In addition to the standard inspiral, non standard events will hopefully allow Advanced LIGO data analysts to garner other interesting pieces of information. In some NS/BH cases the tidal forces will tear the NS apart before the merger. This situation gives astrophysicists a chance to pin down the nuclear equation of state at ten times the density of an atomic nucleus. Also, the spacetime around a BH/BH merger becomes so warped that it will allow Advanced LIGO to see the as yet unstudied nonlinear dynamics of spacetime.

The predicted rates for these mergers to be visible in the Advanced LIGO data are: 1/yr - 800/yr for NS/NS inspirals, 1/yr - 1500/yr for NS/BH events, and 10/yr - 4000/yr for BH/BH mergers. Even at the pessimistic end of these event rates, within a few years of

operation Advanced LIGO will produce a catalog with tens to hundreds of events which will provide a valuable source of information for astronomers and cosmologists.

1.2.2 Spinning Neutron Stars

Another interesting class of signals which should be visible in the Advanced LIGO data come from slightly deformed, spinning neutron stars. The ellipticity of nearby pulsars, such as the Crab and Vela pulsars, will be able to be measured by long integration times down to $\epsilon \sim 2 \cdot 10^{-8} \left(\frac{1 \text{ kHz}}{f}\right)^2 \left(\frac{r}{10 \text{ kpc}}\right)$ [92]. Based on the current understanding ellipticities of $\epsilon \sim 10^{-6}$ are within the realm of possibilities.

An interesting sub-class of these deformed, spinning neutron stars is the low-mass X-ray binaries (LMXB's). These are neutron stars that are being torqued by accretion from a companion but whose spin frequencies are fixed. The standard explanation for this is that gravitational radiation is spinning them down at the same rate that accretion is spinning them up. If this is assumed to be true, then the interferometer can be tuned by signal recycling to specifically target a number of known sources.

The measurement of gravitational waves from these sources, particularly those which have already been identified by electromagnetic observations, will allow astronomers to identify specific details about the system. Some examples are: the nature of the inhomogeneities, the underlying physics of the crust and core, and the angular momentum vector relative to the line of sight.

1.2.3 Supernovae and Other Collapse Phenomena

Although the previously discussed sources will provide interesting information about reasonably well understood systems, the most interesting prospects of the Advanced LIGO era lie in the more poorly understood phenomena. One of the benefits and drawbacks of gravitational wave interferometers is that they act more like an antenna than a telescope. This means that a network of interferometers is required to understand the direction from which a signal came, but at the same time they are sensitive in all directions at the same time (with some antenna pattern). This hopefully means that in

the midst of studying the hundreds of inspirals and asymmetric neutron star events, Advanced LIGO will pick up some rare, ultra violent, poorly understood events at the same time.

A particular instance of these violent events is the Type-II supernova. These are believed to be caused by the violent collapse of a stellar core to form a neutron star or black hole. Optical observations are not very useful for understanding the underlying physics, but gravitational waves are. They carry information about the compactness and density of the core, the asymmetry, and about the final remnant of the collapse. There is also the possibility that exotic bar like objects form during the collapse which would emit gravitational waves in the LIGO band of interest.

1.3 Topology of the Advanced LIGO Interferometers

A schematic layout of the Advanced LIGO detectors is shown in Figure 1-2. In this section I will try to describe qualitatively the reasoning behind the complex design of these devices. I will try to show that the Michelson interferometer forms the heart of the device with a number of high-tech upgrades wrapped around it. We will see that every other piece of the interferometer serves the purpose of enhancing the sensitivity in a particular way.

At the heart of the Advanced LIGO detectors is a Michelson interferometer which, as mentioned in Section 1.1.2, is ideally suited to making measurements of gravitational waves. A perfect Michelson interferometer is highly sensitive to differential changes between its two arms while being completely insensitive to common changes. Since the input beam is common until it is split at the beam splitter, a Michelson interferometer is also insensitive to both the pointing noise and frequency noise of the input beam. The insensitivity to frequency noise is important for gravitational wave detection because

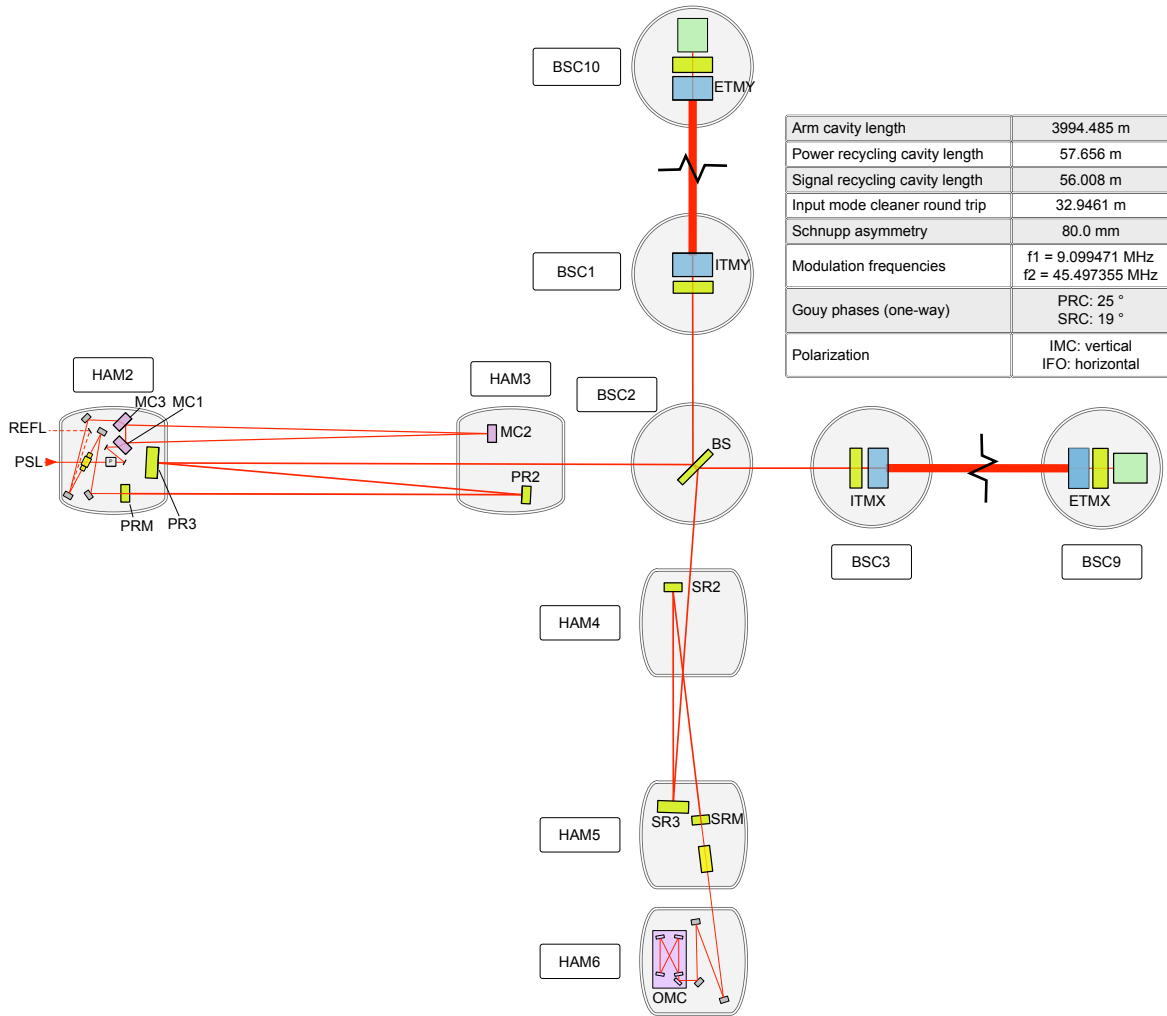


Figure 1-2: The optical configuration of the Advanced LIGO interferometers with the vacuum chambers and major optics labeled. In addition some of the key optical parameters are listed in the table.

even a perfectly stabilized laser has frequency noise due to spontaneous emission¹ of the gain medium within the linewidth of the atomic transition[78].

A typical Michelson interferometer can achieve sensitivity to differential length fluctuations on the order of $\delta l \sim 10^{-15}$ m. Recall that the gravitational wave signal is a

¹ Recall that a laser operates on stimulated emission which creates a photon of identical phase and frequency.

strain signal however so the quantity of interest is

$$h = \frac{\delta\ell}{\ell},$$

where $\delta\ell$ is the sensitivity to length fluctuations and ℓ is the length of the arms of the Michelson. In other words, to achieve astrophysically interesting sensitivities to gravitational waves which have strains on the order of $h \sim 10^{-22}$ we need to effectively extend the arm length of the Michelson to 10^7 meters or increase the sensitivity of the Michelson by 7 orders of magnitude. The aLIGO interferometers use a combination of the two strategies to achieve the sensitivities necessary for gravitational wave detection.

The Michelson interferometer acts as a phase comparator; i.e. it compares the phase of the optical fields returning from each arm and based on that comparison decides how much light to send to each of the two ports. This points to one way of increasing the sensitivity of the Michelson to length fluctuations of the two arms; by increasing the phase change of the reflected light from the two arms. This is accomplished in the Advanced LIGO interferometers by adding resonant Fabry-Perot cavities to each of the arms. For a given change in length of one of the arms the phase of the light reflected from the cavity, when near resonance, is enhanced by a factor of the arm cavity finesse, \mathcal{F} . The finesse of resonant optical cavities is discussed in more detail in Section 3.1.2. This phase sensitivity enhancement can be thought of, in terms of the gravitational wave signal, as extending the effective arm length by the same factor. The arm cavity finesse can not be increased indefinitely however since the bandwidth of the detector drops as the finesse is increased. This decrease in bandwidth can be understood qualitatively as; increasing finesse leads to photons spending longer in the arm cavities, and signals which oscillate faster than this time are averaged out since the photons pick up as many expansions as compressions before carrying the information out to the beamsplitter.

The fundamental limit to any optical experiment which does not employ squeezed quantum states is set by the detection shot noise. Shot noise can be thought of as arising in laser interferometers because the photon emission process of a laser exhibits the randomness typical of quantum systems, and in the end the detection of an optical experiment depends on measuring power which is the same as counting photons. The detection is therefore a counting process which exhibits the noise associated with a typical Poissonian process. Shot noise is discussed in more detail in the next section, but for now it is sufficient to simply accept the fact that, other things equal, the shot noise in a recycled interferometer with Fabry-Perot arms is inversely proportional to the square root of the power at the beam splitter[53];

$$S_{shot}(f) \sim \frac{1}{\sqrt{P_{BS}}}. \quad (1-18)$$

Since gravitational wave signals are differential signals, and the Michelson serves the role of separating the differential signals from the common signals, sending the light which only sees the common signals back into the interferometer has the advantage of increasing the operating power without affecting the differential bandwidth. Hence, modern gravitational wave interferometers typically employ the technique of power recycling which uses a mirror placed at the bright port of the Michelson such that the reflected power interferes constructively with the incident laser power.

The recycling technique can be viewed in a slightly different, but perhaps more illuminating way. Viewed from inside of the arm cavity, the recycling mirrors form a simple Fabry-Perot cavity with the input test masses (ITMs). These compound cavities can have different reflectivities (higher or lower) than the reflectivity of the ITMs alone. For a detuned (slightly off of perfect resonance) recycling mirror, the effective mirror formed with the ITMs can even have a frequency dependent reflectivity. Hence, the power recycling mirror is used to increase the effective reflectivity of the ITMs, which is equivalent to increasing the arm cavity finesse, but only for common signals. Differential

signals still only see the compound mirror formed by the signal recycling mirror and the ITMs. The bandwidth of the detector is still decreased for common signals to ~ 1 Hz, but this is advantageous since it allows very effective sensing of the frequency noise incident on the interferometer.

The signal recycling mirror (SRM), which might be better called the signal extraction mirror, serves the opposite purpose[56]. The compound reflectivity of the signal recycling mirror together with the ITMs is lower (higher transmissivity) than the ITMs alone. Tuning the reflectivity of this mirror can be used to tune the bandwidth of the interferometer to differential signals which include the gravitational wave signals. In fact, positioning the SRM such that the length of the cavity is held slightly off resonance from the carrier light can be used to give the ITM/SRM combo a frequency dependent reflectivity which can be used to tune the interferometer's response for particular sources, a configuration known as narrowband recycling.

Referring to Figure 1-2 shows that both the input and the output of the interferometer employ one more resonant optical cavity. The one at the input is called the input mode cleaner (IMC) and is the topic of a much more in depth discussion in Chapter 6, but for now I will give a brief overview of its purpose. The role of the input mode cleaner is to prepare the laser beam before injection into the main interferometer. It does this by filtering the spatial modes of the laser beam and only allowing the diffraction limited Gaussian mode to pass. It also serves as a first stage of active and passive laser frequency stabilization, in which the ultra-stable length of the IMC is used as a reference to actively quiet the laser's frequency. A few other passive purposes served by the input mode cleaner are; polarization filtering, passive and active beam pointing stabilization, and passive intensity stabilization, all of which are important for the interferometer to reach its design sensitivity.

Finally, in keeping with the symmetry of the rest of the detector, there is another, much smaller optical cavity at the output of the interferometer. This optical cavity is

used when operating the differential sensing by the method known as DC readout. In this scheme the differential length of the arms is held slightly off resonance (on the order of picometers) so that a small amount of the light leaks out of the Michelson to the dark port of the interferometer. Fluctuations in the differential arm length cause the light leaking to the dark port to fluctuate, and this forms the signal for reading out the gravitational wave channel. However, there is other light at the dark port which is not associated with the leaking carrier light from the arm cavities, and this light carries no signal but still adds to the noise of the measurement in the form of shot noise. The job of the output mode cleaner (OMC) is to strip all of this extra light from the light which is leaking from the arm cavities in order to increase the signal to noise ratio of this readout channel.

1.4 Noise Sources in the aLIGO Detectors

In experimental physics it is often easy to devise an experiment which will see what one would like to see. It can be much more difficult however to devise an experiment which will not see what one would like not to see. The class of things which one would like not to see in an experiment is called noise.

As was pointed out above, an interferometric detector is sensitive to the path length differences caused by passing gravitational waves. Let us discuss some of the other sources of test mass displacement which will show up in a terrestrial interferometric detector as well as briefly address some ways to mitigate these noise sources. Figure 1-3 shows the expected Advanced LIGO noise budget, we will use it as our guide for a discussion of the noises beginning at low frequencies on the left of the plot and working our way towards higher frequencies.

It should be noted that there are two ways to decrease the noise from any particular source in an experiment. One may reduce the source of the noise or one may reduce the coupling between the experiment and the noise source. In general it is best to use




Figure 1-3: The theoretical noise budget for the Advanced LIGO interferometers. Notice that, across most of the sensitivity band between 10 Hz and 10 kHz the interferometers will be nearly quantum noise limited. The exceptions to this occur where the thermal noise of the suspensions and optical coatings become comparable with the quantum noise between 10 Hz and 100 Hz.

both methods simultaneously, but in practice one often has no control over the source of the noise and is forced to fight the coupling.

1.4.1 Terrestrial Noise

At very low frequencies the Advanced LIGO detectors will be limited by seismic noise. Seismic noise is motion of the ground which the detector is bolted to, and is caused by internal motion of the Earth, waves crashing on the shore, and human activity among other things. Seismic waves couple most strongly to displacement when they hit the surface of the Earth and become what are known as surface waves. Terrestrial based detectors can therefore reduce seismic noise by going underground. The depth to which one must dig to avoid the surface seismic noise scales with the wavelength of the seismic waves (skin depth), and Rayleigh seismic waves at the microseism

have wavelengths on the order of ~ 50 km so going underground can be prohibitively expensive.

Advanced LIGO sits on the ground and was therefore forced to fight the coupling to seismic noise. It is a straightforward mechanics calculation to show that the transfer function of an undamped pendulum from its support point (X_0) to the suspended mass (X_1) is given by

$$\frac{X_1(\omega)}{X_0(\omega)} = \frac{\omega_0^2}{\omega_0^2 - \omega^2}, \quad (1-19)$$

where $\omega_0 = \sqrt{\frac{g}{\ell}}$. Looking at the limits of this function we see that at frequencies well below the resonance the suspended mass moves exactly as its support point, on resonance there is a large amplification of the motion, but well above the resonance the motion of the support point is suppressed as $1/\omega^2$. Hence, at ten times the resonance frequency the coupling to the support motion is 10^{-2} , and at 100 times the resonance frequency the coupling is 10^{-4} , etc.

The seismic isolation for Advanced LIGO is one of the most complex systems in the detectors, consisting of seven (five) stages of active and passive isolation for the test masses (other major optics). The isolation starts with an out-of-vacuum spring support system for passive isolation together with hydraulic actuators for damping and active isolation[2] (known as HEPI). This system is designed to give some passive isolation above 1 Hz and to be used for active feed-forward at very low frequencies, a technique which was proven during Enhanced LIGO to significantly improve detector uptime and decrease the rate of non-Gaussian transients[21]. Inside of the vacuum system sits an optical platform with two (one) stages of active and passive isolation for the test masses (other major optics)[2]. The passive isolation acts like a simple pendulum in all six degrees of freedom, suppressing noise as f^{-2} per stage above the platform resonances which lie between 1 Hz and 5 Hz. The active isolation of these tables is, however, the true workhorse of the Advanced LIGO seismic isolation system. Each stage of the optical platforms is equipped with low noise inertial sensors and

electromagnet actuators which enables them to suppress the ground motion by up to 3 orders of magnitude in the band from 1 Hz to 20 Hz. This suppresses the RMS motion of the two stage tables at 1 Hz down to a few picometers.

Mounted to these optical platforms are multi-stage pendula suspension which support the mirrors themselves. These multi-stage pendula suspensions use four (three) successive simple pendula to passively isolate the test masses (other major optics), suppressing the ground motion as $\sim f^{-8}$ ($\sim f^{-6}$) above the resonances which are around 1 Hz[74]. Each of these stages is also equipped with the ability to read out the position relative to the support structure and actuate electromagnetically on the position and pointing.

In addition to the seismic noise caused by motion of the supporting structure of the optic, seismic noise (and atmospheric) noise can couple directly to the test masses through gravity. This source of noise is known as gravity gradient noise or Newtonian noise and is one of the fundamental limiting noise sources at low frequencies for terrestrial gravitational wave detectors. This is a fundamental noise source because, since mass only comes in one sign, there is no way to shield from gravity.

There is one difference between gravity gradient noise and astrophysical gravitational waves however. Fluctuating mass distributions generate seismic noise which can be sensed by local inertial sensors. These sensors can be used with optimal feed-forward schemes either on or off line to subtract some of the gravity gradient signal from the gravitational wave signal[25]. As shown in Figure 1-3 gravity gradients are expected to be just below the Advanced LIGO sensitivity and therefore will not be important until the next generation of interferometric detectors is being designed.

1.4.2 Thermal Noise

All systems at finite temperature suffer from stochastic motion of their constituents on the atomic level. This stochastic motion shows up in the interferometer as an apparent displacement of the measured surface of the mirrors. Since the source of

thermal noise is finite temperature, one can fight the source by lowering the temperature of the mirrors and suspensions. Thermal noise only scales with the \sqrt{T} so cooling the mirrors is less advantageous than one might hope. The Advanced LIGO interferometers keep the thermal noise low by using low loss materials whenever possible.

Before discussing the details, let's discuss why low loss materials provide an effective way to reduce the thermal noise. In a seminal paper from the 1950s Callen and Welton[13] showed that it was possible to make a generalization of the Johnson noise in a resistor which was originally explained by Nyquist in the 1920s[66]. The result of this paper has become known as the fluctuation dissipation theorem, which gives a relationship between any lossy system at finite temperature and the fluctuations of a measurement of the system. Speaking specifically of thermal noise in interferometry; we can imagine that if we were to push on the surface of the interferometer mirrors at a given frequency, that the losses in the system would eventually convert the energy which we put in by our push into thermal energy of the mirror or its supporting structure. The fluctuation dissipation theorem merely points out that this process is symmetric; since the mirror is at finite temperature the thermal motion of its constituent atoms shows up as a displacement of the mirror surface. The quantity which relates the flow of energy *in both directions* is the dissipation (losses) in the mirror or its supporting structure.

Thermal noise in a suspended interferometer may be thought about as coming from the losses in four places: the suspension, the substrate of the mirrors, the highly reflective coating on the mirrors, and the interfaces between all three of these. Advanced LIGO, like initial LIGO, uses quartz (fused silica) for its mirror substrates which has a very low dissipation coefficient and therefore contributes negligibly to the overall interferometer noise (see e.g. substrate Brownian noise in Figure 1-3). One of the upgrades for Advanced LIGO was to make the penultimate stage and the wires connecting the penultimate stage and mirror out of quartz so that the final stage of the suspension is made entirely from quartz. This has two advantageous effects, it

reduces the thermal noise of the suspension itself and it reduces the thermal noise due to loss at connections between different materials. The highly reflective coatings on the mirror surfaces are now the largest sources of thermal noise, but materials with lower losses which satisfy the other technical requirements of the coatings are not yet mature enough for large scale applications. The search for better coating materials is an ongoing process within the gravitational wave community, but no new technologies are used in Advanced LIGO.

1.4.3 Quantum Noise

The ultimate noise floor in any measurement system is set by the quantum mechanics of the readout mechanism. In interferometers such as the Advanced LIGO detectors, this readout mechanism is light which is made up of a large number of photons. The final optical readout involves measuring the power of this light which is equivalent to counting the number of photons in the beam. This counting process exhibits the typical Poissonian statistical properties in which during a given period of time the average number of photons will be well defined, but the actual number of photons will exhibit fluctuations around this average. This source of noise is known as shot noise.

Shot noise is itself white, which is to say that it has the same power spectral density at all frequencies. Referring to Figure 1-3 shows that the quantum noise is not flat for the Advanced LIGO interferometer. At high frequencies this is caused by the fact that the detector is less sensitive to strain due to the long, high finesse arm cavities so the contribution of the shot noise to the measured strain is stronger. At low frequencies, however, the interferometer is exhibiting another classic phenomenon of quantum measurement known as quantum back action. It shows up in this case as a fluctuating radiation pressure force on the surface of the mirrors in the arm cavities. This fluctuating force on the mirror causes them to move in a statistically random way, and the slope of the curve at low frequencies is due to the inertia of the free mass. Reducing this

source of noise is possible with typical gravitational wave interferometers by increasing the mass of the mirrors, but it is possible to design a different interferometer which is completely insensitive to radiation pressure noise[18]. In more general language, measurement devices which avoid this radiation pressure effect fall into the category of quantum non-demolition measurement devices.

Although the statistical nature of photons provides an intuitive way of thinking about quantum noise in the interferometers, Caves pointed out a much more powerful way of thinking about this source of noise in 1980[16]. He shows that the quantum noise in an interferometer can be thought of equivalently as vacuum state fluctuations entering the interferometer from all of the loss ports while the laser driving the system can be treated completely classically. This mode of thinking led Caves to point out a way to beat the standard quantum limit in interferometry by injecting squeezed vacuum into the dark port of the interferometer[17]. In the early 2010s, shortly before decommissioning the Enhanced LIGO interferometers, the LIGO collaboration showed that this approach was achievable in a kilometer scale interferometer[90][91]. Injecting squeezed states of light into the dark port of the interferometer will almost certainly be an early upgrade for the Advanced LIGO detectors.

Part II

MATHEMATICAL TECHNIQUES FOR RESONANT OPTICAL INTERFEROMETRY

CHAPTER 2
LASER RADIATION AND LASER BEAM PROPAGATION

2.1 Electromagnetic Radiation in Free Space

Most textbooks[58][85][45] on lasers will have a derivation of the eigenmodes of a laser resonator, but we will recap it here for completeness. Starting from Maxwell's equations in vacuum, it is straightforward to show that all three components of both the electric and magnetic field satisfy the wave equation

$$\left(\nabla^2 - \mu\epsilon\frac{\partial^2}{\partial t^2}\right) u = 0.$$

For monochromatic waves we can factor out the time dependence $e^{-i\omega t}$ which transforms the wave equation into the Helmholtz equation.

$$(\nabla^2 + k^2) u(x, y, z) = 0, \tag{2-1}$$

where $k = \sqrt{\mu\epsilon}\omega = \frac{2\pi}{\lambda}$. If we choose the z axis as our optical axis and make the further assumption that the z dependence will be mostly that of a plane wave

$$u = \psi(x, y, z)e^{-ikz},$$

then the Helmholtz equation reduces to

$$\frac{\partial^2\psi}{\partial x^2} + \frac{\partial^2\psi}{\partial y^2} - 2ik\frac{\partial\psi}{\partial z} = 0, \tag{2-2}$$

where we have made the assumption that the leftover z dependence in $\psi(x, y, z)$ varies slowly enough that we may neglect the second partial derivative of ψ with respect to z . This is known as the paraxial wave equation. We have written it in rectangular coordinates, but the first two partial derivatives are simply the two dimensional Laplacian which can be written in other coordinate systems such as cylindrical.

There are many ways to solve such a differential equation, but the solution will, in general, depend on boundary conditions. The standard method of solution for laser

beams is to observe that the equation admits a simple Gaussian pulse,

$$\psi = \exp \left[-i \left(P + \frac{k}{2q} r^2 \right) \right], \quad (2-3)$$

where $r^2 = x^2 + y^2$. Note that both P and q are functions of z .

We will look more closely at the simple Gaussian solution in a moment, but for now keep in mind that this solution is only one possible solution to (2-2) which may not match the initial boundary conditions. To obtain more solutions to this differential equation we attempt a solution with the Gaussian portion (2-3) factored off

$$\psi(x, y, z) = H(x, y) \exp \left[-i \left(P + \frac{k}{2q} r^2 \right) \right].$$

Doing so and solving the resulting differential equation, we find that $H(x, y)$ is simply a product of the one dimensional Hermite polynomials so that our solution becomes

$$\psi_{nm} = H_n \left(\frac{\sqrt{2}x}{\omega} \right) H_m \left(\frac{\sqrt{2}y}{\omega} \right) \exp \left[-i \left(P + \frac{k}{2q} r^2 \right) \right], \quad (2-4)$$

where H_i represents the i th Hermite polynomial and $n, m \in \mathcal{N}$. The orthogonality of the Hermite polynomials implies that the vector space made up of these solutions is complete. Hence, these function may be thought of as the eigenfunctions of the paraxial wave equation in rectangular coordinates. They provide a complete, orthogonal set of solutions in which the solution to a problem with arbitrary boundary conditions may be expanded.

Had we instead chosen to solve (2-2) in cylindrical coordinates we would have found a similar set of eigenfunctions given by

$$\psi_{\ell p} = L_p^{|\ell|} \left(\frac{2r^2}{\omega^2} \right) \exp \left[-i \left(P + \frac{k}{2q} r^2 \right) \right], \quad (2-5)$$

where $L_p^{|\ell|}$ is the Laguerre polynomial of radial index $p \in \mathcal{N}$ and azimuthal index $\ell \in \mathcal{Z}$.

The eigenfunctions in rectangular coordinates are usually called the Hermite-Gauss

modes while those in cylindrical coordinates are referred to as the Laguerre-Gauss modes.

2.1.1 Propagation of the Fundamental Gaussian Mode

Most laser beams are intentionally prepared such that they are composed almost solely of light in the fundamental mode with only tiny amounts of power in the higher order modes. It is very often the case, therefore, that the most noticeable properties of a laser optical system can be understood by considering its interaction with only the fundamental mode.

Substituting (2-3) into (2-2) and comparing terms of equal powers in r leads to two simple differential equations for q and P ,

$$\frac{dq}{dz} = 1 \quad (2-6)$$

and

$$\frac{dP}{dz} = -\frac{i}{q}. \quad (2-7)$$

We will deal with the physical meaning of P in a moment, but for now let's consider the q parameter which will turn out to be a very useful for characterizing laser beams which are Gaussian or nearly Gaussian.

Integrating (2-6) yields the simple conclusion $q(z) = q_0 + z$. If we now express $\frac{1}{q}$ in terms of its real and imaginary parts

$$\frac{1}{q} = \frac{1}{R(z)} - i \frac{\lambda}{\pi \omega(z)^2}, \quad (2-8)$$

and substitute into (2-3) a clear physical interpretation emerges for R and ω . We see that $R(z)$ represents the radius of curvature of the wavefront, and $\omega(z)$ is a measure of the field amplitude decrease with distance from the optical axis (often called the beam radius). The Gaussian beam contracts to a minimum size at the location where the phase fronts become plane. This location is referred to as the beam waist and has the q

parameter

$$q_0 = i \frac{\pi \omega_0^2}{\lambda}.$$

If we choose to measure z from the beam waist, then the q parameter at any point in the system is given simply by

$$q(z) = i \frac{\pi \omega_0^2}{\lambda} + z. \quad (2-9)$$

Combining this with (2-8) and equating the real and imaginary parts, we find that the beam size and radius of curvature at any point in the system are given by

$$\omega^2(z) = \omega_0^2 \left(1 + \frac{z^2}{z_R^2} \right) \quad (2-10)$$

and

$$R(z) = z \left(1 + \frac{z_R^2}{z^2} \right) \quad (2-11)$$

where

$$z_R = \frac{\pi \omega_0^2}{\lambda}. \quad (2-12)$$

This parameter is known as the Rayleigh range and it acts as a rough division point for the properties of Gaussian beams. Points along the z dimension which are close to the waist with respect to the Rayleigh range are known as the near field and points much farther from the waist than the Rayleigh range are known as the far field. In the near field the beam size does not change much but the radius or curvature varies rapidly over a small region, becoming infinite at the waist. In the far field the expansion of the beam size is linear in the propagation distance and the radius of curvature of the field looks like a point source sitting at the waist.

Lets now turn to the P parameter. Substituting (2-9) into (2-7) and integrating yields

$$\begin{aligned} P(z) &= -i \ln \left(1 - i \frac{\lambda z}{\pi \omega_0^2} \right) \\ &= -i \ln \left(\frac{\omega_0}{\omega} \right) - \arctan \left(\frac{z}{z_R} \right) \end{aligned}$$

The real part of P represents a phase shift between a Gaussian beam and a standard plane wave, commonly known as the Gouy phase. The Gouy phase is an additional phase shift picked up by monochromatic radiation beyond that of a plane wave. The imaginary part represents an amplitude factor which gives the expected intensity decrease along the optical axis due to spreading of the beam. Putting all of this together, we can express our Gaussian solution (2-3) as

$$u(r, z) = \frac{\omega_0}{\omega(z)} \exp \left[-i \left(kz - \arctan \left(\frac{z}{z_R} \right) \right) - r^2 \left(\frac{1}{\omega(z)^2} + \frac{ik}{2R(z)} \right) \right], \quad (2-13)$$

where $\omega(z)$ is given by equation (2-10) and $R(z)$ is given by equation (2-11).

2.1.2 Comments on Propagation of Higher Order Modes

By substituting (2-4) into (2-2) and repeating the analysis of the previous section we find a few important results. First of all, we learn that the two beam parameters $\omega(z)$ and $R(z)$ are the same for the excited modes as for the fundamental mode. This means that, within the limit that $\partial^2 \psi / \partial z^2$ can be neglected in (2-1), the higher order modes do not expand any more quickly than the fundamental mode and that the phase front curvature is the same for all modes. The fact that the higher order modes expand at the same rate as the beam propagates down the optical axis means that the beam structure will not change (modes do not mix) as we track it except for an overall expansion or contraction.

Secondly, we find that the Gouy phase term is different for each mode. Namely,

$$\psi_{nm} = (n + m + 1) \arctan \left(\frac{z}{z_R} \right). \quad (2-14)$$

This term is the reason that optical resonators can trap and isolate the different modes independently. We will see in Chapter 6 that it is key to the functioning of the input mode cleaner. The accumulated phase shift through an optical system also provides a way to track the imaging properties of the beam from one plane to another[27].

2.2 The ABCD Matrices of Geometric Optics

The study of geometric optics is the study of rays of light. These rays are related to wave optics by defining a ray at each point on a wavefront to be perpendicular to the wavefront and pointing in the direction of energy flow. This ray picture ignores, however, the polarization of the optical wave, but, nevertheless, provides a useful picture for understanding the propagation of light through first order, paraxial, optical systems. We will see in a moment that this picture provides a convenient formalism for understanding how a Gaussian beam transforms through such an optical system as well.

2.2.1 Geometric Optics and Ray Matrices

Let us define an arbitrary optical ray by its position, $r(z)$, and slope, $r'(z)$ with respect to the optical axis. The slope of the ray is given by[58][85]

$$r'(z) = \frac{dr(z)}{dz} = \tan \theta \approx \theta, \quad (2-15)$$

where we have made the small angle approximation of $\tan \theta$. This approximation is known in the field of optics as the paraxial approximation and can be summarized in words as ‘rays remain close to the optical axis’. We can use these two quantities to define a vector, \vec{r} , which fully characterizes the ray;

$$\vec{r} = \begin{bmatrix} r \\ \theta \end{bmatrix}. \quad (2-16)$$

The advantage of working within the paraxial approximation is that expressing the transformation of a ray as it propagates through an optical system reduces to a linear problem. As an example, propagation through free space may be expressed as

$$\begin{bmatrix} r_f \\ \theta_f \end{bmatrix} = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix} \begin{bmatrix} r_i \\ \theta_i \end{bmatrix}. \quad (2-17)$$

in the limit that $\sin \theta_i \simeq \theta_i$. Similarly a ray passing through a thin lens can be expressed as

$$\begin{bmatrix} r_f \\ \theta_f \end{bmatrix} \approx \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix} \begin{bmatrix} r_i \\ \theta_i \end{bmatrix}. \quad (2-18)$$

in the paraxial limit.

Many optical systems can be expressed in the paraxial limit in the form of a simple ray matrix;

$$\begin{bmatrix} r_f \\ \theta_f \end{bmatrix} \approx \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} r_i \\ \theta_i \end{bmatrix}. \quad (2-19)$$

The ray matrices for many of these systems can be found in appendix [10.2.4](#). There we simply collect different ray matrices from the literature without derivation. In the next section we will see how these concepts from geometric optics can be applied to the Gaussian beams of Section [2.1](#). We will find surprisingly that the concepts of geometric optics, which studies optical rays, give a nearly complete description of the spatial structure of Gaussian beams which are governed instead by the process of diffraction.

2.2.2 The ABCD Law for Gaussian Beams

In order to show that the ray matrices of geometric optics give a near complete description of the propagation of Gaussian beams we will start from the Huygens-Fresnel diffraction integral which governs the propagation of monochromatic radiation. We will show how this integral through a cascaded, paraxial optical system can be expressed solely in terms of the ABCD matrix elements from geometric optics. From there we will derive a very simple algebraic expression for mathematically propagating a Hermite-Gaussian or Laguerre-Gaussian beam through a cascaded, paraxial optical system.

For simplicity of exposition we will restrict our discussion to one dimension and assume that the propagation begins and ends in vacuum with index of refraction identically equal to one. The reader is referred to the literature for similar derivations with




Figure 2-1: A Gaussian beam propagates through a cascaded, paraxial optical system described in the language of geometric optics by an ABCD matrix. The system has an optical path length along the optical axis of L .

fewer assumptions. In particular; Kogelnik[46] gives one of the original expositions by drawing an analogy with spherical waves from point sources and applies it to resonators with internal lenses. Siegman[85] includes the effects of complex ABCD matrices (systems with loss) and systems which start and finish in different indices of refraction. Baues[10] focuses on 2 dimensional Hermite-Gaussian beams, and Tache[88] gives a similar exposition for Laguerre-Gaussian beams. Finally; Porras, Alda, and Bernabeu[1] derive the relationship for imperfect, non-Gaussian beams. We will generally follow Siegman's exposition here with some simplifications.

Referring to the situation shown in Figure 2-1, we may express the wavefront at the output of the ABCD system in terms of the wavefront at the input through the Huygens-Fresnel integral. In one dimension the Huygens-Fresnel integral may be written as[85]

$$u_2(x_2) = \sqrt{\frac{i}{\lambda L}} \int_{-\infty}^{\infty} u_1(x_1) e^{-ik\rho(x_1, x_2)} dx_1, \quad (2-20)$$

where $\rho(x_1, x_2)$ is the optical path length between $u_1(x_1)$ and $u_2(x_2)$. The square root in front of the integral is due to the fact that we are working in one transverse dimension. This expression says that the field at a point downstream is given by summing coherently the contribution from every point of the field upstream as if each point were the source of a spherical wave (Huygen's principle).

Expressing the output field in terms of the input field thus reduces to finding the optical path length for a ray starting at position x_1 on wavefront u_1 and ending at position x_2 on wavefront u_2 . We start by expressing the angles of incoming and outgoing ray in terms of the ABCD coefficients and the ray's positions. The ABCD relationship gives us

$$x_2 = Ax_1 + B\theta_1 \quad \& \quad \theta_2 = Cx_1 + D\theta_1, \quad (2-21)$$

from which can arrive at

$$\theta_1 = \frac{1}{B}(x_2 - Ax_1) \quad \& \quad \theta_2 = \frac{1}{B}(Dx_2 - x_1). \quad (2-22)$$

Since we are dealing with a paraxial optical system, each point along the wavefront R_1 will be imaged to a corresponding point along the wavefront R_2 . In particular, the ray which travels directly along the optical axis will be imaged to the point on R_2 which lies on the optical axis. We know also from Fermat's principle that these two rays, the one which travels along the optical axis and the one which travels along the path from x_1 to x_2 must have the same optical path length.

Calculating the optical path length (OPL) from the centers of curvature, we find that the OPL of the ray traveling along the optical axis is given by

$$OPL_{axis} = R_1 + L - R_2, \quad (2-23)$$

where the minus sign is due to the sign convention for the radius of curvature. The optical path length along the path between the centers of curvature for the ray connecting $u_1(x_1)$ and $u_2(x_2)$ is given by

$$\begin{aligned} OPL_{x_1 \rightarrow x_2} &= \sqrt{R_1^2 + x_1^2} + \rho(x_1, x_2) - \sqrt{R_2^2 + x_2^2} \\ &\approx R_1 + \frac{x_1^2}{2R_1} + \rho(x_1, x_2) - R_2 - \frac{x_2^2}{2R_2}. \end{aligned} \quad (2-24)$$

Equating the two and solving for ρ gives us the path length through the optical system

$$\rho(x_1, x_2) = L - \frac{x_1^2}{2R_1} + \frac{x_2^2}{2R_2}. \quad (2-25)$$

All that remains is to express the radii of curvature of the point sources in terms of the ABCD elements. Again making use of the paraxial approximation we can write

$$R_1 \approx \frac{x_1}{\theta_1} \quad \& \quad R_2 \approx \frac{x_2}{\theta_2}, \quad (2-26)$$

and using equations (2-22) to simplify these to

$$R_1 = \frac{Bx_1}{x_2 - Ax_1} \quad \& \quad R_2 = \frac{Bx_2}{Dx_2 - x_1}. \quad (2-27)$$

Substituting back into (2-25) gives

$$\rho(x_1, x_2) = L + \frac{1}{2B}(Ax_1^2 - x_1x_2 + Dx_2^2). \quad (2-28)$$

The Huygens-Fresnel integral through a cascaded, paraxial optical system may therefore be written as

$$u_2(x_2) = e^{-ikL} \sqrt{\frac{i}{B\lambda}} \int_{-\infty}^{\infty} u_1(x_1) \exp \left[-i \frac{\pi}{B\lambda} (Ax_1^2 - 2x_1x_2 + Dx_2^2) \right] dx_1. \quad (2-29)$$

The B under the square root out front is necessary to enforce energy conservation. This form of the Huygens-Fresnel diffraction integral tells us that an arbitrary optical wave may be propagated through a cascaded, paraxial optical system with only the knowledge of its ABCD matrix elements. Note that this would not be valid if there were significant apertures in the system, in which case the wave would have to be propagated through the system step by step.

We are now in a position to consider what happens to a Gaussian beam upon propagation through a general ABCD system. To do so we choose

$$u_1(x_1) = \exp \left(-i \frac{\pi x_1^2}{\lambda q_1} \right), \quad (2-30)$$

which is the equivalent of equation (2–3) in one transverse dimension with the arbitrary initial choice of $P = 0$. Recall that the q parameter, or complex radius of curvature is given by (2–8). The Huygens-Fresnel integral expressed in terms of the ABCD matrix elements, (2–29), becomes

$$\begin{aligned} u_2(x_2) &= \sqrt{\frac{i}{B\lambda}} e^{-ikL} \int_{-\infty}^{\infty} \exp \left[-i \frac{\pi x_1^2}{q_1 \lambda} - i \frac{\pi}{B\lambda} (Ax_1^2 - 2x_1x_2 + Dx_2^2) \right] dx_1 \\ &= \sqrt{\frac{q_1}{Aq_1 + B}} e^{-ikL} \exp \left[-\frac{i\pi x_2^2}{\lambda} \frac{q_1(AD - 1) + BD}{B(Aq_1 + B)} \right]. \end{aligned} \quad (2-31)$$

Reading off the new complex q parameter gives

$$q_2 = \frac{B(Aq_1 + B)}{q_1(AD - 1) + BD} = \frac{Aq_1 + B}{Cq_1 + D}, \quad (2-32)$$

where we have used the unitarity relation $AD - BC = 1$.

What is rather amazing is that making the replacements $q_1 \rightarrow R_1$ and $q_2 \rightarrow R_2$ converts this equation into the relationship between the radii of curvature of a spherical wave propagating through an ABCD system.[85] Perhaps we should not be surprised since the propagation through the system of both is governed by Fermat's principle.

Equation (2–32) is extremely powerful. It allows us to propagate a Gaussian beam through an arbitrary optical system using the simple techniques of linear algebra. We will make use of this fact repeatedly in Chapter 3 to understand the resonant modes of an optical cavity.

As stated at the beginning of this section, our derivation was for a centered optical system which starts and ends in materials of the same index of refraction. In the case of a centered optical system which begins and ends in materials of differing indices of refraction we must instead use the relationship

$$\frac{q_2}{n_2} = \frac{A(q_1/n_1) + B}{C(q_1/n_1) + D}, \quad (2-33)$$

the derivation of which can be found in Siegman's book on lasers[85]. In fact, as Siegman points out, it is often helpful when dealing with materials of differing indices




Figure 2-2: A decentered ABCD system. r_i , θ_i and r_f , θ_f are the normal ray coordinates with respect to the centered ABCD system. ρ , φ express the (small) misalignment of the ABCD system with respect to the optical axis which is shown in green.

of refraction to scale the Gaussian q parameter by the index of refraction. In the next section we will look at how one can deal with decentered/misaligned optical systems using a modified version of the ABCD methods of geometric optics.

2.3 Matrix Methods for Treating Misaligned Optical Systems

We have seen the usefulness of the ABCD matrix formalism for understanding the propagation of Gaussian beams through centered paraxial optical systems. In reality, all optical systems are subject to misalignment. Our goal now is to extend the usefulness of the ABCD matrix formalism to handle misaligned optical systems. We will follow the exposition of Wang Shaomin[82], which develops the theory into a 4x4 matrix formalism. Other authors, such as Siegman[85], develop the theory into a 3x3 matrix formalism. The two formalisms are of course identical in function, but the 4x4 methods have a slightly more clear interpretation.

2.3.1 Derivation of the 4x4 Matrix Formalism

Let r_i , θ_i and r_f , θ_f denote the position and angle, with respect to the true optical axis of the system, of the incoming and outgoing rays respectively. Additionally, let ρ and φ denote the position and angle of the misaligned optical axis of the ABCD system with

respect to the true optical axis of the system (see Figure 2-2). We want to derive the equivalent of the input output relation given in equation (2-19). Following Shaomin[82], we will do so in three steps:

1. Express the coordinates of the input ray in terms of the misaligned optical axis of the ABCD system, denoting them as the misaligned coordinates r_{im}, θ_{im} .
2. Propagate the misaligned coordinates through the ABCD system with a simple application of the ABCD law, denoting the output misaligned coordinates as r_{fm}, θ_{fm} .
3. Express the misaligned output coordinates in terms of the true optical axis.

The misaligned coordinates are given by

$$r_{im} = r_i - \rho \quad \& \quad \theta_{im} = \theta_i - \varphi, \quad (2-34)$$

in the paraxial limit. The output misaligned coordinates are given in terms of the input misaligned coordinates by the ABCD relation

$$r_{fm} = Ar_{im} + B\theta_{im} \quad \& \quad \theta_{fm} = Cr_{im} + D\theta_{im}. \quad (2-35)$$

Finally, the output ray coordinates are given in terms of the misaligned output coordinates by

$$r_f = r_{fm} + \rho + L\varphi \quad \& \quad \theta_f = \theta_{fm} + \varphi. \quad (2-36)$$

Putting all of this together gives

$$\begin{aligned} r_f &= Ar_i + B\theta_i + (1 - A)\rho + (L - B)\varphi \\ \theta_f &= Cr_i + D\theta_i - C\rho + (1 - D)\varphi. \end{aligned} \quad (2-37)$$

From here we can see that if we are willing to add a couple of dummy elements to our ray vector, equation (2-16), then we can write the full effect on an optical ray in terms of a 4x4 matrix whose upper left hand quadrant is the original ABCD matrix. In this

formalism the 4x4 matrix will look like

$$\begin{bmatrix} r_f \\ \theta_f \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} A & B & \alpha\rho & \beta\varphi \\ C & D & \gamma\rho & \delta\varphi \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_i \\ \theta_i \\ 1 \\ 1 \end{bmatrix}, \quad (2-38)$$

with

$$\begin{aligned} \alpha &= 1 - A & \beta &= L - B \\ \gamma &= -C & \delta &= \begin{cases} 1 - D & \text{transmission} \\ -1 - D & \text{reflection} \end{cases} \end{aligned} \quad (2-39)$$

Note that the δ term is slightly different for an ABCD system which is used in reflection such as a mirror. A collection of 4x4 ABCD matrices can be found in appendix 10.2.4.

2.3.2 Definition of the Optical Axis

In applying the misaligned ABCD matrices, we must be careful about our definition of the optical axis. For systems with two centers of curvature, such as a lens, the optical axis is unambiguously defined as the line which passes through those two centers of curvature. For a system with only a single center of curvature, however; there is some ambiguity between displacement, ρ , and tilt, φ , of the optical axis. In fact, the only requirement for the definition of the optical axis[82] is that

$$\frac{\rho}{R} + \varphi = C, \quad (2-40)$$

where R is the radius of curvature of the surface and C is an arbitrary constant.

Figure 2-3 shows three possible definitions for a system with a single center of curvature. This choice of definition of the misaligned optical axis factors into the calculation of the misaligned ABCD matrix. As a concrete example, consider the 4x4




Figure 2-3: A system with a single center of curvature is shown as well as three different possible definitions of the misalignment terms. ρ_1, φ_1 is chosen such that the misalignment is strictly displacement; ρ_2, φ_2 is chosen such that the misalignment is mixed between tilt and displacement; and ρ_3, φ_3 is chosen such that the misalignment is strictly tilt.

matrix for a misaligned spherical mirror for the three different coordinate systems shown in Figure 2-3;

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ \frac{2}{R} & 1 & -2\frac{\rho_1}{R} & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 0 & 0 & 0 \\ \frac{2}{R} & 1 & -2\frac{\rho_2}{R} & -2\varphi_2 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 0 & 0 & 0 \\ \frac{2}{R} & 1 & 0 & -2\varphi_3 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2-41)$$

In appendix 10.2.4 we choose the mixed coordinate system corresponding to ρ_2, φ_2 , since the other coordinate systems are special cases of this one.

CHAPTER 3 RESONANT OPTICAL INTERFEROMETERS

In the last chapter we developed a number of techniques for dealing with laser radiation and its propagation through optical systems. In this chapter we want to apply those techniques to resonant optical interferometers. The canonical resonant optical interferometer is the Fabry-Perot interferometer consisting of two partially reflective mirrors facing each other[85]. We will take such a two mirror optical cavity as our example system throughout most of this chapter.

3.1 Quasi-static Interactions Between Laser Radiation and Resonant Optical Interferometers

Lets begin by looking at the quasi-static interactions between laser radiation and resonant optical interferometers. By quasi-static we mean systems where the changes in the laser frequency and the interferometer geometry occur on timescales much longer than the buildup time of the interferometer. We will extend our understanding to dynamic systems in Section 3.2. For now lets start with a plane wave calculation from which we will understand the effect of the reflectivity of the mirrors on this interaction. From there we will move on to see how the curvature of the mirrors defines the spatial mode accepted by the interferometer and affects the stability. We will end with a discussion of the importance of the alignment of the interferometer.

3.1.1 Plane Wave Calculation of Resonant Optical Cavities

As an illustrative example, lets begin by calculating the reflectivity and transmissivity of the Fabry-Perot cavity shown in Figure 3-1 using the infinite sum method. Afterwards we will apply the more straightforward method of using the boundary conditions at the mirrors to solve for the external and internal fields. From this we will be introduced to such concepts as photon storage time, finesse, and power buildup. Throughout this section we will assume that the mirror apertures are infinite and that the mirrors and intra-cavity medium are lossless.




Figure 3-1: A two mirror Fabry-Perot cavity. The cavity is pumped with a laser from the left hand side, and the input and output mirrors have a reflectivity and transmissivity of r_1, t_1 and r_2, t_2 respectively. The one-way phase picked up by the field traversing the cavity is denoted by ϕ . Additionally, the fields in the cavity are shown and labeled in blue.

For both calculations we will employ a couple of conventions which are fairly standard within the optics community. First, we will use the convention that a field picks up a π phase flip upon reflection from a surface with higher index of refraction and no phase upon reflection from an interface with lower index of refraction. This is one of many possible conventions which enforces energy conservation at the mirror, a topic which is discussed in greater detail in Section 3.3. Second, for the boundary condition calculation, we will treat the internal fields as originating at the surface of the mirror from which they are reflected (see Figure 3-1).

For the infinite sum method, we simply sum together all of the fields coming from the multiple internal reflections in the cavity. For the reflected field we have

$$\begin{aligned}
 E_R &= E_0 r_1 - E_0 t_1^2 r_2 e^{-i2\phi} - E_0 t_1^2 r_1 r_2^2 e^{-i4\phi} - E_0 t_1^2 r_1^2 r_2^3 e^{-i6\phi} + \dots \\
 &= E_0 \left[r_1 - t_1^2 r_2 e^{-i2\phi} \sum_{n=0}^{\infty} (r_1 r_2 e^{-i2\phi})^n \right] \\
 &= E_0 \left[r_1 - \frac{t_1^2 r_2 e^{-i2\phi}}{1 - r_1 r_2 e^{-i2\phi}} \right]. \tag{3-1}
 \end{aligned}$$

Similarly, for the transmitted field we have

$$\begin{aligned}
 E_T &= E_0 t_1 t_2 e^{-i\phi} + E_0 t_1 t_2 r_1 r_2 e^{-i3\phi} + E_0 t_1 t_2 r_1^2 r_2^2 e^{-i5\phi} + \dots \\
 &= E_0 t_1 t_2 e^{-i\phi} \sum_{n=0}^{\infty} (r_1 r_2 e^{-i2\phi})^n \\
 &= E_0 \frac{t_1 t_2 e^{-i\phi}}{1 - r_1 r_2 e^{-i2\phi}}
 \end{aligned} \tag{3-2}$$

This method of calculation illustrates clearly that the interference in an optical cavity is from all of the fields simultaneously.

A quicker method of calculating the fields in a resonant optical cavity, which also has the advantage of being more easily applicable to multiple mirror cavities, is to write down the boundary conditions at each of the mirrors in the cavity and solve the simultaneous equations. In our example case of the two mirror Fabry-Perot interferometer the boundary conditions are

$$\begin{aligned}
 E_1 &= t_1 E_0 - r_1 E_2 e^{-i\phi} & E_R &= r_1 E_0 + t_1 E_2 e^{-i\phi} \\
 E_2 &= -r_2 E_1 e^{-i\phi} & E_T &= t_2 E_1 e^{-i\phi}
 \end{aligned} \tag{3-3}$$

Solving these simultaneous equations gives us the fields in and around the cavity in terms of the input field

$$\begin{aligned}
 E_1 &= E_0 \frac{t_1}{1 - r_1 r_2 e^{-i2\phi}} & E_R &= E_0 \left[r_1 - \frac{t_1^2 r_2 e^{-i2\phi}}{1 - r_1 r_2 e^{-i2\phi}} \right] \\
 E_2 &= E_0 \frac{-t_1 r_2 e^{-i\phi}}{1 - r_1 r_2 e^{-i2\phi}} & E_T &= E_0 \frac{t_1 t_2 e^{-i\phi}}{1 - r_1 r_2 e^{-i2\phi}}
 \end{aligned} \tag{3-4}$$

All of the fields in and around the cavity go through a rapid transition as the round trip phase passes through a multiple of 2π . As the phase picked up in a round trip crosses through resonance the promptly reflected field and the cavity leakage field interfere destructively causing the reflected power to drop dramatically. Simultaneously, the circulating power increases drastically and the power transmitted through the cavity goes up in proportion with the circulating power.

What determines the round trip phase? The phase picked up by a plane wave as it propagates is given simply by $\phi = k\ell = \frac{2\pi}{\lambda}\ell = \frac{2\pi f}{c}\ell$, but recall from our discussion in Sections 2.1.1 and 2.1.2 that monochromatic radiation picks up an additional phase shift beyond that of a plane wave, known as the Gouy phase. Denoting the Gouy phase picked up by the fundamental mode as ψ , the one way phase picked up is

$$\phi = kL + (n + m + 1)\psi, \quad (3-5)$$

where n, m denote the number of the higher order mode and L is the length of the cavity. The Gouy phase is larger for the higher order modes than it is for the fundamental Gaussian mode; this property is what allows an optical resonator to act as a mode cleaner. We will run into this fact in Chapters 6 and 8 when we discuss the eigenspectrum of the Advanced LIGO input mode cleaner.

3.1.2 Finesse, Buildup, and Storage Time of a Two Mirror Optical Cavity

The condition for resonance in an optical cavity is

$$kL + (n + m + 1)\psi = 2\pi z,$$

where z is any integer. If we scan the length of the cavity we will therefore find that the spacing between any two resonances is equal to the wavelength of the light; $\delta\ell = L_2 - L_1 = \lambda$. If we instead chose to scan the frequency of the laser then we would find that the frequency separation between any two resonances is

$$\nu_0 = \frac{c}{L}, \quad (3-6)$$

where c is the speed of light and L is the *round-trip* length of the cavity.

Another useful parameter for characterizing an optical cavity is a parameter known as the finesse. The finesse (denoted \mathcal{F} and not to be confused with the coefficient of finesse) is a measure of the ability of a resonator to distinguish between different wavelengths, or equivalently frequencies, of light. It is defined as the free spectral range

of the cavity divided by the full width at half max (FWHM) of the resonance peak;

$$\mathcal{F} = \frac{\nu_0}{2\nu_c}. \quad (3-7)$$

Here we have used ν_c to denote the half width at half max¹. Squaring any one of the equations from (3-4) and calculating where it becomes half of its maximum gives

$$\cos(2\phi_c) = 1 - \frac{(1 - r_1 r_2)^2}{2r_1 r_2} \Rightarrow \sin \phi_c = \frac{1 - r_1 r_2}{2\sqrt{r_1 r_2}}. \quad (3-8)$$

The finesse of the cavity is therefore

$$\mathcal{F} = \frac{\pi}{2 \sin^{-1} \left(\frac{1 - r_1 r_2}{2\sqrt{r_1 r_2}} \right)} \simeq \frac{\pi \sqrt{r_1 r_2}}{1 - r_1 r_2}. \quad (3-9)$$

Figure 3-2 shows the transmission and reflection spectrum of two mirror cavities with varying finesse. The left hand side of that plot shows the effect of changing the finesse of a cavity with a fixed reflectivity of the rear mirror. This plot also illuminates an effect known as coupling; as we increase the reflectivity of the front mirror relative to the rear mirror we couple more light into and through the cavity. In the right hand plot the reflectivity of the two mirrors is set to be equal, and the reflectivity of both mirrors is varied simultaneously to adjust the finesse of the cavity. This type of cavity where the reflectivity of the two mirrors is equal is known as impedance matched (drawn from an analogy with circuits) or critically coupled. It should be remembered that we have been working in the limit of a lossless cavity. For high finesse optical cavities, the losses cannot be omitted in considering how they will act. The power recycling mirror of the Advanced LIGO interferometers is near impedance matched with a transmissivity of 3% because the tiny losses inside of the arm cavities add up to roughly this after so many round trips.

¹ We will see later that this parameter, known as the cavity pole, is useful in characterizing the dynamic response of an optical cavity.

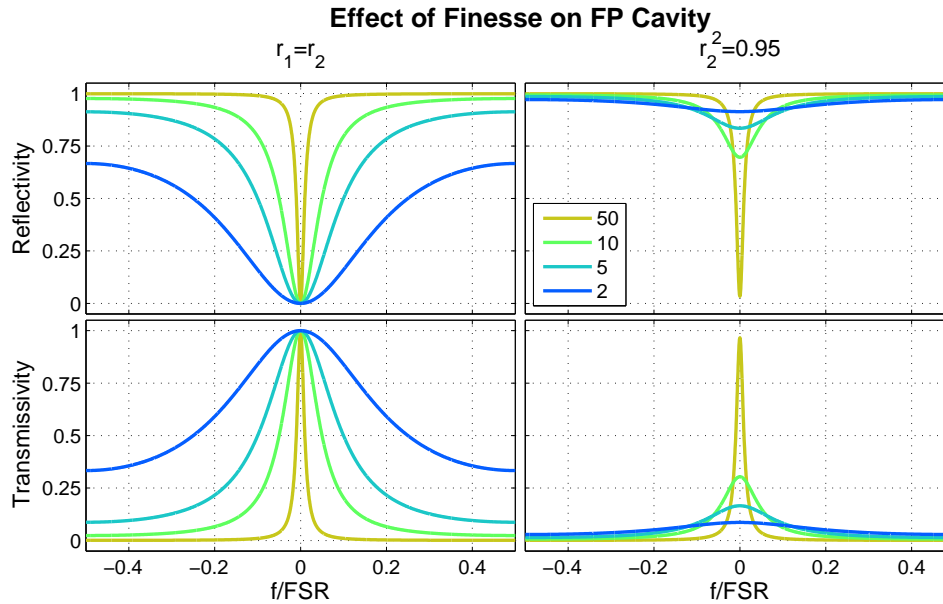


Figure 3-2: The effects of differing finesse on the transmission and reflection coefficients of a two mirror Fabry-Perot resonator. The right hand plot chooses the reflectivity of the rear mirror to be 95% and varies the reflectivity of the front mirror to set the Finesse. Notice that the amount reflected and transmitted on resonance varies with the finesse. The left hand plot sets the front and rear reflectivities to be equal and varies the common reflectivity to set the Finesse. This case is known as the impedance matched case; notice that the reflectivity always drops to zero on resonance and the transmissivity always goes to 1.

There is another interesting parameter which we can derive from simple plane wave considerations. The photon storage time tells us how long, on average, a photon spends inside of the cavity. Knowing the photon storage time is important for understanding the dynamical response of a resonant optical interferometer. Processes which happen on much shorter timescales, such as passing gravitational waves at high frequencies, will not be sensed by the cavity because the effects are averaged out by the photon storage time.

To derive the photon storage time we imagine monitoring the power in the field at the front mirror of an optical cavity beginning immediately after the input field has been extinguished. After m round trip times the power in the field will be

$$P\left(m\frac{L}{c}\right) = (r_1^2 r_2^2)^m P_0,$$

where r_1 and r_2 are the amplitude reflectivity of the input and output mirrors, L/c is the round trip time of the cavity, and P_0 is the initial power in the cavity. The number of photons in the cavity will therefore be

$$\Phi\left(m\frac{L}{c}\right) = (r_1^2 r_2^2)^m \Phi_0 = e^{-m\frac{L}{c\tau}} \Phi_0 \quad (3-10)$$

where Φ_0 is the initial number of photons in the cavity and

$$\tau = -\frac{L}{c \ln(r_1^2 r_2^2)} \quad (3-11)$$

In the limit that the losses are small on each round trip (high finesse limit) we may approximate this discrete formula with the continuous version

$$\Phi(t) = e^{-\frac{t}{\tau}} \Phi_0 \quad (3-12)$$

Now we can see that τ is the average lifetime of a photon inside of the optical cavity. We can expand this in the high finesse limit to obtain

$$\tau = -\frac{L}{c \ln(r_1^2 r_2^2)} \simeq \frac{L r_1 r_2}{c(1 - r_1^2 r_2^2)} \simeq \frac{L \mathcal{F}}{c 2\pi} \quad (3-13)$$

Finally, one other parameter derivable from the plane wave view of a two mirror cavity is the buildup. Looking at the power in the field E_1 from equations 3-4 on resonance we find that the buildup is

$$\left(\frac{E_1}{E_0}\right)^2 = \frac{t_1^2}{(1 - r_1 r_2)^2} \simeq t_1^2 \left(\frac{2\tau c}{L}\right)^2 \quad (3-14)$$

where the approximations are in the high finesse limit.

All of our calculations so far of the properties of resonant optical cavities have used plane waves, but as we saw in Section 2.1 laser beams spread as they propagate due to diffraction. A stable optical cavity must therefore account for this spreading by using curved mirrors or the interference that we assumed in the plane wave calculations would not occur.

3.1.3 Spatial Modes, Stability, and Gouy Phase in a Two Mirror Optical Cavity

To understand the spatial modes of optical resonators, it is helpful to first understand what happens to a Gaussian beam upon reflection from a curved mirror. Most laser mirrors are ground with spherical instead of the more ideal parabolic surfaces². This is for two reasons; 1) A sphere has no true center, meaning that the alignment tolerances for centering on a spherical mirror are set strictly by the clipping of the beam at the edges of the mirror and 2) spheres are more easy to grind than more complex shapes like parabolas. To see how this affects a Gaussian beam upon reflection we need to calculate the phase picked up by the beam.

The height of a sphere relative to the optical axis is given by

$$h(r) = R \left(1 - \sqrt{1 - \frac{r^2}{R^2}} \right), \quad (3-15)$$

where r is the distance from the optical axis, R is the radius of curvature, and h is the height of the sphere with $h = 0$ at the optical axis. In the limit $r/R \ll 1$ we are justified in expanding the square root with a Taylor expansion and dropping all but the lowest order terms;

$$h(r) = R \left[1 - \left(1 - \frac{r^2}{2R^2} - \frac{r^4}{8R^4} - \dots \right) \right] \simeq \frac{r^2}{2R}. \quad (3-16)$$

The higher order terms which we dropped will lead to the most common type of aberration known as spherical aberration.

² Parabolic surfaces are more ideal because they match the curvature of the phase fronts of the diffraction limited Gaussian beam.

Lets now consider a Gaussian beam, (2-13), which reflects off of a mirror with a spherical profile and picks up the phase associated with it

$$\begin{aligned}
 u(r, z) &= \frac{\omega_0}{\omega} \exp \left[-r^2 \left(\frac{1}{\omega^2} + i \frac{k}{2R_{bi}} \right) \right] * e^{-i2k \frac{r^2}{2R_m}} \\
 &= \exp \left[-r^2 \left(\frac{1}{\omega^2} + i \frac{k}{2} \left(\frac{1}{R_{bi}} + \frac{2}{R_m} \right) \right) \right]
 \end{aligned}
 \tag{3-17}$$

The new radius of curvature of the wavefront of the beam is therefore

$$\frac{1}{R_{bf}} = \frac{1}{R_{bi}} + \frac{2}{R_m},
 \tag{3-18}$$

where R_{bf} is the final radius of curvature of the beam, R_{bi} is the initial radius of curvature of the beam, and R_m is the radius of curvature of the mirror. Notice that if we choose $R_m = -R_{bi}$ (the minus sign is due to our sign convention), then the radius of curvature of the new beam is $R_{bf} = -R_{bi}$. I.E. we have reflected the beam back onto itself. This sets the requirement for the spatial modes of a two mirror resonator; the radius of curvature of the wavefront must match the radius of curvature of the two end mirrors.

Following Siegman[85] we can determine the characteristics of the beam which will resonate inside of the two mirror cavity by demanding, through equation (2-11), that the radius of curvature of the beam match the radius of curvature at each end mirror;

$$-R_1 = z_1 + \frac{z_R}{z_1} \quad \& \quad R_2 = z_2 + \frac{z_R}{z_2} \quad \& \quad L = z_2 - z_1.
 \tag{3-19}$$

Here R_i is the radius of curvature of the each mirror and z_i is its location. The minus sign in the first equation is because of our sign convention for the radii of curvature of the mirrors (which are viewed from inside the resonator) and the beam (which is viewed

along the optical axis). Solving this system of equations for z_1 , z_2 , and z_R yields

$$\begin{aligned} z_1 &= -\frac{L(R_2 - L)}{R_1 + R_2 - 2L} \quad \& \quad z_2 = \frac{L(R_1 - L)}{R_1 + R_2 - 2L} \\ \& \quad z_R^2 &= \frac{L(R_1 - L)(R_2 - L)(R_1 + R_2 - L)}{(R_1 + R_2 - 2L)^2}. \end{aligned} \quad (3-20)$$

It is instructive for a two mirror cavity to recast these equations in terms of the cavity g parameters defined by

$$g_i = 1 - \frac{L}{R_i}. \quad (3-21)$$

Doing so gives

$$\begin{aligned} z_1 &= -L \frac{g_2(1 - g_1)}{g_1 + g_2 - 2g_1g_2} \quad \& \quad z_2 = L \frac{g_1(1 - g_2)}{g_1 + g_2 - 2g_1g_2} \\ \& \quad z_R^2 &= L^2 \frac{g_1g_2(1 - g_1g_2)}{(g_1 + g_2 - 2g_1g_2)^2}. \end{aligned} \quad (3-22)$$

Looking at z_R we can see that the criterion for the cavity mode to have a real Rayleigh range, i.e. for a cavity mode to exist, is that

$$0 \leq g_1g_2 \leq 1. \quad (3-23)$$

This is the condition on resonator stability.

Resonator stability is also closely tied to the round trip Gouy phase in an optical cavity. To see this with the two mirror cavity, lets calculate its round trip Gouy phase from the parameters in equation (3-22). Substituting into equation (2-14) we find that the Gouy phase at the first and second mirrors respectively is

$$\eta_1 = -\tan^{-1} \left(\frac{g_2(1 - g_1)}{\sqrt{g_1g_2(1 - g_1g_2)}} \right) \quad \& \quad \eta_2 = \tan^{-1} \left(\frac{g_1(1 - g_2)}{\sqrt{g_1g_2(1 - g_1g_2)}} \right).$$

The round trip Gouy phase is given by $2(\eta_2 - \eta_1)$ which can be simplified using the trigonometric identity $\tan^{-1}(a) + \tan^{-1}(b) = \tan^{-1} \left(\frac{a+b}{1-ab} \right)$ giving

$$\Delta\eta = 2(\eta_2 - \eta_1) = \tan^{-1} \left(\sqrt{\frac{1 - g_1g_2}{g_1g_2}} \right) = 2 \cos^{-1}(\sqrt{g_1g_2}) \quad (3-24)$$

Recalling the range of \cos^{-1} we see that the round trip Gouy phase in a two mirror cavity must be between 0 and 2π .

3.1.4 Spatial Modes, Stability, and Gouy Phase in Complex Resonant Optical Interferometers

Matching the radii of curvature of the fields to the radii of curvature of the mirrors is a useful way of understanding the spatial modes of a two mirror cavity; in this section we want to look at how one can understand the spatial modes of more complex optical interferometers. The resonant mode in a multi-element linear cavity will still match the radii of curvature to the end mirrors, but this is not necessarily true in a ring cavity³.

In Section 2.2.2 we saw that the q parameter, which fully characterizes a Gaussian beam, can be propagated through a paraxial optical system simply by knowing its ABCD matrix

$$q_f = \frac{Aq_i + B}{Cq_i + D}. \quad (3-25)$$

In order to obtain perfect interference in a complex resonant optical cavity, the beam must be identical to itself after a round trip through the cavity. If we write down the *round trip* ABCD matrix of the optical cavity the ideal resonant mode of the cavity must satisfy

$$q = \frac{Aq + B}{Cq + D}. \quad (3-26)$$

³ The recycling cavities of the aLIGO interferometer provide an example of a multi-element linear cavities in which the phase front curvature matches the curvature of the two end mirrors while the input and output mode cleaners provide an example of a ring cavity in which the field curvature and mirror curvature need not have any particular relationship

Solving the quadratic equation for q gives the resonant q parameter as

$$\begin{aligned} q &= \frac{1}{2C} \left(A - D \pm \sqrt{(A - D)^2 + 4BC} \right) \\ &= \frac{1}{2C} \left(A - D \pm 2\sqrt{\left(\frac{A + D}{2}\right)^2 - 1} \right), \end{aligned} \quad (3-27)$$

where the plus or minus sign is chosen to give a positive imaginary part by convention and the second form of the equation uses the unitarity condition of the ABCD matrix, $AD - BC = 1$.

Recalling the meaning of the imaginary part of the q parameter, (2-9), as the Rayleigh range of the laser beam, we can see that in order for the q parameter to be complex we must have

$$-1 < \frac{A + D}{2} < 1. \quad (3-28)$$

Rearranging puts this into a more recognizable form

$$0 < \frac{A + D + 2}{4} < 1, \quad (3-29)$$

which shows that this quantity plays the same role for complex resonant optical cavities as the product of the g factors plays for a two mirror cavity. For a somewhat more rigorous derivation involving raising the round trip ABCD matrix to the n th power and letting n go to infinity the reader is referred to Siegman[85].

Indeed, we can even calculate the round trip Gouy phase in the cavity (up to a multiple of 2π) from the ABCD matrix. We will not go through the derivation here, but it can be found in the literature [6][27]. The equation for the round trip Gouy phase is

$$\Delta\eta = 2 \cos^{-1} \left(\text{sign}(B) \sqrt{\frac{A + D + 2}{4}} \right), \quad (3-30)$$

where $\Delta\eta$ refers to the round trip Gouy phase in the cavity. Flipping back to equation (3-24) shows that $(A + D + 2)/4$ truly plays the same role as the product of the g factors in dealing with a complex multi-element cavity.

3.1.5 Alignment Modes of Resonant Optical Cavities

In the last section we understood how to calculate the resonant spatial modes in optical cavities. In this section we want to ask how the resonant optical axis moves as the mirrors get misaligned inside of the cavity. The strategy for doing so is to apply the same idea as the last section, except with the 4x4 matrices of Section 2.3.1. Since the misalignment terms are tied up inside of the 4x4 matrices this method does not present a simple closed form solution like the spatial modes of the previous section. The best way to understand its utility is therefore to look at an example.

Lets consider how the optical axis moves in a two mirror cavity with radii of curvature R_1 and R_2 for the input and output mirror respectively, and a distance L between them. It is conventional in this case to think of movement of the optical axis at the waist of the cavity, and we will do the same here. If we define the 4x4 matrix for reflection from a misaligned spherical mirror to be $M(R, \rho, \phi)$ and the 4x4 matrix for translation over a distance d to be $T(d)$ (refer to appendix 10.2.4 for the form of these matrices), then the round trip matrix beginning at the waist is

$$\mathcal{M}(\phi_1, \phi_2) = T\left(L - \frac{L(R_1 - L)}{R_1 + R_2 - 2L}\right) M(R_1, 0, \phi_1) T(L) M(R_2, 0, \phi_2) T\left(\frac{L(R_1 - L)}{R_1 + R_2 - 2L}\right). \quad (3-31)$$

The distance from the waist to the rear mirror, used in the first and last matrices, comes from our derivation of z_2 in Section 3.1.3. We have chosen not to include the displacement terms of the two mirrors since we are interested in alignment effects.

We might naively expect that the eigenvalues of this matrix will tell us something about the resonant axis of the cavity. The problem with this method is that the vectors of optical rays must have ones as the last two elements, and the eigenvalue calculation of the full matrix does not respect this requirement. In addition, the elements which we want to solve for, ϕ_1 and ϕ_2 , are buried inside of the matrix. The idea in this case is to solve the eigenvalue equation for our own particular choices of eigenvectors by finding conditions on the misalignment terms ϕ_1 and ϕ_2 . The eigenvectors we will consider are

the typical choice in this case; strict translation and strict tilt at the waist of the cavity. These are sometimes called the hard and soft modes of a two mirror cavity because translation at the waist is dynamically unstable (soft) when radiation pressure is taken into account while tilt at the waist is dynamically stable (hard).

To solve for the soft mode (translation at the waist) we consider the equation

$$\begin{bmatrix} r \\ 0 \\ 1 \\ 1 \end{bmatrix} = \mathcal{M}(\phi_1, \phi_2) * \begin{bmatrix} r \\ 0 \\ 1 \\ 1 \end{bmatrix} \quad (3-32)$$

Solving this system of equations we find the relatively simple answer that the soft mode of the cavity is given by

$$\phi_1 = -\frac{r}{R_1} \quad \& \quad \phi_2 = -\frac{r}{R_2}. \quad (3-33)$$

As a check on our calculation consider the case where $R_1 = R_2 = R$ in which we find that the soft mode of the cavity is caused by common tilt of the two mirrors by equal amounts. For the hard mode we instead consider the equation

$$\begin{bmatrix} 0 \\ \theta \\ 1 \\ 1 \end{bmatrix} = \mathcal{M}(\phi_1, \phi_2) * \begin{bmatrix} 0 \\ \theta \\ 1 \\ 1 \end{bmatrix}. \quad (3-34)$$

Solving the system of equations gives

$$\begin{aligned} \phi_1 &= -\theta \frac{(L - R_1)(L - R_1 - R_2)}{R_1(R_1 + R_2 - 2L)} = \theta \frac{(L - R_1 - R_2)z_1}{R_1 L} \quad \& \\ \phi_2 &= \theta \frac{(L - R_2)(L - R_1 - R_2)}{R_2(R_1 + R_2 - 2L)} = \theta \frac{(L - R_1 - R_2)z_2}{R_2 L}, \end{aligned} \quad (3-35)$$

where z_1 and z_2 are the positions of the front and rear mirrors derived in Section 3.1.3.

These are slightly more obfuscated than the soft mode; it is helpful to consider the ratio

of the two coordinates

$$\frac{\phi_1}{\phi_2} = -\frac{R_2(L - R_1)}{R_1(L - R_2)}, \quad (3-36)$$

which becomes -1 in the limit $R_1 = R_2 = R$. This is exactly what we would expect for a symmetric two mirror cavity; the soft mode is caused by common tilt of the two end mirrors while the hard mode is caused by differential tilt of the two end mirrors.

In this section we looked at how one can understand through mathematics the static effects of resonant optical interferometers. In particular, we saw how the photon storage time and power buildup are both related to the finesse of the cavity; we understood how to calculate the resonant spatial modes of these cavities, and we saw how to understand how the optical axis of the cavity changes with misalignments. In the next section we want to address dynamical questions related to these cavities. In particular, we want to ask what happens when the frequency of the light is changing at rates comparable to or greater than the photon storage time. We would also like to know how the cavity responds if the mirrors are vibrating at frequencies comparable to the photon storage time since this will tell us how the arms of the Advanced LIGO interferometers see passing gravitational waves.

3.2 Dynamics of Two Mirror Optical Cavities

We now want to relax the quasi-static condition enforced in the previous section and look at the dynamical effects in resonant optical cavities. In particular, we will be interested in how the light in reflection and transmission of the resonant optical interferometer is affected by changes in the laser frequency or motions of the mirrors which happen on timescales comparable to or much quicker than the photon storage time of the interferometer. Since the dynamics of complex systems such as the full Advanced LIGO interferometers are the topic of entire Ph.D. theses[70][72] [94] we will focus on the simple two mirror Fabry Perot cavity. Understanding this simple system will give us an understanding of the more complex systems since they consist simply of coupled Fabry-Perot cavities.

3.2.1 Laser Frequency Noise and a Two Mirror Cavity

We will start by looking at how laser frequency variations interact with a cavity of fixed length in the sideband picture. Note that thinking in the sideband picture usually consists of expanding an exponential of the sort $e^{i\delta \cos(\omega t)}$. There is an exact expansion of these terms given by the Fourier-Bessel identity

$$e^{iz \cos \phi} = \sum_{k=-\infty}^{\infty} i^k J_k(z) e^{ik\phi}. \quad (3-37)$$

We generally drop all but the lowest order terms, $k = 0, \pm 1$, in which case the small angle Taylor expansion will be found to be nearly identical for $z \ll 1$.

Our derivation of the interaction between a static cavity and laser frequency noise will follow that of Camp et. al[14] and Regehr[72]. We will let the electric field of the noiseless carrier be given by

$$E = E_0 e^{i\omega_0 t}, \quad (3-38)$$

where E_0 and ω_0 are the carrier amplitude and frequency, respectively. If we write the laser frequency noise as

$$\nu = \nu_0 + \delta\nu \cos \omega t, \quad (3-39)$$

then the phase of the electric field at time t is given by

$$\begin{aligned} \phi(t) &= 2\pi\nu_0 t + \frac{2\pi\delta\nu}{\omega} \sin \omega t \\ &= \omega_0 t + \frac{2\pi\delta\nu}{\omega} \sin \omega t. \end{aligned} \quad (3-40)$$

Lets take a step back and think about the quantities in this equation for a moment. There are three different variables with the units of frequency contained in it: ω_0 , $\delta\nu$, and ω .

The mean laser frequency is given by ω_0 and is of the order of 100 THz. The *magnitude* of the frequency variations is given by $\delta\nu$ in units of $\frac{1}{s}$ as opposed to $\frac{rad}{s}$ like the laser frequency. Finally, ω is the *frequency* of the laser frequency variations. In the end we

will be interested in asking questions of the type: what is the amplitude $\delta\nu$ of the laser frequency variations at a given frequency ω ?

Putting this together, the laser field is given by

$$E = E_0 e^{i\omega_0 t} \exp\left(i \frac{2\pi\delta\nu}{\omega} \sin \omega t\right).$$

If the amplitude of the frequency variations $\delta\nu$ is small, then we are justified in expanding the second exponential in the small angle expansion or the single sideband expansion mentioned above;

$$\begin{aligned} E &= E_0 e^{i\omega_0 t} \left(1 + i \frac{2\pi\delta\nu}{\omega} \sin \omega t\right) \\ &= E_0 e^{i\omega_0 t} \left[1 + i \frac{\pi\delta\nu}{\omega} (e^{i\omega t} - e^{-i\omega t})\right] \\ &= E_0 \left[e^{i\omega_0 t} + \frac{\pi\delta\nu}{\omega} (e^{i(\omega_0+\omega)t} - e^{i(\omega_0-\omega)t})\right] \end{aligned} \quad (3-41)$$

Hence, the laser frequency variations may be treated as simple sidebands at frequency ω around the laser frequency ω_0 . These are what are typically referred to as audio sidebands in order to distinguish them from the radio frequency (RF) sidebands added to the beam by phase modulation. The transmissivity and reflectivity of a two mirror cavity to these frequencies is therefore given by our result for the quasi-static case (equations (3-4)).

It is often convenient to have a linear approximation for these nonlinear transfer functions. We may arrive at an approximation of this nature by noting that if the carrier is on resonance, then the one way phase can be written as

$$\phi = kL = 2\pi N + \frac{L}{c}\omega \quad (3-42)$$

for $N \in \mathbb{N}$. In this regime we may expand the exponentials in (3-4) as

$$e^{-i2\phi} \simeq 1 - \frac{2L}{c}i\omega, \quad (3-43)$$

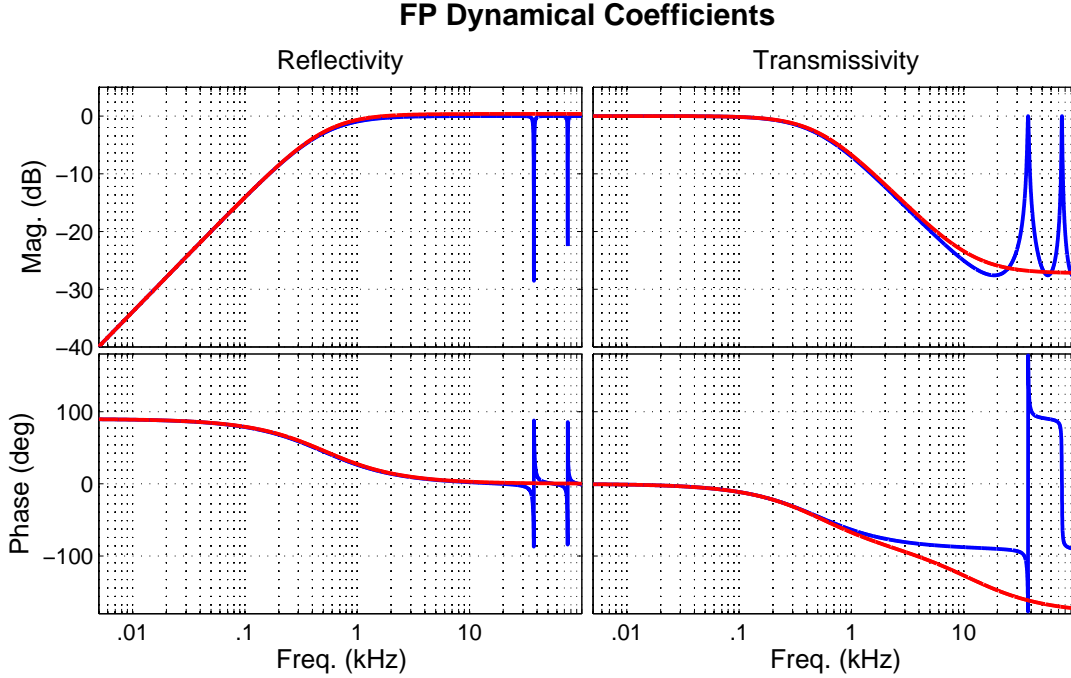


Figure 3-3: The reflection and transmission coefficients of a two mirror Fabry-Perot cavity together with their low frequency approximations. The full reflection and transmission coefficients, shown in blue, are given by (3-4), and the approximations, shown in orange, are given by (3-44). The parameters chosen for this plot are: $R_1 = 92\%$, $R_2 = 92\%$, and $L = 4\text{km}$.

which transforms the transmission and reflection coefficients into

$$r(\omega) \simeq \frac{1}{r_1} \frac{i\omega + \frac{r_1 - r_2}{r_2} \frac{c}{2L}}{i\omega + \nu_0} \quad \& \quad t(\omega) \simeq -\frac{t_1 t_2}{2r_1 r_2} \frac{i\omega - \frac{c}{L}}{i\omega + \nu_0} \quad (3-44)$$

where we have defined the quantity

$$\nu_0 = \frac{1 - r_1 r_2}{r_1 r_2} \frac{c}{2L} = \frac{1}{\tau(1 - r_1 r_2)} \quad (3-45)$$

usually known as the cavity pole. Recall that this quantity showed up in Section 3.1.2 during our derivation of the finesse of a Fabry-Perot cavity. Figure 3-3 shows the actual reflection and transmission coefficients of an impedance matched cavity together with these approximations. There we can see that the approximations do not start to deviate significantly until roughly a few cavity poles away from the center of the FSR.

3.2.2 Length Noise in a Two Mirror Cavity

We would now like to understand how vibrations of one of the mirrors (or equivalently fluctuations in the length of the cavity) show up in the laser frequency. Since a gravitational wave passing by our detector will cause audio band length fluctuations in the length of the arm cavities, this is a very important issue to understand. We will follow Regehr[72] with a few modifications (i.e. we won't worry about the RF sidebands) as we did in the last section.

Thinking solely about the rear mirror of the cavity for a moment, lets consider how length fluctuations of a mirror will affect the beam reflected off of it. We will assume that the length fluctuations are described by $x(t) = X \cos(\omega t)$ where X is the amplitude of the fluctuations and ω is the angular frequency. Let the incident beam coming from the input mirror be given by $E_{1,0} e^{-i\omega_0 t}$ where we are using two subscripts to distinguish the fields; the first number is used to denote the field as labeled in Figure 3-1 and the second number is used to distinguish the carrier from the audio sidebands. The reflected beam with the phase variations of the moving mirror therefore has the form

$$\begin{aligned} E_{2,:} &= -r_2 E_{1,0} e^{-i\phi_0} e^{-i\omega_0 t} e^{-i2kX \cos(\omega t)} \\ &\simeq -r_2 E_{1,0} e^{-i\phi_0} \left(J_0(2kX) e^{-i\omega_0 t} + iJ_{-1}(2kX) e^{-i(\omega_0+\omega)t} - iJ_1(2kX) e^{-i(\omega_0-\omega)t} \right). \end{aligned} \quad (3-46)$$

So the effect of the moving mirror is to add audio sidebands at frequencies $\pm\omega$ to the reflected beam. Note that if we were to use the small angle expansion of the Bessel functions; $J_0(x) \simeq 1$, $J_1(x) \simeq \frac{x}{2}$, and $J_{-1}(x) \simeq -\frac{x}{2}$; then we would recover the small angle expansion of the exponential. In general, the incident beam will already have three frequencies on it since the input mirror will reflect this beam back. Taking

$$E_{1,:} = E_{1,0} e^{-i\phi_0} e^{-i\omega_0 t} + E_{1,1} e^{-i\phi_1} e^{-i(\omega_0+\omega)t} + E_{1,-1} e^{-i\phi_{-1}} e^{-i(\omega_0-\omega)t} \quad (3-47)$$

as the beam reflected off of the mirror we now find the reflected beam to be

$$\begin{aligned}
E_{2,:} = & -r_2 \left\{ \left[J_0(-2kX) e^{-i\phi_0} E_{1,0} + iJ_1(-2kX) e^{-i\phi_1} E_{1,1} - iJ_{-1}(-2kX) e^{-i\phi_{-1}} E_{1,-1} \right] e^{-i\omega_0 t} \right. \\
& + \left[J_0(-2kX) e^{-i\phi_{-1}} E_{1,-1} + iJ_1(-2kX) e^{-i\phi_0} E_{1,0} \right] e^{-i(\omega_0 - \omega)t} \\
& \left. + \left[J_0(-2kX) e^{-i\phi_1} E_{1,1} - iJ_{-1}(-2kX) e^{-i\phi_0} E_{1,0} \right] e^{-i(\omega_0 + \omega)t} \right\}. \tag{3-48}
\end{aligned}$$

Looking at the upper sideband; the effect of the moving mirror is to add a small amount of power $-iJ_{-1}(-2kX)$ to the sideband from the carrier. From this we can write down a set of field equations just as we did in the static case (equations 3-4). As an example lets consider the upper sideband for which the equations will be

$$\begin{aligned}
E_{2,1} &= -r_2 e^{-i\phi_1} J_0(-2kX) E_{1,1} - r_2 e^{-i\phi_0} iJ_1(-2kX) E_{1,0} \\
E_{1,1} &= -r_1 e^{-i\phi_1} E_{2,1}. \tag{3-49}
\end{aligned}$$

Solving the simultaneous equations yields

$$\begin{aligned}
E_{1,1} &= \frac{-r_1 r_2 e^{-i\phi_1}}{1 - r_1 r_2 J_0(2kX) e^{-i2\phi_1}} iJ_1(2kX) e^{-i\phi_0} E_{1,0} \\
E_{2,1} &= \frac{r_2}{1 - r_1 r_2 J_0(2kX) e^{-i2\phi_1}} iJ_1(2kX) e^{-i\phi_0} E_{1,0} \tag{3-50}
\end{aligned}$$

Lets simplify these by making the small X approximations described above and by assuming that the carrier $E_{1,0}$ is on resonance such that $\phi_0 = 2\pi m$ for some $m \in \mathbb{N}$. Doing so gives the fields in the forward and backward directions (note that it does not make sense to talk about the fields being reflected since they are sourced inside of the cavity), referred to as $E_{T,1}$ and $E_{R,1}$ respectively, as

$$\begin{aligned}
E_{T,1} &= E_0 \frac{t_1 t_2 r_1 r_2^2 e^{-i2\phi_1} ikX}{(1 - r_1 r_2 e^{-i2\phi_1})(1 - r_1 r_2)} \\
E_{R,1} &= E_0 \frac{-t_1^2 r_2^2 e^{-i\phi_1} ikX}{(1 - r_1 r_2 e^{-i2\phi_1})(1 - r_1 r_2)}, \tag{3-51}
\end{aligned}$$

where E_0 is now the carrier field at the input to the cavity, i.e. the same E_0 as in equations (3-4). The one way phase picked up by the audio sideband ϕ_1 is going to

be the phase picked up by the carrier frequency ϕ_0 plus the phase picked up by the audio frequency. Since the carrier is assumed resonant, we can set $\phi_0 = 0$ and the one way phase by the audio sideband will be simply

$$\phi_1 = \frac{L}{c}\omega. \quad (3-52)$$

Notice that in this case, both the fields at the reflection and transmission ports both have a very similar transfer function to the frequency noise in transmission which is given by (3-44) and plotted in Figure 3-3. We might have expected this because, as pointed out earlier, the fields are sourced inside of the cavity.

3.3 Symmetry in Optical Systems

A person who works in the field of optics quickly catches on that high quality laser systems are generally symmetric. That is to say, if one places a retro-reflecting mirror at some point in the system, then the beam will propagate directly back along the entire system to the laser unless a special type of element known as an optical diode is included. To understand why this is lets begin by looking at what is demanded by energy conservation in lossless optical systems without optical activity (without polarization effects).

3.3.1 Energy Conservation in Lossless Systems with no Optical Activity

We want to consider a lossless optical system such as the ones shown in Figure 3-4 where the initial and final indices of refraction are the same. In particular, we are interested in any two port optical system which is described by the relations

$$E_3 = r_{31}E_1 + t_{32}E_2 \quad \& \quad E_4 = t_{41}E_1 + r_{42}E_2. \quad (3-53)$$

In previous sections we treated the reflectivity and transmissivity coefficients of a system as being real, positive numbers and included the phase as a minus sign in front of the appropriate terms. In this section we will treat these coefficients as complex numbers which have both phase and magnitude. In fact, we will understand at the end where this




Figure 3-4: Three different examples of lossless optical systems; a beamsplitter, a mirror, and a black-box system which could contain e.g. an optical cavity. What is important for the discussion of symmetry in optical systems is that the system obey the relations given in equations 3–53. It is shown that energy conservation demands symmetry in this case, in addition to demanding particular phase relationships between the various fields.

phase relationship arises from. We will, in general, follow the derivation of Loudon[52] in his section on lossless beamsplitters, though we will draw some broader implications.

Applying energy conservation to this system yields

$$\begin{aligned}
 |E_1|^2 + |E_2|^2 &= |E_3|^2 + |E_4|^2 \\
 &= |r_{31}E_1 + t_{32}E_2|^2 + |t_{41}E_1 + r_{42}E_2|^2 \\
 &= (|r_{31}|^2 + |t_{41}|^2)|E_1|^2 + (|t_{32}|^2 + |r_{42}|^2)|E_2|^2 + 2\Re[(r_{31}t_{32}^* + r_{42}^*t_{41})E_1E_2^*].
 \end{aligned}
 \tag{3-54}$$

The only way to satisfy this equation for all possible E_i is if

$$|r_{31}|^2 + |t_{41}|^2 = |r_{42}|^2 + |t_{32}|^2 = 1 \quad \& \quad r_{31}t_{32}^* + r_{42}^*t_{41} = 0. \tag{3-55}$$

The first set of equations states a fact which was obvious from the outset; if a system is lossless, then the amount of power reflected and transmitted by it must sum to the total incident power. Writing the terms in the last equation in terms of their amplitude and phase, e.g. $r_{31} = |r_{31}|e^{i\phi_{31}}$, leads to the relations

$$|r_{31}||t_{32}| = |r_{42}||t_{41}| \quad \& \quad \phi_{31} - \phi_{32} + \phi_{42} - \phi_{41} = \pm\pi. \tag{3-56}$$

The first of these equation together with (3–55) implies that

$$|r_{31}| = |r_{42}| \quad \& \quad |t_{32}| = |t_{41}|. \quad (3-57)$$

We have shown that for any lossless black box system which can be described in terms of reflectivity and transmissivity coefficients that the magnitude of the reflectivity and transmissivity must be the same from both directions, i.e. that the system must be symmetric when viewed from the two ports. This applies to beamsplitters, mirrors, and optical cavities as well as more complex optical systems. In addition, we have shown that energy conservation demands that these reflectivity and transmissivity coefficients obey a particular phase relationship.

There are two common conventions for this phase relationship, one is typically used in classical optics, and the other is typically used in quantum optics. For simplicity let us define $|r_{31}| = |r_{42}| = r$ and $|t_{32}| = |t_{41}| = t$ where r and t are positive real numbers. The classical choice keeps all of the numbers real and includes all of the phase shift in the reflectivity coefficient off of the higher index of refraction

$$r_{31} = -r, \quad r_{42} = r, \quad t_{32} = t, \quad t_{41} = t. \quad (3-58)$$

The quantum optics choice includes the phase shift symmetrically in the transmissivity coefficients;

$$r_{31} = r, \quad r_{42} = r, \quad t_{32} = te^{-i\frac{\pi}{2}}, \quad t_{41} = te^{-i\frac{\pi}{2}}. \quad (3-59)$$

There are often very useful reasons to break the symmetry of an optical system. For instance, we might want an optical component which transmits the light traveling forwards along the optical axis while reflecting the light traveling backwards. This is useful, for example, in collecting the light reflected off on an optical cavity or interferometer to generate locking signals. Introducing losses into an optical system can break the symmetry in some cases, but it is generally a bad idea in an optical

experiment to throw away photons. Lets examine how we can break the symmetry by exploiting the polarization of the beams.

3.3.2 Optical Symmetry Breaking with Optical Diodes

In this section we want to relax our assumption about optical inactivity from the last section. We will see that the possibility of optical activity allows us to build a device which is known as an optical diode. These optical diodes will allow us pick up the reflected light from an optical apparatus which can be used for sensing and control. In particular, we will look at the two most common types of optical diodes which are constructed from a quarter wave plate (QWP) and a Faraday rotator (FR).

In the last section we saw that optically inactive systems exhibited symmetry when viewed from the two different ports. Lets refer to this type of symmetry as two port symmetry. We will see in a moment that the optical diodes still exhibit a type of symmetry in that if we place retro-reflectors at enough of the ports, then the beam will still return to its originating source. We will refer to this type of symmetry as three and four port symmetry for the QWP based and FR based optical diodes respectively.

Figure 3-5 a) shows the quarter wave plate based optical diode. In this diode system the light is passed through a quarter wave plate (QWP) before interacting with the first optical system, OS_1 . This converts the linear polarized light to circularly polarized light, and upon reflecting off of OS_1 the QWP converts the light to be linear polarized in the opposite plane as before. This allows the polarizing beam splitter (PBS) to pick up the reflected light so that it can interact with a separate optical system, OS_2 instead of returning to the laser. Light reflected by this system returns to OS_1 with circular polarization of the opposite handedness and after reflection returns back to the laser.

This system exhibits three port symmetry; meaning that there are two usable ports other than the port at which the laser is injected. Even if only three ports are needed this can be a disadvantage since the scattered light off of the optical system in the reflected

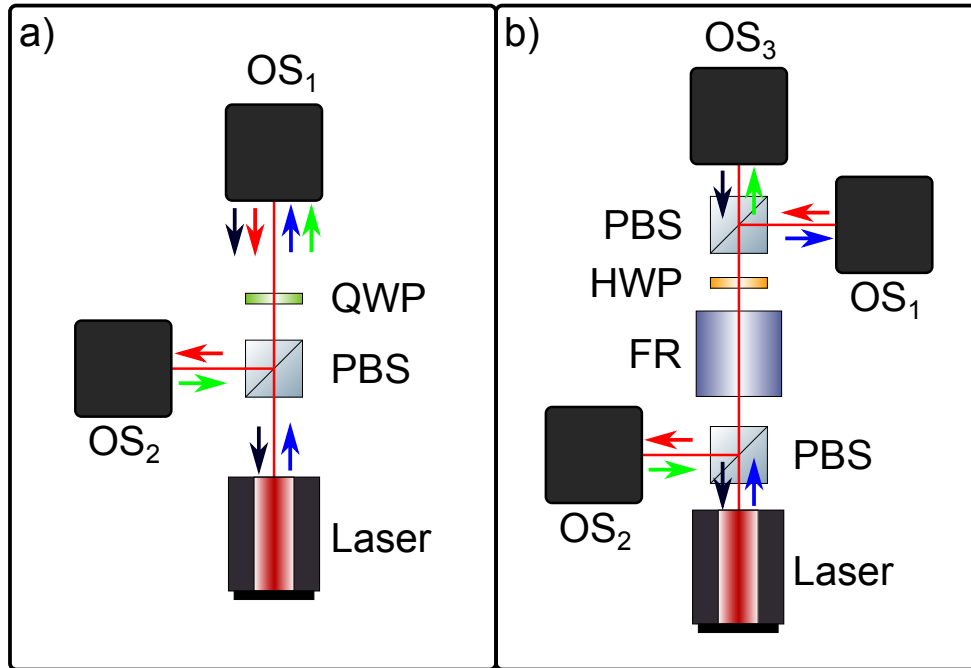


Figure 3-5: Two common types of optical diodes. In a) a diode based on a quarter wave plate is shown which exhibits 3-fold symmetry while in b) a diode based on a Faraday isolator is shown which exhibits 4-fold symmetry. OS: Optical system; any black box optical system defined strictly by a reflection and transmission coefficient. HWP: Half wave plate. QWP: Quarter wave plate. PBS: Polarizing beam splitter. FR: Faraday rotator with 45° of rotation.

port, OS_2 , will be scattered back into OS_1 albeit in a different polarization. Another disadvantage of this type of optical diode is that the transmitted beam which interacts with OS_1 is circularly polarized. Working with circular polarization can be difficult since typical optics have different reflectivities for the S and P polarizations when used at an angle which leads to circularly polarized light becoming elliptical.

Figure 3-5 b) shows a schematic of the Faraday rotator based optical diode which is known as a Faraday isolator. The Faraday effect is discussed in more detail in Section 5.2, but for now all we need to know is that a beam passing through picks up 45° of non-reciprocal polarization rotation. The half wave plate (HWP) after the FR adds another 45° of polarization rotation, but this will be undone when the beam comes back through in reflection (reciprocal). Hence, the transmitted beam shows up at OS_1 and

the reflected beam at OS_2 just like the QWP based diode. The difference between the two shows up in the reflected beam from OS_2 which is not sent back to OS_1 with the FR diode but instead shows up in a new location at OS_4 . Finally, the beam reflected from OS_4 will head back to the laser.

This system exhibits four port symmetry; meaning that there are three usable ports other than the port at which the laser is injected. It has the advantage that scattered light at the reflected port is not injected back into the main system of interest. In addition, it also has the advantage of the QWP based system of having linear polarized light at all ports which is easier to work with in an optical system. These advantages are the reason that a Faraday isolator was chosen as the in-vacuum optical diode for the Advanced LIGO interferometers (discussed at length in Chapter 5).

3.4 Analytical Models of Thermal Deformations

As we saw in Section 3.1.2 the power buildup in high finesse optical cavities can lead to very large circulating powers. In the Advanced LIGO interferometers, circulating powers up to 1 MW are possible in the arm cavities during high power operation. Even with mirrors of very low absorption this leads to a significant amount of heat being deposited in the mirrors. This heat deposition causes two deleterious effects in the mirrors. First, it causes the surfaces of the mirrors to deform in shape which changes the effective radius of curvature, changes the mode structure, and scatters some of the power out of the fundamental Gaussian mode. Second, since the index of refraction of most materials is temperature dependent, transmissive optics develop gradients in their indices of refraction which cause them to act like a lens. These two effects combined mean that the coupling efficiency between different parts of the interferometer typically drops with higher input powers as the mirrors deform and transmissive optics begin to lens.

In this section we want to see how one can understand these effects mathematically. The thermal effects in solid bodies are governed by the heat equation so we will

begin with a quick discussion of its general properties. We'll move on to look at some analytical solutions of the heat equation in a simplified case which we'll use to get an understanding of its general properties. Finally, we'll use the Hello-Vinet[36] calculation to calculate the effective focal length of the thermal effects in the Advanced LIGO interferometers as well as understand the timescales for the onset of these effects.

In order to gain some understanding of the thermal effects in high power optical systems it will be instructive to begin with some analytical models of these effects before discussing numerical results. We will start with a general discussion of the heat equation which governs the thermal effects in the mirrors of resonant interferometers. From there we will discuss the thermal deformation caused by small point absorbers as a way to gain intuition for the general structure of the solutions to the heat equation. Finally, we will end the section with some results from the Hello-Vinet[36] calculation for realistic Gaussian heat sources in the interferometer.

3.4.1 The Heat Equation

The heat equation, or diffusion equation, is[101]

$$\frac{\partial}{\partial t}u(\vec{x}, t) - \frac{k}{\sigma\rho}\nabla^2u(\vec{x}, t) = \frac{1}{\sigma\rho}f(\vec{x}, t). \quad (3-60)$$

Here \vec{x} represents a coordinate in three dimensional space, e.g. (x, y, z) , and t represents time. $u(\vec{x}, t)$ is the temperature of the system and $f(\vec{x}, t)$ describes the heat sources and sinks. The material is characterized by k : the thermal conductivity, ρ : the mass density, and σ : the specific heat. In general all of these material parameters can be dependent on direction and position in the material, but we will assume from here forward that our material is isotropic and homogeneous.

The heat equation is similar to the wave equation except that the time derivative is only first order instead of second order. This changes the solutions significantly from those of the wave equation. In the case of the unforced wave equation the eigenfunctions can be thought of as the traveling waves, $e^{-i(kx \pm \omega t)}$, while the eigenfunctions

of the heat equation can be thought of as spatial waves which die out over time, $e^{-ikx \pm \alpha t}$. Indeed, this is a general property of the heat equation that if there are no heat sources or sinks, the initial temperature distribution will die down to a uniform temperature over time.

For time independent solutions, the heat equation reduces to the Poisson equation

$$\nabla^2 u(\vec{x}) = -\frac{1}{k} f(\vec{x}), \quad (3-61)$$

whose properties are well known from the field of electrostatics[40]. The unforced equation reduces even further to the Laplace equation. As a way of gathering some intuitive feeling for this equation in the next section we will consider how it acts for a point absorber sitting on the surface of a mirror.

The solutions of the heat equation will tell us the temperature distribution, $u(\vec{x}, t)$ throughout the optic, but the quantities which we are truly interested in are the deformation of the surface of the optic and the gradient in the index of refraction. For small temperature variations we can get a good approximation to both of these by integrating the temperature distribution through the optic along its axis of symmetry. To calculate the expansion of the surface we integrate the temperature distribution from the center of the optic to the surface and multiply by the coefficient of thermal expansion. To calculate the effective lens of the optic, we integrate through the whole optic and multiply by the change in the index of refraction with temperature, $\frac{dn}{dT}$.

3.4.2 The Deformation Caused by a Point Absorber on the Mirror Surface

We can get an idea of the time independent properties of the heat equation by considering how it acts with a delta function source. For this purpose we want to approximate the optic as an infinite medium with cylindrical symmetry which fills the half plane ($\rho \geq 0, z \geq 0$) and has no internal sources or sinks of heat. We will first consider a delta function source of heat sitting at the origin heating the surface of the optic.

The Poisson equation sourced by a delta function at the origin is given by

$$\nabla^2 u(\rho, z) = -\frac{\epsilon P}{k} \left(\frac{1}{\rho} \delta(\rho) \delta(z) \right). \quad (3-62)$$

One solution to this equation is the Green function of the Poisson equation which is given by Constantin and Pavel[20] (see also Jackson[40]) as

$$\begin{aligned} u(\rho, z) &= \frac{P}{4\pi k} \left(\frac{1}{|\vec{x} - \vec{x}'_{+z'}|} + \frac{1}{|\vec{x} - \vec{x}'_{-z'}|} \right) \\ &= \frac{P}{4\pi k \sqrt{\rho^2 + \rho'^2 - 2\rho\rho' \cos(\theta - \theta') + (z - z')^2}} \\ &\quad + \frac{P}{4\pi k \sqrt{\rho^2 + \rho'^2 - 2\rho\rho' \cos(\theta - \theta') + (z + z')^2}} \end{aligned} \quad (3-63)$$

$$= \frac{P}{2\pi k \sqrt{\rho^2 + z^2}}, \quad (3-64)$$

where we have chosen the Green function which satisfies the Neumann boundary conditions on the infinite half plane, i.e. no heat flow across the front face of the optic. This form of the Green function necessitates the two terms in the first line above similar to an image charge problem in electrostatics. The last line takes the source point to be at the origin and assumes cylindrical symmetry. As a check on the physicalness of our solution observe that switching to spherical coordinates and integrating over a half sphere gives the total power flow out of any half sphere to be P , independent of ρ , i.e. all of the power deposited in the optic is leaving it. This solution is obviously unphysical because it is divergent right at the origin, but our source is unphysical in the exact same way so this is not surprising.

If we denote the thickness of our optic by h , then integrating (3-64) from $-\frac{h}{2}$ to 0 and multiplying by the coefficient of thermal expansion of the optic, κ , gives the deformation of the surface from a point absorber

$$d(\rho) = \kappa \int_{-h/2}^0 u(\rho, z) dz = \frac{\kappa P}{2\pi k} \log \left(\frac{\sqrt{h^2 + 4\rho^2} + h}{2\rho} \right). \quad (3-65)$$

Deformation of Optic from a 100mW Point Absorber

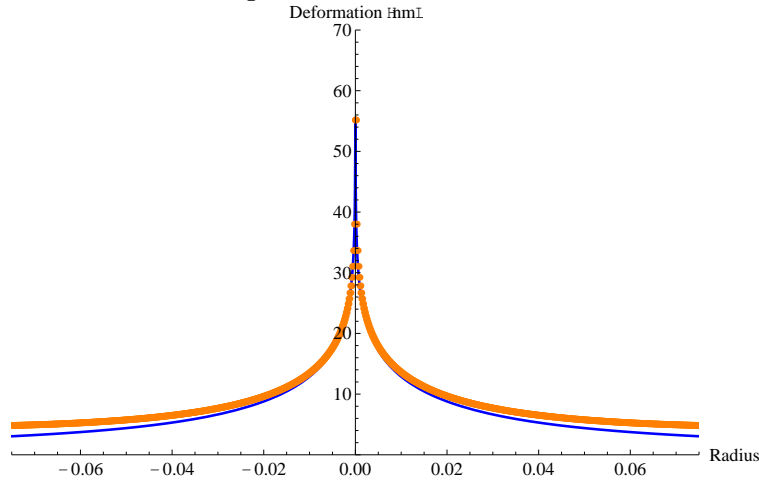


Figure 3-6: The green function solution to the heat equation together with the result of the Hello-Vinet calculation with a delta function absorber. The Green function solution is given by (3-65). The agreement is generally good; diverging a bit towards the edges of the optic due to the different boundary conditions. The Hello-Vinet calculation uses 2000 terms in the power series expansion.

Figure 3-6 shows this deformation together with the Hello-Vinet solution (to be discussed shortly) for a delta function source.

We can come up with a more realistic model of a small absorber by calculating the temperature distribution of the optic for a small square of power instead of a delta function. This calculation can be accomplished by using Green's theorem together with the Green function given in (3-63). In our case Green's theorem[40] states that the temperature distribution may be obtained from the Neumann boundary conditions by

$$u(\vec{x}) = \langle U \rangle_S + \frac{1}{4\pi} \oint_S \frac{\partial u(\vec{x}')}{\partial n'} G_N(\vec{x}, \vec{x}') da', \quad (3-66)$$

where $\langle U \rangle_S$ is the average temperature on the boundary and $G_N(\vec{x}, \vec{x}')$ is the Green function (3-63). Since the Green function goes to zero at infinity, we only have to be concerned with the heat flow across the face of the optic at $z = 0$. If we consider a uniform square source with side length a , then the power flow across the face of the optic is simply a uniform constant. We therefore need to integrate (3-66) with (3-63)

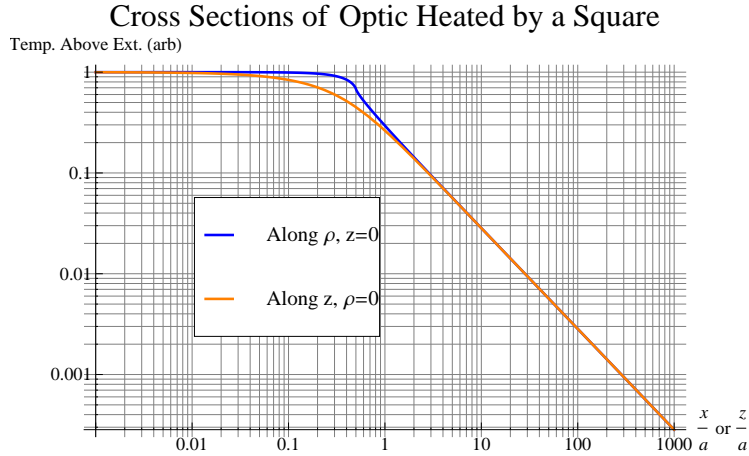


Figure 3-7: The thermal cross sections of an optic of infinite extent heated by a small square. On the vertical axis is shown the temperature above the external temperature in arbitrary units. The horizontal axis shows the cross-section along the x and z directions of the optic, both normalized to a , the square size. Note that the axes are in log-log scale. We can see that in both directions the temperature roll off far from the source is $\frac{1}{d}$ where d is the coordinate of interest. This feature is the same as that of the Green function solution presented in Figure 3-6. The one difference in this case is that the temperature is roughly flat up until reaching the characteristic length of the system, in this case a . This is approximately true for all optics which are heated solely at the face.

expressed in rectangular coordinates. Doing so yields an analytical solution that is too long to write directly in this dissertation; instead we show the plots in Figure 3-7.

From Figures 3-6 and 3-7 we can understand a few general properties of the static solutions to the heat equation. Namely, that the temperature profile at the surface will take the shape of the applied heat and will roll off roughly as one over distance after you get one length scale away from the center of the applied heat. In the case of our applied square of heat the length scale is the side length of the square, but we will see shortly that the same observation applies roughly to heating by a Gaussian beam where the length scale is set by the beam size ω .

3.4.3 Realistic Heat Sources and Time Dependence with the Hello-Vinet Bessel Series Solution

The Hello-Vinet calculation[36] of the thermal state of an optic is a nearly complete analytic solution. Its only assumptions are cylindrical symmetry and linearized radiative

boundary conditions, but the solution is otherwise exact. We won't rehash the entire solution here but instead point out a few of the key steps and show some of the results for the Advanced LIGO optics.

Since the optics are in vacuum, the boundary conditions on the heat equation at the edges of the optic are radiative boundary conditions of the form

$$\frac{\partial u}{\partial n} = \sigma'(u^4 - u_{ext}^4) \quad (3-67)$$

where σ' is the emissivity corrected Stefan-Boltzmann constant, n is the normal coordinate at the boundary of interest, and $u_{ext} \approx 300$ K is the external temperature to which the optic is radiatively synced. Under the assumptions that

$$\frac{u_{max} - u_{ext}}{u_{ext}} = \frac{\delta u}{u_{ext}} \ll 1, \quad (3-68)$$

where u_{max} is the maximum temperature reached at the surface, we can linearize this radiative boundary condition as

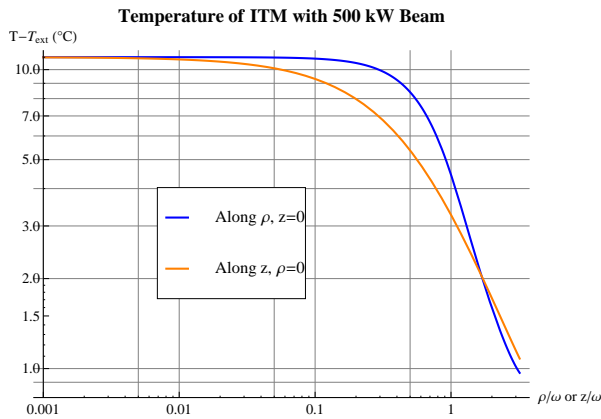
$$\frac{\partial u}{\partial n} = 4\sigma' u_{ext}^3 \delta u. \quad (3-69)$$

If we denote the optic radius as a and its thickness as h , then we may write the full set of Neumann boundary conditions as

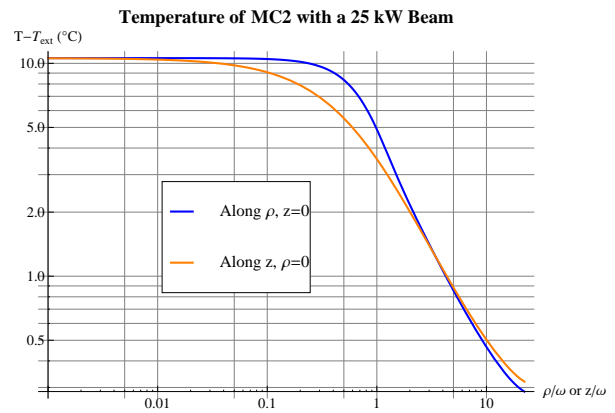
$$\begin{aligned} -k \frac{\partial}{\partial \rho} \delta u(a, z) &= 4\sigma' u_{ext}^3 \delta u(a, z) \\ -k \frac{\partial}{\partial z} \delta u(\rho, \frac{-h}{2}) &= \epsilon I(\rho) - 4\sigma' u_{ext}^3 \delta u(\rho, \frac{-h}{2}) \\ -k \frac{\partial}{\partial z} \delta u(\rho, \frac{h}{2}) &= 4\sigma' u_{ext}^3 \delta u(\rho, \frac{h}{2}), \end{aligned} \quad (3-70)$$

where $I(\rho)$ is the intensity applied at $z = \frac{-h}{2}$, the front face of the optic, and ϵ is the absorption at that face. In the case of bulk absorption we must also include a forcing term on the right hand side of the heat equation so that $f(\vec{x}, t)$ in equation 3-60 is given by

$$f(\vec{x}, t) = \alpha I(\rho) \quad (3-71)$$



(a) Temperature profiles for the ITM, where only surface heating at the HR face is assumed.



(b) Temperature profiles along for MC2, one of the three input mode cleaner mirrors.

Figure 3-8: The temperature profiles of two of the aLIGO optics calculated from the Hello-Vinet calculation. The Hello-Vinet Calculation is described in Section 3.4.3. The power of the beams used on the two optics is typical of the Advanced LIGO interferometer operating at the full design power of 200 W at the input. The x axis in both plots is normalized to the beam radius ω and only surface heating is used. Notice the similarities between these temperature curves and those of Figure 3-7 even though the boundary conditions are more realistic and the heat source is not uniform.

where α is the coefficient of bulk absorption.

With the appropriate boundary conditions in place the solution proceeds in the physicist's favorite fashion, by separation of variables. We end up with an infinite series with \sin and \cos for the z dimension, Bessel functions for the ρ dimension, and exponentials for the time coordinate. For the gory details, the reader is referred to the original paper[36], for our purposes here it will suffice to look at a couple of examples with Gaussian heating of the optic.

Figure 3-8 shows the temperature profiles due to surface absorption of two different optics in the Advance LIGO interferometers. Notice the similarity to Figure 3-7 in that the temperature remains roughly constant until it reaches the length scale of the intensity profile, in this case a Gaussian beam, after which it falls off roughly as $1/d$ as we might expect from the Green function solution of Section 3.4.2.

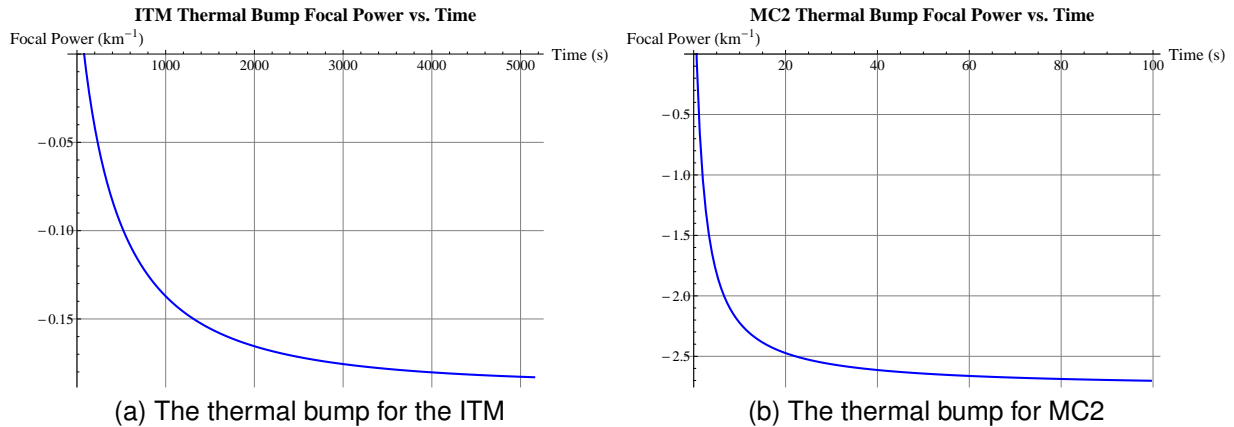


Figure 3-9: The development of the thermal bump on two different optics in the Advanced LIGO interferometers with only surface heating assumed. The thermal bump is expressed in terms of focal power at the center of the optic obtained by integrating the temperature profile in z over half of the optic, multiplying by the coefficient of thermal expansion, and extracting the r^2 coefficient in the Taylor expansion. Notice that the timescales are significantly longer for the ITM than for MC2.

Figure 3-9 shows the timescales for the thermal bump to form on the surface of these same optics. The vertical axis in those plots shows the focal power of the thermal bump on the optic which is given by $2/R$ where the radius of curvature, R is calculated by extracting the r^2 component of the curvature calculated with the Hello-Vinet formalism. Notice that the timescale of the smaller MC2 optic is quicker by a factor of almost 200 than the thermal timescale of the test masses. For optics of the same material (such as MC2 and ITM) this timescale scales roughly as the square of the mass ratio.

Part III

THE ADVANCED LIGO INPUT OPTICS

CHAPTER 4 THE ADVANCED LIGO ELECTRO-OPTIC MODULATOR

In Part II of this dissertation we discussed methods for calculating effects in resonant optical interferometers. In this part we are going to shift modes to a much more grounded topic, namely we will discuss the input optics of the Advanced LIGO interferometers. The University of Florida group was almost solely responsible for the design, production, and installation of the Advanced LIGO input optics for all three interferometers.

When I arrived at the LIGO Livingston Observatory in January of 2012 the design of the input optics had already been completed by the University of Florida team. Early during my visit I assisted with assembly and installation of much of the input optics, but my key role was as a lead in the testing and commissioning phase. Most of the experiments to collect the data shown in this part of the dissertation were designed and built by me with a few notable exceptions. The measurements of the magnetic field of the Faraday isolator magnets, shown in Figure 5-2, were made by Rodica Martin of the University of Florida. Although I did the initial modeling for the mode tracking measurements discussed in Section 6.4.3, the experiment was designed and run by William Korth of the California Institute of Technology. The scatter measurement of Section 6.4.2 were made by David Feldbaum of the University of Florida and Eleanor Jean King of the University of Adelaide with some assistance from me in the interpretation of the results. Finally, the angular controls scheme, described in Section 6.3.2, was devised and perfected by Suresh Doravari and Valery Frolov of the LIGO Livingston Observatory.

The input optics subsystem underwent a significant upgrade for the Advanced LIGO era. It is the task of the input optics to prepare the laser beam from the Pre-Stabilized Laser (PSL) before injection into the main interferometer. The upgraded sensitivity of the aLIGO interferometers places stringent requirements on the stability of the input light

and therefore on the input optics. In particular, the input optics must supply multiple low amplitude and phase noise RF sidebands. They must provide a frequency reference stable enough to not spoil the interferometer's gravitational wave sensitivity. They must provide mode matching and beam pointing actuation into the interferometer. Finally, they must provide in-vacuum isolation between the interferometer's reflected light and the input optics chain. In this particular chapter we will focus on the design and testing of the electro-optic modulators which were designed for aLIGO.

There are a number of much smaller tasks of the input optics which are equally as important as the tasks of these three major pieces[93][89][7] but which will not be discussed in this dissertation. These smaller but nonetheless important components of the input optics include:

- Mode matching between the pre-stabilized laser (PSL) and the input mode cleaner (IMC).
- Remote control over the power injected into the interferometer.
- Diagnostics of the beam quality and power level before injection into the IMC.
- Mode matching between the IMC and the main interferometer.
- Beam steering and seismic isolation of the beam between the IMC and the main interferometer.
- Diagnostics of the beam quality and power level before injection into the main interferometer.
- Stray light control in the IMC, recycling cavities, and in the region between the IMC and main interferometer.

4.1 Requirements

Sensing and controlling most of the seven degrees of freedom of the interferometer is accomplished by some variant of the Pound Drever Hall technique[24]. This technique uses optical heterodyne detection to sense the phase of the returning light from an optical cavity by sensing the phase modulation to amplitude modulation conversion

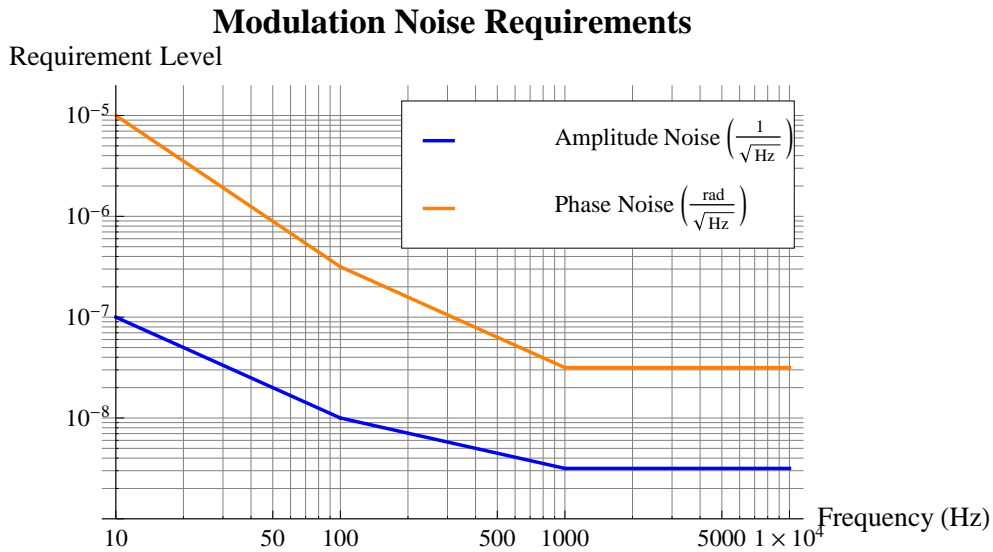


Figure 4-1: The amplitude and phase noise requirements of the aLIGO electro-optic modulator. They are taken from the input optics design requirements document[89] and are based on an early noise model of the full interferometer.

which occurs when the cavity is slightly off resonance. It depends on having low noise radio frequency (RF) sidebands to generate detectable beat notes with the carrier light.

The aLIGO electro-optic modulator (EOM) is required to add three RF sidebands to the beam, one to sense the length and alignment of the input mode cleaner and two others to sense the various longitudinal and angular degrees of freedom of the main interferometer. The initial modulation depth requirement necessitated that the phase modulator be able to produce phase modulation depths up to $\beta = 0.8$, but that Figure was later reduced to $\beta = 0.4$ with a drive strength of 24 dBm. These modulation depths must be attained while maintaining a residual amplitude modulation level below $1 \cdot 10^{-4} \frac{\text{AM}}{\text{PM}}$ ratio[89]. In addition the phase and amplitude noise of the modulator must be below the curves shown in Figure 4-1. The EOM must also be able to survive continuous exposure to the 200 W laser beam generated by the PSL, and it must show minimal thermal lensing while doing so.

4.2 Design

The design of the aLIGO EOM was motivated largely by two factors; the high optical power throughput and the deep modulation depths required. It is hard to find specialized optical crystals, such as those with high electro-optic coefficients, with cross-sectional sizes much larger than a few mm. Hence, with 200 W of laser power, the optical power density at the center of a beam which is not significantly clipped is of the order of $0.1 \frac{\text{MW}}{\text{cm}^2}$. In addition to the possibility of damage to the crystal by such power densities, it is important that the optical crystal not create a significant thermal lens. Another practical consideration arises from the fact that deep modulation depths require voltage levels across the crystal of hundreds of volts which gives a power dissipation on the order of Watts, a level above the rating of common inductors.

Table 4-1 lists some of the key optical properties for the materials considered for aLIGO. Lithium Niobate, LiNbO_3 , is the most common optical modulator material and was the material used in initial LIGO, but its thermal properties were not sufficient for the aLIGO optical power because the high optical absorption combined with the radically different refractive index shifts in the y and z planes leads to significant elliptical thermal lensing. In the end a rubidium titanyl phosphate (RTP) crystal with dimensions of 4 mm x 4 mm x 40 mm was chosen for its high damage threshold, low optical absorption, and low thermal lensing ellipticity.

In addition to its excellent damage and thermal properties, the modulation depths attainable with RTP are not significantly below those of the other materials considered. All of the materials considered make use of the r_{33} coefficient which is utilized by sending light through the crystal with the same polarization as the applied electric field. This has clear advantages over older forms of modulators which required the electric field to be applied along the optical axis necessitating clear electrodes. The phase picked up while traversing through the part of the crystal submerged in the electric field

Table 4-1: The key properties of a number of different materials considered for the aLIGO phase modulators.

Property	Units/Conditions	RTP	RTA	KTP	LiNbO ₃
Damage Threshold	MW/cm ² (10 ns, 1064nm)	>600	400	600	280
n_z	1064 nm	1.82	1.89	1.84	2.16
Absorption	cm ⁻¹ 1064 nm	< 5 10 ⁻⁴	< 5 10 ⁻³	< 5 10 ⁻³	< 5 10 ⁻³
r_{33}	pm/V	39.6	40.5	36.3	30.8
$n_z^3 r_{33}$	pm/V	239	273	226	306
Conductivity	$\frac{1}{\Omega cm}$, 10 MHz	$\sim 10^{-9}$	3 10 ⁻⁷	-	6 10 ⁻⁶
dn _y /dT	10 ⁻⁶ /K	2.79	5.66	13	5.4
dn _z /dT	10 ⁻⁶ /K	9.24	11.0	16	37.9

Source: UF LIGO Group “Upgrading the Input Optics for High Power Operation.” LIGO DCC T060267 (2006)[93] and Mueller, G. and R. Amin “EO-Modulators for Advanced LIGO Part I.” LIGO DCC T020025 (2003)[65]

Note: Rubidium titanyl phosphate, RTP, was chosen for its thermal properties and high damage threshold.

is given approximately in this case by[98]

$$\Delta\phi = -2\pi\frac{d}{\lambda}\Delta n \approx -\pi\frac{d}{\lambda}n_z^2r_{33}E_z = -\pi\frac{d}{\lambda}n_z^3r_{33}\frac{V}{a}, \quad (4-1)$$

where d is the length of the material in the electric field, n_z is the nominal index of refraction, r_{33} is the electro-optic coefficient, and a is the distance between the electrodes. As we can see, the important material property is the product $n_z^3r_{33}$.

Referring to Table 4-1 we see that this value is not significantly less for RTP than for LiNbO₃. The crystals are also wedged at 2.85 degrees on both ends which avoids a parasitic interferometer and helps to prevent RFAM by using the natural birefringence of the material to dump any residual polarization.

Having selected RTP for the optical material it is still necessary to design an electrical circuit which can give the voltages required across the crystal for deep modulation depths. In order to simplify the circuit design the aLIGO EOM has three




Figure 4-2: A schematic of the EOM crystal and one of the resonant circuits. The crystal is wedged to reduce RFAM and prevent parasitic interferometers. There are three separate pairs of electrodes, with one being longer than the other two and each having its own resonant circuit (only one is shown). The resonant circuit uses an inductor, L , in series with the crystal capacitance to get a build up in the voltage on resonance, with the buildup being limited by the losses in the inductive path, R_{loss} . A parallel capacitor, C , allows the impedance to be set to 50Ω on resonance.

separate pairs of electrodes across the crystal; two 15 mm long electrodes for the 9.1 MHz and 45.5 MHz sidebands and one 7 mm long electrode for the 24.0 MHz sideband. Each one of these electrodes has an individual LC circuit wrapped around it as shown in Figure 4-2 (only one circuit is shown). The LC circuit has a capacitor in parallel adjusting the value of which allows the impedance to be set to 50Ω on resonance. The inductors in the circuit are hand wound air coils around polyurethane cores because commercial coils were not able to handle the power load required by the resonant buildup.

The total impedance of the circuit is straightforward to calculate and is given by

$$Z_t = \frac{Z_c(R + Z_{cc} + Z_L)}{Z_c + R + Z_{cc} + Z_L} = \frac{1 + RC_c s + LC_c s^2}{(C + C_c)s + RCC_c s^2 + LCC_c s^3}, \quad (4-2)$$

where $R = R_{loss}$ is the resistive loss in the circuit which is dominated by the hand wound inductors, and $s = i\omega$ is the standard s of the Laplace transform. From this we can calculate the voltage gain across the crystal as

$$G = \frac{V_{cc}}{V} = \frac{1}{1 + RC_c s + LC_c s^2} = \frac{\omega_0^2}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2}. \quad (4-3)$$

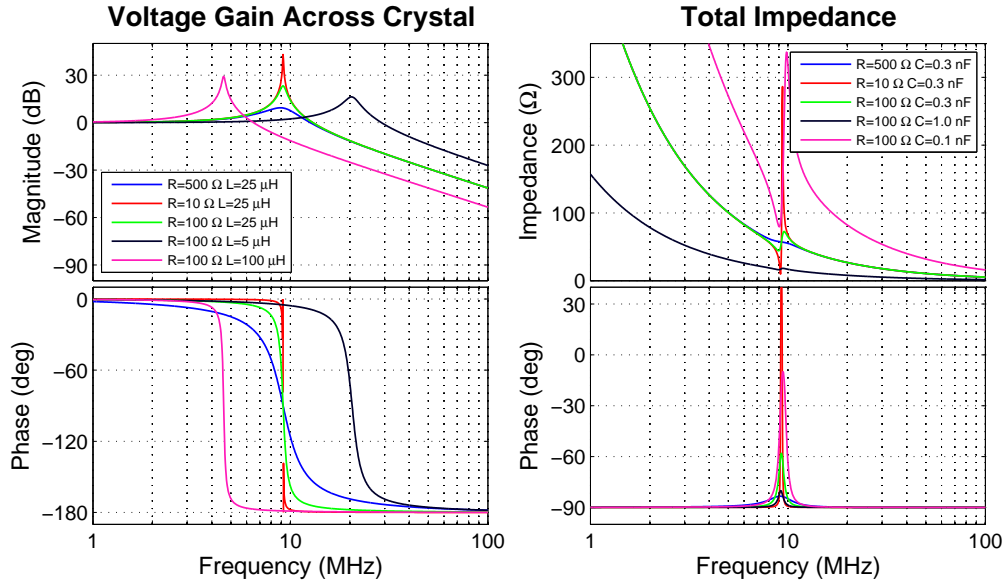


Figure 4-3: The response of the EOM resonant circuit. The left hand plots show the magnitude and phase of the voltage gain across the crystal while the plots on the right hand side show the total impedance presented to the voltage source. On the left hand side the effects of changing the inductance and the resistive losses is shown. On the right hand side the tuning of the capacitor and resistive losses is shown. These plots assume a crystal capacitance of $C_c = 12$ pF.

In the last statement we have made the replacements

$$\omega_0 = \frac{1}{\sqrt{LC_c}} \quad \& \quad Q = \sqrt{\frac{L}{C_c R^2}} \quad (4-4)$$

in analogy with the simple harmonic oscillator equation. Notice that when the circuit is driven on resonance, $s = i\omega_0$, then the voltage gain becomes

$$G_{res} = -iQ \quad (4-5)$$

so that the quality factor sets the achievable voltage gain. Figure 4-3 shows Bode style plots of these two key parameters of this circuit. The left hand plots show the resonant voltage gain across the crystal as the inductance and losses are varied while the right hand plots show the total impedance as the parallel capacitance and losses are varied. All of the plots assume a crystal capacitance of $C_c = 12$ pF.

We can make a few observations from these two plots and from equations (4-4). First, notice that, as expected, the losses in the circuit determine the Q of the resonance for a given center frequency. Notice also that increasing the value of the inductance increases the resonant frequency, as expected. Additionally, and perhaps less expected, decreasing the value of the inductance decreases the Q of the resonance. Looking at the total impedance, which is shown in the right hand set of plots, we can see that increasing the value of the parallel capacitor, C , indeed increases the impedance overall. This plot also shows that a high Q resonance has the disadvantage of making the impedance vary wildly near the resonance, making it difficult to hit 50Ω right on resonance.

4.3 Measured Phase Modulation Indices

A phase modulated beam can be written as

$$E = E_0 e^{-i(\omega t + \beta \sin(\Omega t))} \quad (4-6)$$

where we have neglected to include the spatial dependence since it is not important in this calculation. In this equation ω represents the carrier frequency of the laser beam, β represents the modulation index, and Ω the frequency of the RF modulation. As in Section 3.2.1 we can expand this with the Fourier-Bessel identity (eq (3-37)) or the small angle expansion to get

$$E \approx E_0 \left[e^{i\omega t} + \frac{\beta}{2} e^{i(\omega+\Omega)t} - \frac{\beta}{2} e^{i(\omega-\Omega)t} \right] \quad (4-7)$$

so that, to first order, the phase modulated beam looks like three overlapped beams with three different frequencies. The two beams at frequencies $\omega \pm \Omega$ are referred to as RF sidebands. Calculating how this phase modulated beam interacts with an optical system is therefore as simple as calculating how all three beams interact individually with the system and summing them up.

As we saw in Chapter 3 an optical cavity has the ability to discriminate between different frequencies because they pick up a different amount of round trip phase when traversing the cavity. This therefore gives us a method for measuring the strength of the phase modulation of a laser beam. The one difficulty is that the quantity of interest is the total power in the sidebands versus the power in the carrier, but an optical cavity typically also separates out the different spatial modes of the carrier and sidebands. This can be avoided by using a confocal Fabry-Perot cavity which has end mirrors with equal radii of curvature and a separation between its two end mirrors equal to the radius of curvature of the mirrors. This configuration gives g parameters of $g_1 = g_2 = 0$ which allows all of the even modes to oscillate simultaneously and all of the odd modes to oscillate simultaneously with the odd resonances falling exactly halfway between the even resonances. Aligning a laser beam to such a cavity and scanning the length tells you the relative spectral content of the beam regardless of the spatial structure. Devices of this nature are referred to as optical spectrum analyzers (OSA).

An optical spectrum analyzer was used to measure the phase modulation index of the aLIGO EOM installed at the Livingston observatory. The results together with fits to equation 4-3 are shown in Figure 4-4. The data for the mode cleaner sideband and upper sideband do not extend far enough into the tails to constrain the fit of the Q value so those two values are likely inaccurate. Additionally, the drive strengths were not equal for the three measurements; the lower and upper interferometer sidebands were measured with a 36 dBm drive while the mode cleaner sideband was measured with a 24 dBm drive.

Figure 4-4 shows that modulation depths of 0.8 are attainable for the lower interferometer sideband, but that a factor of two increase in drive strength would be necessary for such a modulation depth on the upper interferometer sideband. Increasing the drive strength by a factor of two is likely too much RF power dissipation for the inductors to handle. Fortunately, sideband strengths of 0.1-0.3 have been found more

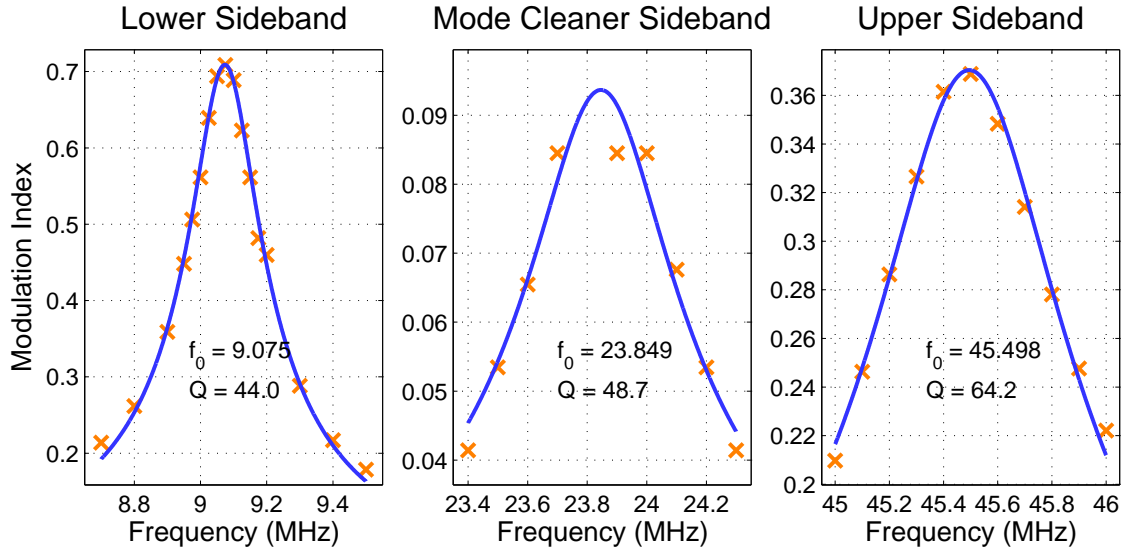


Figure 4-4: The measured phase modulation index of the LLO EOM as a function of frequency. Also shown are the fits to equation (4-3) with the fitted frequency, f_0 , and Q value. Note that the Q value of the 24 MHz and 45.5 MHz sidebands is probably overestimated since the data does not constrain the tails of the fit as well as the 9.1 MHz data. The drive strength for the lower and upper sideband measurements was 36 dBm while the drive strength for the mode cleaner sideband was 24 dBm.

than sufficient for interferometer locking so far at the LIGO Livingston Observatory. In fact, the mode cleaner sideband has been operating with a modulation index near 0.05 for almost two years now.

4.4 Residual Amplitude Modulation

Due to the previously described relationship between phase and amplitude sidebands, it is an unavoidable effect of electro-optic modulators that they produce some residual amplitude modulation, and the aLIGO EOMs are no different. We will see that there are many effects which couple some of the phase modulation to amplitude modulation but that the amount of residual amplitude modulation (RAM) induced by the aLIGO modulators is not expected to be high enough to be a limiting noise source for gravitational wave detection.

4.4.1 Causes and Effects of RAM in Advanced LIGO

The difference between a typical ideal phase modulator and amplitude modulator is that the amplitude modulator has a polarizer at its output and the polarization of the input beam is not aligned along one of the principal axes of the crystal. Since the input mode cleaner is an excellent polarizer, as we will see in Chapter 6, the only requirement for generating amplitude modulation is that the polarization of the input beam be misaligned with respect to the principal axes of the RTP crystal. Wedging the ends of the RTP crystal should avoid this effect by exploiting the natural birefringence of the crystal to separate the two polarizations regardless of the polarization state of the input beam. Wedging the ends of the crystal also has the advantage of preventing etalon effects which could otherwise produce low frequency amplitude modulation.

There are a number of less obvious ways to generate residual amplitude modulation in electro-optic phase modulators[76][95]. We will list some of these effects here but refer the reader to the literature for more in-depth study. Thermal effects in the electro-optic medium have been witnessed to have a strong effect on the amount of RAM, and the first line of defense in aLIGO if RAM is found to be too high will be to stabilize the temperature of the crystal. Other sources known to produce RAM are: scattering by impurities in the medium, spatial inhomogeneities in the electric field, and piezo-electrically induced vibrations of the electro-optic medium. Additionally, Sathien et al.[76] describe a RAM source caused by gradually evolving photorefractive scattering centers randomly distributed throughout the medium. They show that this effect scales with intensity, and that it can be reduced by introducing a second, more intense optical pump beam to erase the photorefractive scatter.

In addition to these effects in the electro-optic modulator, it is possible to convert phase modulation to amplitude modulation in the interferometer itself. We can see this by considering the first order expansion of a phase modulated and amplitude modulated beam. A phase modulated beam is shown in equation (4-7) while an

amplitude modulated beam is given by

$$E = E_0 e^{i\omega t} (1 + \gamma \cos \Omega t) = E_0 \left(e^{i\omega t} + \frac{\gamma}{2} e^{i(\omega+\Omega)t} + \frac{\gamma}{2} e^{i(\omega-\Omega)t} \right), \quad (4-8)$$

where γ is referred to as the amplitude modulation index. Comparing the two shows that the only difference is that with phase modulation the sidebands are out of phase while in amplitude modulation they are in phase. If the two sidebands pick up a differential phase shift or a differential loss somewhere inside of the interferometer, then the resulting beam can be expressed as the sum of a phase modulated and an amplitude modulated beam. I.E. any differential effect on the two sidebands converts phase modulation to amplitude modulation. One place where it is possible for this to happen is the input mode cleaner. If the carrier is perfectly on resonance in the input mode cleaner, then the symmetry of the sidebands about the carrier ensures that there will be no differential effects. Any offset in the locking point of the input mode cleaner will introduce differential losses in the two sidebands and convert phase modulation to amplitude modulation.

Kokeyama et al.[47] discuss the effects of RAM on an aLIGO-like interferometric gravitational wave detector. In particular, they show that RAM introduces offsets into the error points of the control loops which are used simultaneously to keep the interferometer at its operating point and extract the gravitational wave measurement. These offsets cause variations in the sensing matrix and degrade the precision noise subtraction schemes used in the multi-input, multi-output (MIMO) control system. In addition, they show that the offset in the locking point of the optical cavities leads to optical spring effects which further perturb the sensing matrix. In the end they conclude however that an AM/PM ratio of 10^{-4} puts the RAM induced noise below the design sensitivity of the aLIGO interferometers.

4.4.2 Measurements of RAM at the LIGO Livingston Observatory

The RAM level of the Livingston aLIGO EOM was measured during installation. Measuring RAM at RF frequencies is not a completely trivial task. It is important that

Table 4-2: The phase and amplitude modulation indices measured on the aLIGO EOM at the Livingston Observatory.

Sideband	9.1 MHz	24.0 MHz	45.5 MHz
AM Index	$3.9 \cdot 10^{-5}$	$1.0 \cdot 10^{-6}$	$6.2 \cdot 10^{-6}$
PM Index	0.39	0.085	0.15
AM/PM Ratio	$1.0 \cdot 10^{-4}$	$1.2 \cdot 10^{-5}$	$4.1 \cdot 10^{-5}$

the photodiode have a well known transfer function so that the RF power level can be compared to the DC power level. It is also important that the lensing solution to get the beam on the photodiode not be overly sensitive to pointing noise since inhomogeneities in the electric field of the crystal can show up as such. From equation 4-8 we can see that the amplitude modulation index is given by

$$\gamma = \frac{P_{\Omega}}{2P_{DC}}, \quad (4-9)$$

where P_{Ω} is the power measured at the modulation frequency and P_{DC} is the total power of the beam. The values measured for each of the three sidebands are shown in Table 4-2.

The measurements shown in Table 4-2 were taken by putting the entire beam onto an RFPD, giving an integrated measurement of the amplitude modulation index. The spatial uniformity of the amplitude modulation index across the cross section of the beam was also characterized. This was done by using a telescope which simultaneously made the beam significantly larger than the photodiode and placed the PD at an image plane of the EOM so that any induced pointing noise at the EOM did not show up as displacement at the PD. The PD was mounted on a translation stage with micrometers attached so that it could be translated across the beam profile.

The results for the 9.1 MHz sideband are shown in Figure 4-5. We can see that the measured amplitude modulation index varies over the cross-section of the beam by almost a full order of magnitude. It is possible that the phase modulation index varies

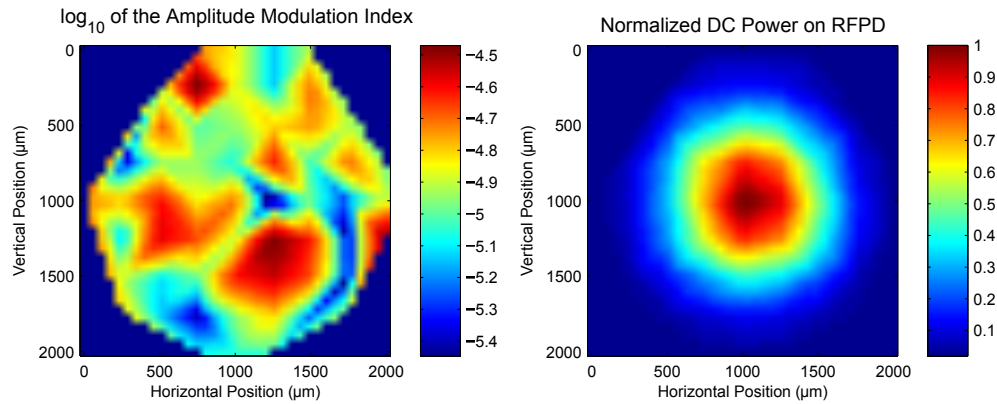


Figure 4-5: The spatial variation of the amplitude modulation index of the aLIGO EOM. This was measured by using a telescope to simultaneously make the beam large on the photodiode and place the photodiode in an image plane of the EOM to reduce pointing noise contamination. The PD was mounted on a translation stage controlled by micrometers, and the AM index measurement was made on a $25 \mu\text{m}$ grid. The left hand plot shows the measured amplitude modulation index and the right hand plot shows the normalized power on the photodiode to give the scale of the beam size.

in the same way, but this effect was not measured. Some possible explanations are inhomogeneities in the crystal and/or in the applied electric field or slowly evolving photorefractive scattering centers as proposed by Sathien et al.[76]. The source of this effect was not fully understood, but the ramifications are minimal for two reasons. First, the input mode cleaner (discussed in Chapter 6 will significantly filter the spatial non-uniformity of the RF sidebands. Additionally, measurements of the sideband beat notes in the interferometer integrate large portions of the beam so that any residual non-uniformity is heavily averaged.

In addition to studying the spatial homogeneity of the RAM from the EOM, the low frequency time variation was also studied. For this study, a spectrum analyzer was used to record the RF notes of the 9.1 and 45.5 MHz sidebands roughly once per minute for close to 48 hours. The results, shown in Figure 4-6 together with two temperature sensors inside of the PSL enclosure, indicate that the dominant factor in the drift of the RAM is temperature fluctuations of the enclosure. The green line in the lower plot of that figure shows the feedback to the temperature loop of the pre-modecleaner (PMC). This

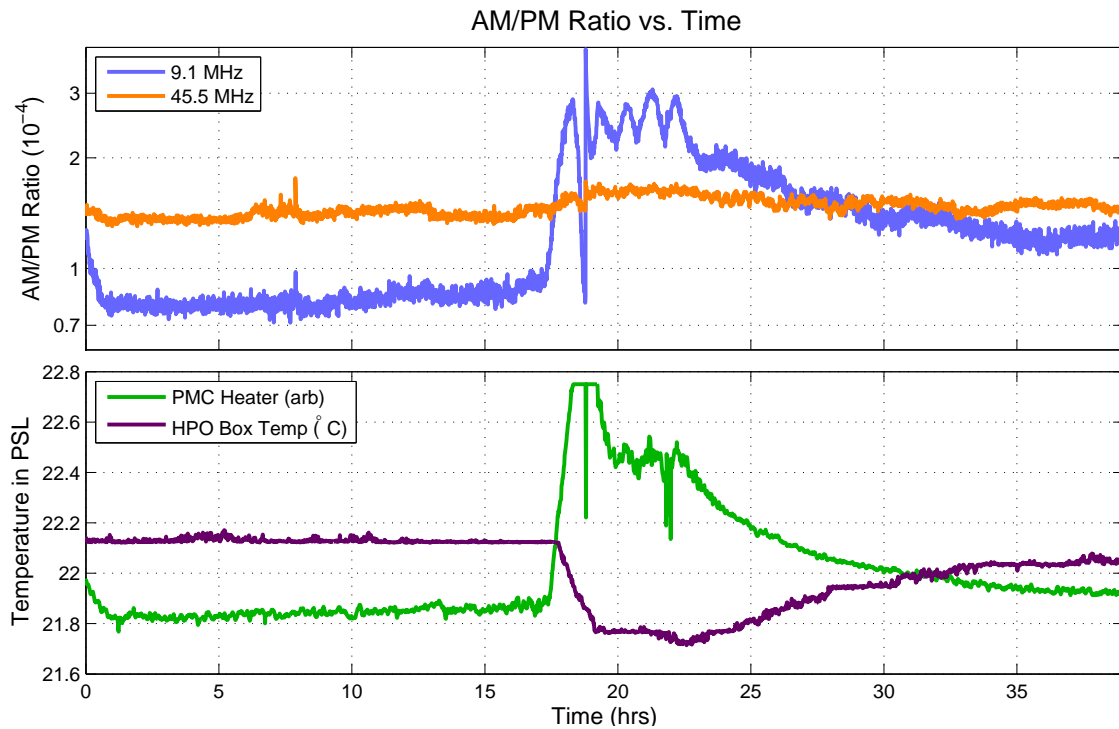


Figure 4-6: A long term study of the RAM directly after the EOM. During the study the air conditioner in the PSL was switched on between hours 18-23. The lower plot shows two different temperature sensors inside of the PSL, the HPO Box Temp is a temperature sensor inside of the box for the high power oscillator and is lowpassed by the enclosure. The PMC Heater is the feedback signal from the pre-mode cleaner to the heater; as the temperature in the room drops the heater had to apply more power to keep the cavity length constant. It is included because it has a faster response time than the HPO Box Temp sensor. The PMC lost lock for a short time around 19 hours because the heater was railed and this shows up as a spike in the data.

signal shows how much power is being applied to the heater on the aluminum body of the PMC in order to keep its length fixed. It acts in essence as an accurate readout of the inverse of the temperature of the PSL enclosure at frequencies below a few hundred millihertz.

4.5 Conclusion

In conclusion, the Advanced LIGO electro-optic modulator is ready for science mode operation of the interferometers. There are two outstanding issues which may need to be addressed in the future. Firstly, the modulation index is not quite able to reach the initially required level of $\beta = 0.8$. This has not been found to be a problem during

initial commissioning of the interferometers during which the modulation index was typically reduced to $\beta \sim 0.1$. If it is found in the future that the required modulation index is larger, then the easy fix will be to increase the power handling ability of the resonant circuitry and increase the drive strength.

The other issue which may be a problem in the future is temperature related drift of the residual amplitude modulation of the EOM which, as shown in Figure 4-6, can drift above the required 10^{-4} level which was found to be sufficient by Kokeyama et al.[47]. Addressing this issue requires a fairly straightforward temperature stabilization servo to be attached to the EOM. This method has been shown[50] to be effective in maintaining low RAM levels from EOMs.

CHAPTER 5 THE ADVANCED LIGO FARADAY ISOLATOR

In this chapter we will look at one of the other major components of the input optics; the in-vacuum Faraday isolator (FI). As with the electro-optic modulators in the previous chapter we will start by looking at the requirements placed on the Faraday isolator and how it was designed to meet these requirements. After that we will look at the testing results of the first aLIGO FI which was built, tested, and installed at the LIGO Livingston Observatory (LLO).

5.1 Requirements

Building an in-vacuum Faraday isolator which is capable of dealing with high powers is not an easy task, so it is reasonable to question why one is necessary at all. The first reason is that the length between the IMC and the main interferometer is uncontrolled so any back-scattering of the IMC creates an uncontrolled, low finesse Fabry Perot cavity. This effect causes the power incident at the input of the interferometer to be modulated at low frequencies, an effect which was witnessed in the initial commissioning of initial VIRGO[3] before an in-vacuum FI was installed. The power modulation observed was at low enough frequencies to not spoil the sensitivity of the instrument, but it would cause spurious kicks in the IMC control loops which were responsible for most of the downtime during initial commissioning. In addition to this power modulation effect, the light reflected off of the PRM drops significantly when the interferometer acquires lock which will cause a radiation pressure kick to the already locked IMC. Rodica Martin, who played a large role in the design of the aLIGO FI, explored the need for an in-vacuum FI in the design stages of aLIGO[55] and found that these effects were likely not to be a problem. The other advantage of having an in-vacuum FI for aLIGO, the one which tips the scales in favor of keeping it, is that the reflected light from the interferometer is picked off while still on the seismically isolated in-vacuum tables. Since aLIGO employs

Table 5-1: The requirements of the aLIGO Faraday isolator.

Parameter	Goal	Comment
Optical Throughput (%)	> 95%	Limited by residual surface reflection (16 surfaces)
Isolation Ratio (dB)	> 30 dB	Limited by extinction ratio of calcite wedge polarizers
Thermal Lens Power (m^{-1})	< 0.02	Leads to a 2% reduction in mode matching
Thermal Beam Drift (μrad)	< 100	Based on dynamic range of tip-tilt actuators

Source: Tanner et al. "Input Optics Subsystem Design Requirements Document." LIGO DCC T020020 (2009)[89] and UF LIGO Group "Upgrading the Input Optics for High Power Operation." LIGO DCC T060267 (2006)[93]

specially designed low noise in-vacuum photodetectors it is important that the reflected light remain in-vacuum before detection.

The key requirements on the aLIGO FI are listed in Table 5-1. In particular, the requirement on isolation is placed at 30 dB based on power scaling arguments from parasitic interferometric effects seen in initial LIGO. The optical throughput is required to be greater than 95%. The thermal lens power should be such that it causes less than a 2% drop in mode matching to the interferometer at full power. Finally, the thermal beam steering should be less than 100 μrad .

5.2 Design

The Faraday effect is a well known optical effect in solids first discovered and explored by Michael Faraday in 1845[77]. It is understood now to arise from the motion of the electrons in the material which are pumped by the optical electrical field and curved by the magnetic field. At first order the effect of submerging an optical material in a magnetic field is to show different indices of refraction to the two different circular polarizations. To a linear polarized field this appears as a rotation of angle θ given by[75][29]

$$\theta = VB\ell, \quad (5-1)$$

where B is the magnetic flux density in the direction of propagation of the wave, ℓ is the length of the crystal over which the light is immersed in the magnetic field, and V is the Verdet constant which contains the material parameters. The Verdet constant for the Northrup Grumann produced terbium gallium garnet (TGG) crystals used in the aLIGO Faraday isolators is $V \simeq -40 \frac{\text{rad}}{\text{T m}}$ at 1064 nm.

The advantage of the Faraday effect, as discussed in Section 3.3 is that it is non-reciprocal. This means that, unlike other optical elements with optical activity, when a wave travels through a Faraday rotator, reflects back onto itself, and travels back through the Faraday rotator it undergoes twice the rotation of the single pass instead of zero net rotation. A Faraday isolator exploits this effect by choosing magnetic field strength and crystal length such that the optical field picks up 45° of rotation inside of the Faraday crystal. By adding a half wave plate (HWP) which rotates it 45° more, then the polarization is switched on one direction through the assembly and stays the same when it passes through in the other direction.

The limiting factors for the amount of isolation attainable are the homogeneity of the rotation in the Faraday rotator and the extinction ratio of the polarizers at the input and output. Barring thermal effects in the Faraday rotator, which we will discuss post haste, the homogeneity is generally not the limiting factor for isolation of a Faraday isolator since the beam size is typically small compared to the magnet bore size. The limiting factor for isolation ratio is therefore the extinction ratio of the input and output polarizers. It is for this reason that the aLIGO FI uses calcite wedge polarizers (CWP) which have an extinction ratio on the order of 10^5 . This comes at the added cost of difficulty in picking up the reflected beam which is only deflected by ~ 10 mrad[93].

The enhanced LIGO (eLIGO) design was optically identical[23] to the aLIGO design except that it included an intermediate thin film polarizer (TFP) to facilitate picking up the interferometer reflected beam. This was dropped in aLIGO because it was found to be the dominant source of loss in the FI, throwing away a few percent of the power

in the forward pass. One problem caused by removal of the TFP is that any incorrect polarization at the input of the FI is dumped inside of the Faraday magnet because of the small deflection angle of the CWPs. One non-optical change was also made to the aLIGO FI; the HWP is now mounted on a rotation stage which can be controlled from outside of the vacuum system. This allows one to optimize the isolation ratio in situ for different powers.

The rest of the design of the FI is in place to try and compensate thermal effects caused by the high optical power planned for the aLIGO interferometers. At high optical powers the amount of rotation through a magneto-optical element, such as the terbium gallium garnet crystals used in the aLIGO FI, becomes non uniform. This is caused by two effects; the temperature dependence of the Verdet constant and the photoelastic effect caused by thermal strains. Khazanov et al.[44] found that the thermal depolarization is dominated by the photoelastic effect and subsequently proposed and investigated some methods of compensation[86]. The aLIGO FI employs one of these methods in which two separate Faraday rotators with 22.5° of rotation are used with an intermediate quartz rotator. The two Faraday rotators are oriented identically but the polarization of the beam going through the two crystals is rotated such that the thermal birefringence of one crystal is largely compensated by the other.

In addition to thermal birefringence in the magneto-optical crystals thermal lensing due to a non-negligible change in the index of refraction of most materials with temperature, dn/dT , causes the optical crystals of the Faraday isolator to act as a power dependent lens[43] (a phenomenon known as thermal lensing). This effect is compensated in the aLIGO FI by adding an optical element whose dn/dT coefficient has the opposite sign of the TGG crystals. In particular, a piece of deuterated potassium dihydrogen phosphate (DKDP) is added. The thickness of this crystal must be matched to each FI individually since the total amount of thermal lensing in the rest of the crystals

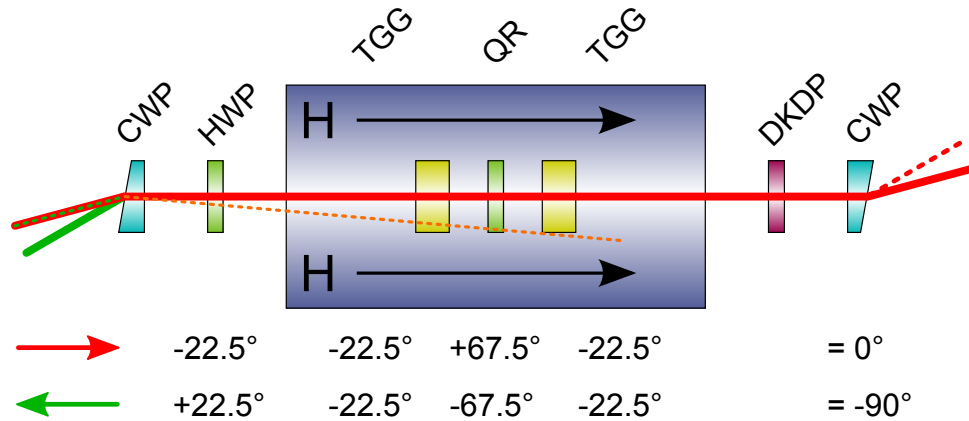


Figure 5-1: A schematic of the aLIGO Faraday isolator. The components are: two calcite wedge polarizers (CWP), a half wave plate (HWP), two terbium gallium garnet crystals (TGG), a quartz rotator (QR), and a deuterated potassium dihydrogen phosphate crystal (DKDP). The TGG crystals have a high Verdet constant and are used as the Faraday rotation medium which is why they are submerged in the magnetic field. The QR is used to flip the polarization of the beam between the two TGG crystals for thermal depolarization compensation. The DKDP is a negative dn/dT element used to compensate the thermal lensing of the rest of the optics in the FI.

is difficult to determine ahead of time. For the LLO FI the installed crystal thickness is 3.5 mm.

A schematic view of the aLIGO FI is shown in Figure 5-1 which includes the total theoretical polarization rotation in both directions through the FI. The optical elements and rotation in the forward direction are: CWP(0°) \rightarrow HWP(-22.5°) \rightarrow TGG(-22.5°) \rightarrow QR($+67.5^\circ$) \rightarrow TGG(-22.5°) \rightarrow DKDP (0°) \rightarrow CWP(0°) = 0° . In the backward direction we have: CWP(0°) \rightarrow DKDP (0°) \rightarrow TGG(-22.5°) \rightarrow QR(-67.5°) \rightarrow TGG(-22.5°) \rightarrow HWP($+22.5^\circ$) \rightarrow CWP(0°) = -90° .

In addition to the optical design of the aLIGO FI, the magnet was carefully researched during the design phase.[67][83][100] The magnet is made up of seven magnetic annuli composed eight wedged Nd-Fe-B permanent magnets. Six of the rings have radially directed fields at the center, three pointed towards and three pointed away from the center of the assembly. The center magnetic annulus is one solid ring with the magnetization vector directed along the optical axis. The axial and radial repulsive

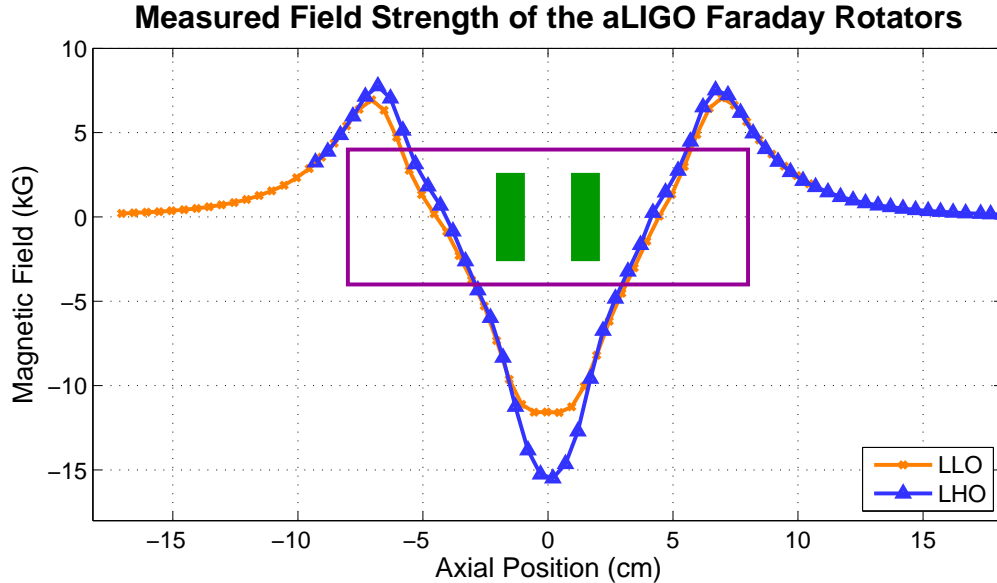


Figure 5-2: The measured magnetic field strength of the LLO and LHO Faraday rotators. The green boxes represent the length and approximate location of the TGG crystals and the magenta box represents the approximate length of the magnet assembly. The vertical dimension is meaningless in both cases.

forces are substantial; 500 MPa in the axial direction and 250 MPa in the radial direction. For this reason all of the magnets are mounted inside of a solid aluminum cup with 9 mm wall thickness.

The axial field strength of the Livingston and Hanford Faraday rotator magnets is shown in Figure 5-2. The use of modern ND-Fe-B magnets arranged in alternating polarities allows the permanent magnet of the Faraday rotator to reach very high field strengths, on the order of 1 Tesla (10 kG). The field of the composite magnet is also specifically designed to have a large axial gradient, while being uniform radially, so that the amount of Faraday rotation obtained through the TGG can be finely tuned by sliding the crystals towards or away from the center of the magnet.

5.3 Measured Optical Losses

In the following sections we will show the results of as-built testing of the aLIGO Faraday isolator at the LIGO Livingston Observatory (LLO). The results for the Hanford observatory are similar but will not be reported here. In particular, we will start by

Table 5-2: The possible optimizations of the LLO aLIGO FI at 20 W.

Optimized For	R_s (W)	R_p (mW)	T_p (W)	T_s (mW)	Isol. (dB)	Trans. Rej. (dB)
Isolation	19.7	0.550	20.2	15.6	45.5	31.1
Trans. Rejection	19.9	12.7	20.2	0.735	32.0	44.4

Note: The meaning of the headers is: R_s and R_p are the s and p polarizations in the backward direction and T_s and T_p are the s and p polarization in the forward direction. The HWP in the Faraday was optimized in each case to minimize the amount of power dumped in either the transmission (forward) or isolation (backward) direction. These numbers tell one about the non-uniformity of the magnetic field in the Faraday isolator as well as the amount of uniform rotation not equal to 45° .

looking at the losses inside of the Faraday isolator since this is one of the dominant loss points in the input optics. Afterwards we will look at the isolation ratio measurements both in and out of the vacuum system. Finally, we will see the results of thermal lensing measurements.

The losses in the eLIGO Faraday isolator totaled to $\sim 7\%$ with the dominant source of loss being the thin film polarizer which contributed 4% to the total[22]. This was already an improvement over the iLIGO Faraday isolator which suffered from 14% losses. As previously mentioned, the only major change to the Faraday isolator between eLIGO and aLIGO was the removal of the thin film polarizer. The measured aLIGO losses were roughly as expected with the removal of the TFP at $3.3 \pm 0.5\%$.

These losses are dominated by residual reflectivity of the surfaces within the Faraday isolator. This is because anti-reflective (AR) coatings are not as good for the exotic optical materials as they are for the fused silica used for the interferometer mirrors. In addition, the sheer number of surfaces (a total of 14) makes it difficult to obtain a total loss significantly lower than was obtained. For optical squeezing[90] of the aLIGO interferometers, low loss Faraday isolators are essential for insertion of the squeezed beam, but it is not important that these FI be able to handle high power. If we make the naive assumption that the loss is evenly distributed across all surfaces and

assume that a similar low power isolator will only have 8 surfaces, then we can estimate the losses for the low power isolator to be 1.9 %.

In addition to the losses from components in the FI, there is another source of loss which is negligible in scale but interesting in its own right. Namely, any non-uniformity in the rotation of the Faraday rotator or any amount of uniform rotation not equal to 45° will lead to some power being dumped at the output polarizer when the Faraday is optimized for isolation. The possible optimizations at 20 W of input power (single pass) are shown in Table 5-2. From the transmission data measured when optimized for minimum transmission rejection we can compute that the inhomogeneities in the magnetic field are at the level of 36 ppm. From the transmission data measured when optimized for isolation we can compute that the amount of missing uniform rotation is 1.6° .

5.4 Measured Thermal Lensing

As discussed above, the Advanced LIGO Faraday isolator (FI) employs a scheme to compensate for self induced thermal lensing inside of the Faraday rotator. The scheme employed, investigated by Khazanov et al.[43], uses a piece of deuterated potassium dihydrogen phosphate (DKDP) whose thickness is matched to the thermal lensing in the rest of the FI. DKDP is able to compensate for the thermal lenses in the rest of the FI because DKDP has a shift in its index of refraction with respect to temperature (dn/dT) which is negative whereas the other elements in the FI, indeed most optical materials, have a positive dn/dT . So as the FI heats up due to absorption in its optical elements, the negative lens in the DKDP compensates for the positive lenses in the rest of the FI.

The expected thermal lens developed in one of the TGG crystals is shown in Figure 5-3. This figure shows the thermal lens inside one of the TGG crystals calculated from the Hello Vinet [36] formalism discussed in Section 3.4. It also shows the parabolic approximation (approximation as a single focal length) to the thermal lens as well as the beam profile inside of the TGG. The key point of this figure is to show that the self

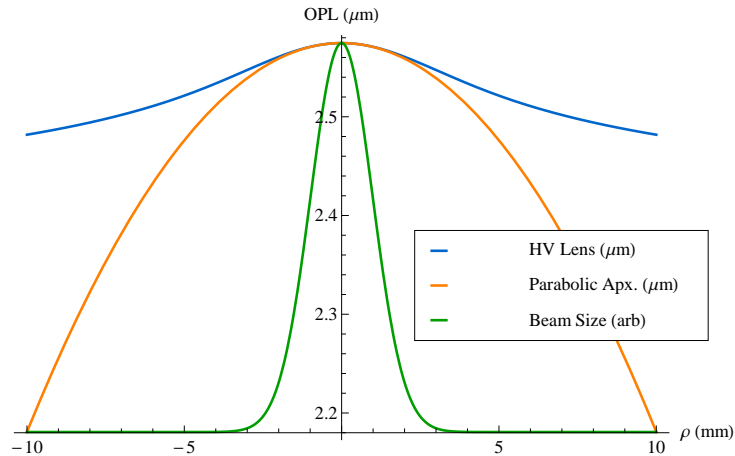


Figure 5-3: The theoretical self induced thermal lens in one of the TGG crystals in the FI at the full 165 W of input power. The profile of the thermal lens is calculated with the Hello Vinet[36] formalism discussed in Section 3.4. Together with the HV thermal lens is shown the parabolic approximation which has a focal length of $f = +126$ m. Also shown is the beam profile in arbitrary units. Notice that the parabolic approximation is fairly good in this case as it stays close to the actual phase profile over the full width of the beam. This is because of the high thermal conductivity of TGG.

induced thermal lens caused by optical absorption in materials is not truly parabolic, and any talk, therefore, of the focal length of a thermal lens neglects this fact. One of the advantages of using a negative dn/dT material to compensate a thermal lens is that to first order, the full thermal profile of the material is compensated rather than just the parabolic component as long as the two crystals are close together with respect to the Rayleigh range of the beam.

Measuring the thermal lensing of the aLIGO FI at LLO was accomplished by using a high power beamsan. The measurements were taken in the PSL enclosure using the full high power beam of the PSL with a telescope which was designed to reproduce the expected mode inside of the FI in-vacuum. Thermal lensing measurements can be a bit difficult because all of the optics in the path before and after the FI will also lens. For this reason a reference measurement was taken with the full optical path but with the FI removed. All of the thermal lensing data was then compared to this reference measurement to extract the focal power of the FI. The full report for testing of the aLIGO

Table 5-3: The measured focal length of the LLO FI with the three different thicknesses of DKDP that were tried.

DKDP Thickness	Foc. Length @ Meas. Power	Foc. Length @ 120 W
4.5 mm	-23.5 m @ 62 W	-12.1 m
4.0 mm	-50.0 m @ 65 W	-27.1 m
3.5 mm	-31.0 m @ 120 W	-31.0 m

For technical reasons the full power of the PSL laser was not available for two of the measurements. The installed item at LLO uses the 3.5 mm DKDP crystal.

FI can be found in [64], but the final results will be reported here. The measured focal powers of the LLO FI for three different thicknesses of DKDP are shown in Table 5-3. The LLO FI was installed into the vacuum system with the 3.5 mm DKDP crystal. The focal power of the FI with this crystal at 120 W is $\sim -0.032 \text{ m}^{-1}$ which is just slightly outside of the requirements.

A set of example beamscans together with their fits are shown in Figure 5-5. Note that, in order to get the beam to go through a waist, a lens of focal length $f = +1146 \text{ mm}$ was placed in the beam after the FI. ABCD methods, described in Section 2.2, were then used to infer the focal length of the FI. The blue data are the reference data taken with the Faraday isolator removed from the beam path, and the red data was taken with the 3.5 mm DKDP crystal at 120 W of power. Notice that the red data does not match as well with the fit as the blue data; this is likely evidence of higher order thermal effects caused by the FI.

The data from Table 5-3 are shown graphically in Figure 5-4. The green horizontal line in that Figure shows the inverse of the measured focal power of the Faraday isolator without the DKDP in place. The orange circles show the measured focal lengths of the various thicknesses of DKDP. Note that the DKDP focal lengths were not measured independently; they were inferred from the measurements of the total focal length of the FI and the measurement of the FI without the DKDP. The blue line shows the results of a finite element model calculated with the commercial software multi-physics

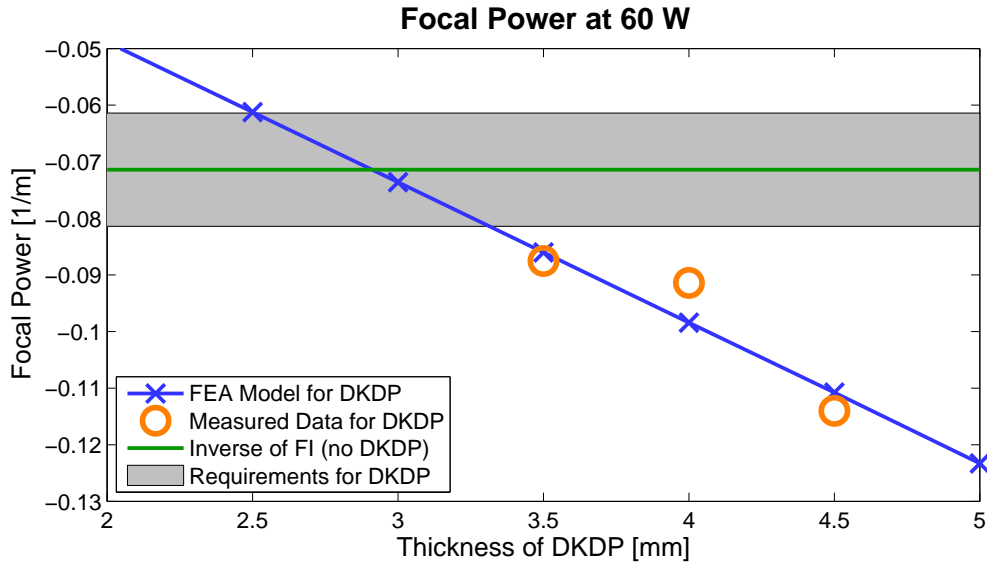


Figure 5-4: The measured focal powers for the FI without the DKDP and the different thicknesses of DKDP. Also shown is a finite element model of the thermal lensing in the DKDP as well as the region in which the DKDP focal length must fall to be within the requirements. Note that all numbers are scaled to 60 W of input power. The LLO FI was installed with the 3.5 mm DKDP crystal, but this plot indicates that a 3 mm crystal would have been a better choice. If thermal lensing in the FI is found to be a problem in the future, the DKDP crystal in the FI should be replaced.

package COMSOL. Finally, the grey shaded region shows the region in which the DKDP focal length must fall to fit within the requirements. This plot indicates that the proper choice of DKDP thickness for the LLO FI was probably 3 mm, rather than the 3.5 mm crystal which was installed. The 3.5 mm crystal which leaves the FI just outside of the requirements was installed because the lead time on ordering a 3.0 mm crystal would have pushed back the commissioning schedule by at least one month. If thermal lensing in the FI is found to be an issue during high power operation of the Advanced LIGO detector, then the crystals can be swapped fairly easily.

5.5 Measured Isolation Ratio

In this section we will present measurements of the key property of the Faraday isolator, its isolation ratio. For technical reasons the full power of the PSL laser was not available during these isolation ratio measurements so the highest measured powers

Thermal Lensing 3.5 mm DKDP @ 120 W Vertical Axis Blue=Cold Red=Hot
 Beam Radius (μm)




Figure 5-5: A set of example beamscaans together with the fits to the data. The blue data is the reference data taken at high power with the FI removed, and the red data is the measurement taken at 120 W of power with the 3.5 mm DKDP crystal. The focal lengths quoted in this section are determined by using the beam parameters extracted from beamscaans such as this one together with an ABCD model to determine the focal length of the FI. Notice that the hot data shows some signs of higher higher order thermal effects.

are at 70 W. The isolation ratio is typically quoted in units of dB and is defined by

$$I = -10 \log_{10} \left(\frac{P_{leak}}{P_{ref}} \right), \quad (5-2)$$

where P_{leak} is the power that leaks back along the same axis as the main beam and P_{ref} is the power which is sent along the reflected axis.

As discussed above, the amount of optical rotation in the Faraday rotator is not perfect. Hence, some amount of power is dumped at the output polarizer in transmission. In analogy with the isolation ratio we adopt a similar ratio for quantification of this trait of the FI, the transmission rejection ratio. It is defined by

$$TR = -10 \log_{10} \left(\frac{P_{rej}}{P_{trans}} \right), \quad (5-3)$$

where P_{rej} is the amount of power rejected into the wrong polarization and P_{trans} is the amount of power transmitted in the proper polarization.

The measured in-air isolation and transmission rejection ratios are shown in Figure 5-6. At the time of measurement the Faraday isolator was optimized for isolation at 20 W of input power since this is the expected operational power of the initial science runs. The figure shows that although the isolation ratio does start to roll off at higher power, it is within the requirements at all measured powers in air. One difficulty with going to vacuum is that the lack of convective cooling makes the thermal effects in the Faraday isolator stronger so it is expected that the high power roll off in isolation ratio will be somewhat stronger. Unfortunately, we have not yet been able to measure this effect because the full power of the PSL was not available to be injected into the vacuum system for technical reasons.

The previously mentioned addition to the aLIGO FI of a rotation stage for the HWP which can be controlled from outside vacuum system allows one to reoptimize the isolation ratio at different operating powers. So, in principle, the isolation ratio can be improved over that shown in Figure 5-6 at higher powers. This functionality has not yet been tested since the aLIGO FIs have not yet seen high operating powers in-vacuum.

The low power isolation ratio was measured in vacuum after the installation of the LLO FI. This was accomplished by placing matched beamsplitters angled at $\pm 45^\circ$ in the input beam on the PSL table. The matched beamsplitters allows for insertion without changing the alignment of the beam being injected into the interferometer; only one of the beamsplitters is needed for the measurement. The downstream beamsplitter was calibrated with a power meter for its specific angle of incidence, and the power going into the beamsplitter was increased until 3W was transmitted through towards the interferometer. By measuring the reflected power off of this beamsplitter and accounting for all known losses the in-vacuum, low power isolation ratio was measured to be 29 dB. The details and results are documented in LLO alog 8608[63].

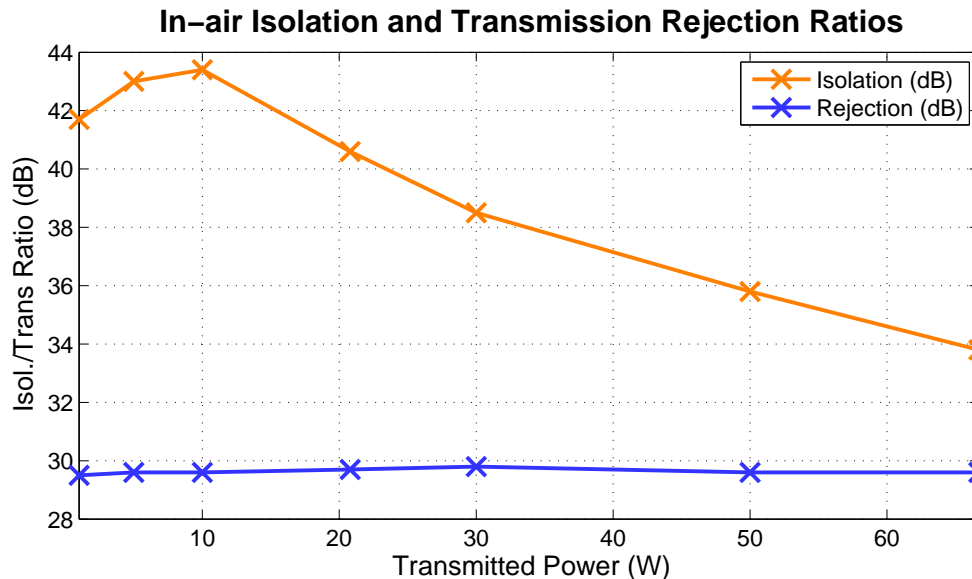


Figure 5-6: The measured isolation and transmission rejection ratios as a function of power. The isolation ratio is defined in equation (5-2) and the transmission rejection ratio in (5-3) respectively. These measurements were made in air, and at the time the Faraday isolator was optimized at 20 W of input power.

5.6 Conclusion

Overall, the in-vacuum Faraday isolator is performing at a level which sufficient for gravitational wave detection at low input powers. Out of vacuum testing of the high power performance indicates that the Faraday isolator is within the requirements except for the marginally unacceptable thermal lensing. Since the major drawback of thermal lensing in the Faraday isolator is the decrease in the coupling of light into the interferometer, this level may still be acceptable for high power operation of the aLIGO interferometers. If it is found to be unacceptable, then the fix is to simply swap the DKDP crystal for a slightly thinner version; an operation which does not require significant downtime other than the time needed to vent and evacuate the vacuum system.

CHAPTER 6 THE ADVANCED LIGO INPUT MODE CLEANER

The input mode cleaner (IMC) is the heart of the input optics, serving simultaneously as a spatial filter, polarization filter, frequency reference, and pointing reference. It is an in-vacuum, suspended, three mirror cavity with mirrors hanging from the LIGO small triple suspensions[73]. It has a free spectral range of 9.099 MHz and a finesse of 515. The beam is injected along one of the long arms and extracted along the other (see Figure 6-1). The reflected beam is outfitted with an RF photodiode for Pound-Drever-Hall length sensing [24] and two differential wavefront sensors for angular sensing. In addition, a pickoff of the intra-cavity light is extracted behind the curved mirror, MC2, and sent to a quadrant photodiode for additional angular information.

In this chapter, as with the last two, we will begin by discussing the requirements of the input mode cleaner (IMC) and how it was designed to meet those requirements. Afterwards we will describe results of the commissioning of the input mode cleaner at the LIGO Livingston Observatory, where the first IMC was installed and tested. In particular, we will discuss how the IMC is kept at the operating point by using active control in both the length and angular degrees of freedom. We will examine the length noise of the IMC to try and understand how well it is meeting the requirements. Finally, we will look at the losses and thermal effects in the LLO IMC and discuss the importance of these effects to the full interferometer.

6.1 Requirements

The design requirements for the input mode cleaner derive from the designed sensitivity of the full Advanced LIGO interferometer. Figure 6-2 shows the design sensitivity of the Advanced LIGO interferometers together with the measured sensitivity of one of the Initial LIGO interferometers during the sixth science run. The Advanced LIGO interferometers are designed to be limited by seismic noise below ~ 10 Hz, radiation pressure noise in the 10-50 Hz range, coating thermal noise between 50




Figure 6-1: A schematic layout of the input optics. The input mode cleaner is represented by the three mirrors MC1, MC2, and MC3. It is an in vacuum, triply suspended, triangular cavity used as a frequency and pointing reference. It also passively filters the spatial structure and polarization of the transmitted beam.

Hz and 200 Hz, and photon shot noise above 200 Hz. The mantra for setting design requirements for the interferometer subsystems is therefore to keep the 'technical noise sources' below this limit. In fact the subsystem requirements were set by demanding that the technical noise sources be a factor of 10 below the design sensitivity[32].

6.1.1 Frequency Noise at the Interferometer Input

A perfect Michelson interferometer is completely insensitive to frequency noise; a fact which is a large motivator in choosing the Michelson topology for gravitational wave interferometers. However any asymmetry between the two arms of the Michelson interferometer will couple frequency fluctuations to the dark port. The asymmetries between the two arms depend on many parameters such as: the imbalance between the reflectivity and transmissivity of the beam splitter, the reflectivity imbalance of the arm cavity input mirrors, different amounts of losses in the two arm cavities, and the static

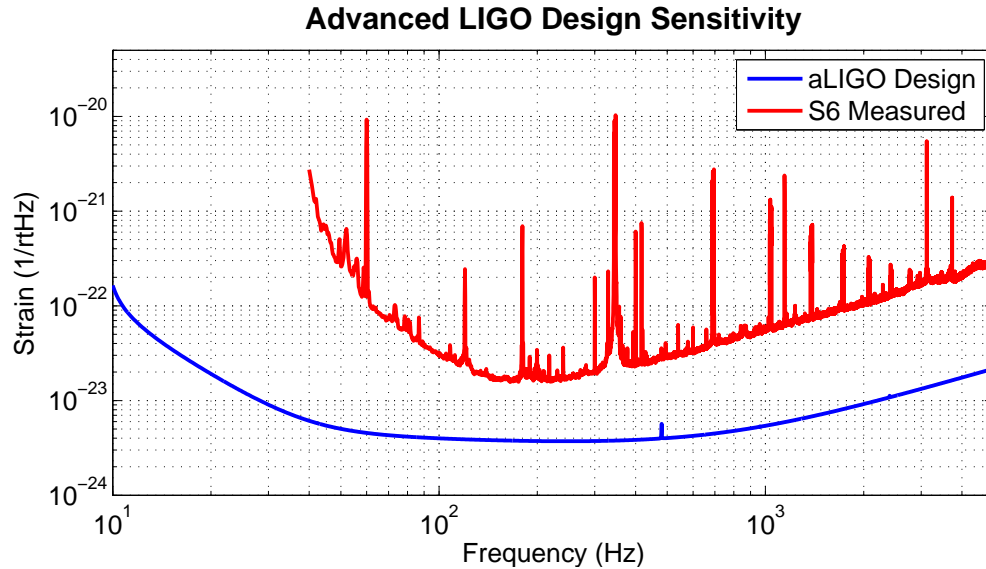


Figure 6-2: The Advanced LIGO design sensitivity[84] is shown together with the measured sensitivity from the LIGO Hanford Observatory during the sixth LIGO science run[69]. The design sensitivity includes the fundamental limiting noise sources (quantum, thermal, and seismic) which set the requirements for the interferometer subsystems such as the input optics.

distance asymmetry between the beam splitter and the two arm cavities (known as the Schnupp asymmetry).

Using experience from Initial LIGO and a reasonable estimate of the Advanced LIGO parameters, the Interferometer Sensing and Control (ISC) group put together a detailed numerical simulation of the full Advanced LIGO interferometer [32]. With this simulation they derive the frequency noise requirement at the input to the interferometer to be roughly the solid red curve shown in Figure 6-3. The frequency noise stabilization for the interferometer is not solely the job of the input optics; the average length of the two arms acts as another, more stable frequency reference. Figure 6-3 also shows the frequency noise requirements before suppression by the common arm length (CARM) feedback servo. Also shown are the modeled length noise of the input mode cleaner as well as the expected frequency noise out of the pre-stabilized laser (PSL).

The job of the input mode cleaner is therefore to stabilize the laser frequency by roughly a factor of 20 between 10 Hz and 10 kHz. In reality the frequency noise

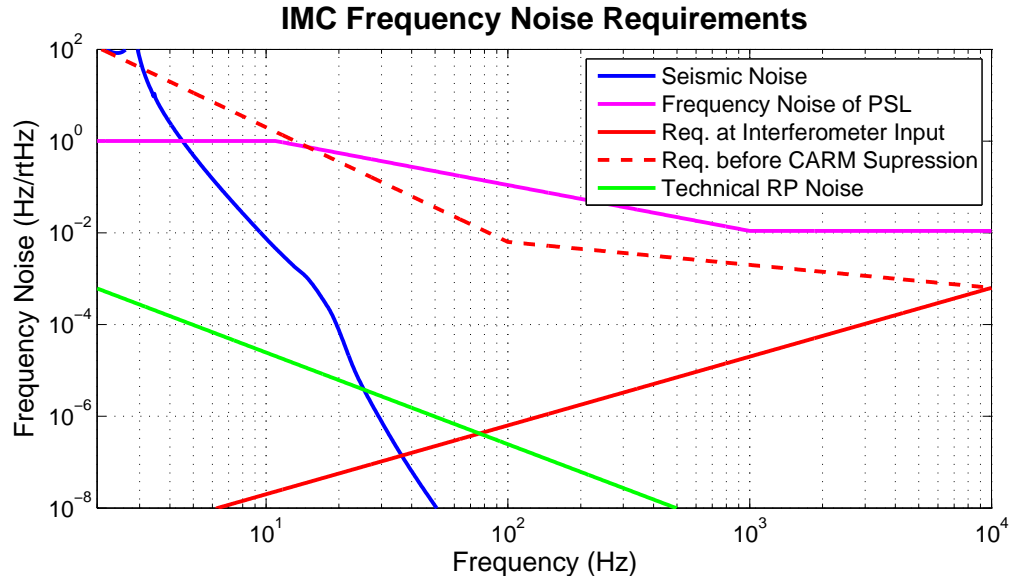


Figure 6-3: The frequency noise requirements at the input to the interferometer both with and without the common arm loop suppression. Also shown are the expected length noise of the input mode cleaner from seismic motion and from radiation pressure noise as well as the expected frequency noise out of the pre-stabilized laser (PSL). Ostensibly the job of the input mode cleaner is to be able to suppress the frequency noise out of the PSL to the requirements level after the common arm gain is taken into account. In reality vibrations in the injection optics between the PSL and the in-vacuum input optics will require a higher level of suppression.

suppression necessary for the input mode cleaner is higher because vibrations in the injection chain will add frequency noise before the light is injected to the isolated in-vacuum optics as we will see when we discuss the noise budget below.

6.1.2 Beam Alignment Noise at the Interferometer Input

As with frequency noise, the contrast of a perfect Michelson interferometer is first order insensitive to pointing noise. There are two types of imperfections which couple pointing noise to the gravitational wave readout channel. The first, considered in [89], is caused by residual misalignments of the core interferometer optics causing the misaligned beam to rejoin with the main interferometer mode and show up in the gravitational wave channel. The second coupling is caused by the beam spot motion on the core interferometer optics causing the residual angular motion of the optics to couple to the gravitational wave channel. These two effects lead to similar noise requirements

at the input to the interferometer; a beam stability of $4 \cdot 10^{-10} \frac{\text{rad}}{\sqrt{\text{Hz}}}$ at 100 Hz and above and decreasing (getting less stringent) as f^2 below 100 Hz.

6.2 Design

The ability of the input mode cleaner to filter is due to the ability of an optical resonator to distinguish very precisely the round trip phase as we saw in Section 3.1. This fact, coupled with the fact that the different higher order modes of a laser beam pick up a different amount of Gouy phase (see Section 2.1.2) as they propagate, gives rise to the ability of the input mode cleaner to selectively transmit only the fundamental Gaussian mode with the correct polarization.

To see how this works, let's calculate the transmission and reflection coefficients of the input mode cleaner in terms of the Hermite-Gauss modes. We will use the subscripts 1, 2, and 3 to denote the properties of MC1, MC2, and MC3 respectively so that r_1 and t_1 are the amplitude reflectivity and transmissivity of MC1. In addition, we will denote the phase picked up when traversing between two mirrors as e.g. ϕ_{12} for the phase picked up between MC1 and MC2. In this notation, the coefficients are given by

$$r_{IMC} = r_1 + \frac{r_2 r_3 t_1^2 e^{-i(\phi_{12} + \phi_{23} + \phi_{31})}}{1 + r_1 r_2 r_3 e^{-i(\phi_{12} + \phi_{23} + \phi_{31})}} \quad \& \quad t_{IMC} = -\frac{r_2 t_1 t_3 e^{-i(\phi_{12} + \phi_{23})}}{1 + r_1 r_2 r_3 e^{-i(\phi_{12} + \phi_{23} + \phi_{31})}}. \quad (6-1)$$

If we make the simplifications: $r_2 = 1$, $t_2 = 0$, $r_1 = r_3 = r$, $t_1 = t_3 = t$, $r^2 + t^2 = 1$, and $\phi_{12} + \phi_{23} + \phi_{31} = \phi$, then the transmissivity and reflectivity coefficients become

$$r_{IMC} = \frac{r(1 + e^{-i\phi})}{1 + r^2 e^{-i\phi}} \quad \& \quad t_{IMC} = -\frac{t^2 e^{-i(\phi - \phi_{31})}}{1 + r^2 e^{-i\phi}}. \quad (6-2)$$

As we saw in the case of a simple two mirror cavity in Section 3.1.1, this is a sharply peaked function around the resonances.

The round trip phase in the input mode cleaner is the sum of the plane wave propagation contribution and the Gouy phase

$$\phi = kL + (n + m + 1)\Psi + \text{mod}_2(n) \pi, \quad (6-3)$$

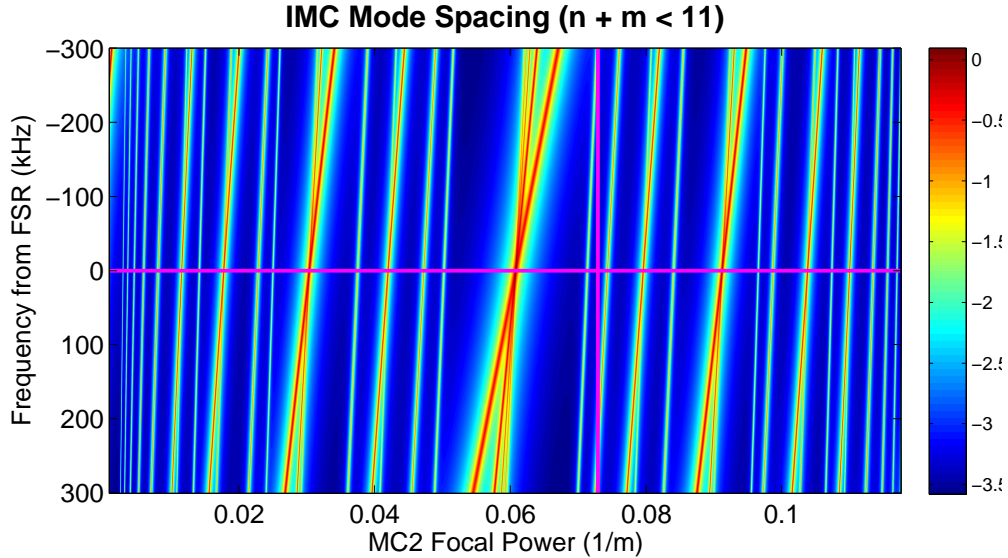


Figure 6-4: The mode spacing of the input mode cleaner. The color axis is the base ten logarithm of the power in transmission, not including the carrier, where the power incident upon the IMC is given by $P(n, m) = 10^{-(n+m)/5}$ and the HOMs up to order 10 are included. The bounds of the MC2 ROC (horizontal axis) are chosen such that the beam size on all of the mirrors is between 1 mm and 10 mm. The magenta lines show the carrier location and the chosen value for the MC2 ROC.

where the final term comes from the fact that the odd higher order modes (HOMs) pick up an extra π of phase upon reflection from a mirror. This feature of selecting the odd HOMs is unique to an optical cavity with an odd number of mirrors; such as the input mode cleaner. Each of the HOMs will resonate when the round trip phase is equal to π modulo 2π , and this fact allows us to derive the frequencies at which each mode will resonate. If we assume that the carrier is on resonance, then the frequency difference of each HOM from the carrier is given by

$$\Delta f = \frac{c}{2\pi L} \text{mod}_{2\pi} [(n + m)\psi + \text{mod}_2(n)] \quad (6-4)$$

For modal filtering, it is important that the modes which need to be filtered out be significantly further from the carrier resonance than the width of the carrier resonance (the cavity pole).

Figure 6-4 shows the modal spacing of the IMC near the carrier as a function of the curvature of MC2. The specific geometry chosen leaves the nearest HOMs up to order 10 at least 100 kHz from the carrier resonance. As we will see shortly, the FWHM of the carrier resonance is only ~ 9 kHz so this is a fine geometry for mode cleaning.

In addition to filtering the spatial mode of the input beam, the choice of a three mirror cavity also allows it to act as a polarization filter. In fact, the theoretical extinction ratio of the IMC is almost as high as that of a calcite wedge polarizer, which is one of the highest extinction ratio polarizers on the market. The reason that the IMC is able to filter the polarization of the input beam is because the two different polarizations of an optical beam pick up a different phase shift upon reflection off of a mirror at an angle, with one picking up π more than the other. This can be most easily understood by thinking about the two different polarizations in terms of circular polarization. Thinking in this basis, the choice that the positive direction of the optical axis points in the direction of energy flow implies that right circular polarization becomes left circular polarization and vice versa. This is equivalent to adding 180° of phase to one of the linear polarizations.

The input mode cleaner also passively filters the laser frequency above the cavity pole as was shown in Section 3.2. Choosing a higher finesse for the IMC is advantageous in that it moves the cavity pole to lower frequency which makes the IMC a better passive frequency filter, but has the disadvantage of storing more power inside of the cavity (see Section 3.1.2) which leads to stronger thermal effects. The cavity pole of the input mode cleaner can be expressed within the simplifications of equations (6-2) as

$$\Omega_0 = \frac{1 - r^2}{r^2} \frac{c}{L}, \quad (6-5)$$

where L is the *round-trip* length of the cavity. In the end a finesse of 515 was chosen which gives a cavity pole of 8.7 kHz and a power buildup of 167.

One more consideration which must go into the design of the input mode cleaner is a proper choice of the mass of the mirrors. It is important that the mirrors be massive

enough that the intensity fluctuations of the input beam do not cause the mirrors to move more than the required amount of frequency suppression. If we assume that the frequencies we are concerned about are well above the suspension resonances so that the mirrors act as free masses and that they are low enough that the intensity fluctuations act in phase on all three mirrors, then the length variation of the IMC, $\delta x(f)$, due to intensity fluctuations $\delta P(f)$ is given by

$$\delta x(f) = \frac{4 \delta P(f)}{(2\pi)^2 M c} \quad \Rightarrow \quad \delta f(f) = \frac{4 \delta P(f)}{(2\pi)^2 M \lambda L} \quad (6-6)$$

where $\delta f(f)$ gives the equivalent frequency fluctuations caused by the IMC length variations. Here M is the mass of the mirrors, c is the speed of light, L is the round-trip length of the IMC, and we have included a factor of 3 to account for the three mirrors of the cavity. The mass of the IMC mirrors was chosen to be 2.9 kg, and this curve is plotted in Figure 6-3 under the assumption of 165 W of input power and a RIN of 10^{-8} . Note that this RIN is an order of magnitude larger than the expected aLIGO RIN. From that figure one can see that the choice of $M = 2.9$ kg puts the technical radiation pressure noise well below the required level, and indeed there were no radiation pressure effects observed during commissioning of the LLO IMC.

The mirrors of the IMC are suspended from the aLIGO small triple suspensions[73] (HSTS). Figure 6-5 shows a front and side view of the main components of these suspensions. The top mass is a rectangular shaped metal stage while the intermediate mass has similar dimensions to the optic but is also made of metal. The three stages provide isolation from seismic motion proportional to f^{-6} above the three resonances which are near 1 Hz for all degrees of freedom except vertical. Vertical isolation is provided by blade springs to which the wires of the upper two stages are attached and provide isolation proportional to f^{-4} above the resonances near 1 Hz.

All three stages are outfitted with permanent magnets which can be actuated on by small electromagnets attached to the suspension cage (not shown in Figure




Figure 6-5: A front and side view of the LIGO small triple suspensions[73]. The suspension have three stages of passive isolation in all degrees of freedom except vertical, in which they have two. Each stage has magnets attached to it which can be pushed with small electromagnets which are attached to the frame (not shown). This allows the bottom two stages to be controlled in length, pitch, and yaw while the top can be controlled in all six degrees of freedom.

6-5). This allows actuation in length, pitch, and yaw at the bottom two stages (mirror and intermediate) while the top stage can be actuated in all 6 degrees of freedom. In addition, these electromagnets employ a shadow sensor which senses the location of the magnets as a readout of the position of each stage. The magnet strengths are chosen at each stage such that the middle stage has ~ 25 times more actuation authority and the upper stage has ~ 1500 times more actuation authority than the bottom stage. This necessitates the use of hierarchical control in which feedback at lower frequencies is offloaded to the upper stages of the suspension as discussed in Section 6.3.

6.3 Control Loops

Operating the input mode cleaner, as with the other key parts of the aLIGO interferometers requires maintaining active control of both the length and angular degrees of freedom. The length degree of freedom must be controlled in order to keep the light resonant inside of the cavity so that it is transmitted through to the interferometer. Active control of the length degree of freedom is also necessary to use the IMC as an active frequency stabilizer at frequencies below the cavity pole. Control of the angular degrees of freedom is important to keep the input beam and the cavity aligned to a common axis so that, among other things, the transmitted power is stable.

6.3.1 Length Control

Looking at the reflectivity coefficient of equations (6–2), the generic properties of which were explored in Section 3.1, we can see that as the cavity swings through resonance the reflected phase undergoes a very rapid transition from $-\pi/2$ to $\pi/2$. This phase change is sensed with the Pound-Drever-Hall[24] sensing scheme which adds an RF sideband to the incident laser beam and uses the phase modulation to amplitude modulation conversion near resonance (discussed in Section 4.4.1) to sense this rapid phase change.

Detection of a signal in this way is highly sensitive to the round-trip phase of the cavity, but there is some ambiguity in choosing an actuator for feedback. If we write out the round trip phase of the cavity in terms of the length of the cavity and frequency of the laser and include the possibility of variations in these two quantities we find

$$\frac{2\pi}{c} f \ell = \frac{2\pi}{c} (f_0 + \delta f)(\ell_0 - \delta \ell) \quad (6-7)$$

where we have assumed that we will compensate for variations in the laser frequency by changing the length or vice versa. Expanding the right hand side to first order and

rearranging gives us an equivalence relationship

$$\frac{\delta f}{f_0} = \frac{\delta \ell}{\ell_0} \quad (6-8)$$

which tells us that we can compensate for variations in the round-trip phase of the cavity equivalently by adjusting the length or the frequency by an amount given by this relationship.

Looking at Figure 6-3 we can see that the length noise of the cavity, caused by seismic and technical radiation pressure noise, dominates below ~ 5 Hz while the laser frequency noise should dominate above ~ 5 Hz. For this reason the ideal place to cross the feedback between the length and the frequency of the cavity is around 5 Hz so that below this crossover frequency the length of the cavity is servoed to the quiet laser frequency and above the crossover frequency the laser frequency is servoed to the quiet cavity length.

Figure 6-6 shows the modeled length control loops actually implemented in the IMC feedback system. The Advanced LIGO triple suspensions from which the mirrors of the IMC hang allow for actuation at each stage of the triple pendulum with progressively stronger actuation able to be applied at higher stages. In order to get enough actuation range it was necessary to feedback to the middle stage, M2, above the suspension resonances at ~ 1 Hz. This necessitated pushing the bottom stage, M3, crossover up from the ideal value of 5 Hz discussed above to ~ 15 Hz. The M2 to M3 crossover is around ~ 7 Hz and the feedback to the top stage, M1, is only used at very low frequencies, crossing over with M2 at ~ 100 mHz. The unity gain frequency of the IMC length loops is ~ 70 kHz and is limited by the delay in the frequency stabilization servo of the PSL.

Notice also in Figure 6-6 that the feedback to the frequency path is not rolled off at low frequencies but is DC coupled instead. This is because the common arm length (CARM) is used to control the laser frequency in a similar way to the IMC as




Figure 6-6: A model of the open loop transfer functions of the various actuators in the IMC length/frequency control servo. The triple pendulum suspension of the MC2 mirror is used in a hierarchical scheme to control the length of the cavity at low frequency using the laser frequency as a reference. At higher frequencies the length of the IMC is more stable and the feedback signal is instead used to stabilize the laser frequency.

discussed above. Since the laser must be simultaneously resonant in the IMC and the arms, the CARM feedback is sent to the length of the IMC and the DC coupling of the frequency feedback allows this scheme to adjust the laser frequency rather than just be canceled out by the IMC length control loops. At high frequencies where the CARM loop can not push the mirrors of the IMC fast enough the feedback is sent directly to the laser, offsetting the frequency from the length of the mode cleaner by a small amount. Modeling of the CARM control loops with expected noise sources[60] showed that the RMS offset from the ideal mode cleaner resonance should be less than 1 Hz which is significantly below the IMC cavity pole of 8.7 kHz.

6.3.2 Angular Control

Sensing of the angular degrees of freedom of the input mode cleaner is achieved through the use of differential wave front sensing[4][33]. The differential wavefront

sensing scheme relies on sensing the beat note between the first higher order mode and the fundamental Gaussian mode which is accepted by the cavity. When a tilt or translation exists between the optical axis of the input beam and the optical axis of the input mode cleaner the reflected beam from the cavity has a significant fraction of its power in the TEM₀₁/TEM₁₀ modes[59].

The TEM₀₁/TEM₁₀ modes beat with the sideband light in a unique way due to the fact that the field strength has opposite signs on opposite sides of the beam. This means that the beat note produced with the TEM₀₀ mode of the sideband light has opposite signs on opposite sides of the beam. These signs would cause cancellation if the entire beam were collected on a single detector, but a split photodetector is sensitive to this effect. For this reason the aLIGO differential wavefront sensors use a quadrant detector with the four quadrants demodulated individually at the RF frequencies which is sensitive simultaneously to the pitch and yaw modes by using the correct combinations of the four signals. In addition to the two differential wavefront sensors a quadrant detector (QPD), operating at DC, sitting behind MC2 gives absolute information about the input beam relative to the in-vacuum optical table. Figure 6-1 shows the out-of-vacuum wavefront sensors labeled WFSA and WFSB as well as the quadrant detector behind MC2 labeled QPD.

The differential wavefront sensors (WFS) are placed in the IMC reflected beam with 90° separation in Gouy phase so that the signals are orthogonal. They are not placed with any particular Gouy phase relationship with respect to the cavity mirrors, however. Indeed, the total accumulated Gouy phase from the cavity waist is roughly 55° for WFSA and 155° for WFSB.

In a cavity with more than two mirrors, the number of degrees of freedom of the cavity exceeds the two degrees of freedom of an optical beam. The WFS are only sensitive to the *relative* alignment between the optical axis of the input beam and the optical axis of the cavity. This leaves one degree of freedom unsensed in the input mode

Table 6-1: The measured sensing matrix of the angular control loops of the input mode cleaner.

	Mirror	WFSA ($\frac{W}{rad}$)	WFSB ($\frac{W}{rad}$)
	MC1	-46	300
Pitch	MC2	-863	377
	MC3	-91	291
	MC1	-413	51
Yaw	MC2	72	687
	MC3	453	-80

Note: The units of the sensing elements are in W/rad. The accumulated Gouy phase shift from the cavity waist is approximately 55° for WFSA and 155° for WFSB.

cleaner. This unsensed degree of freedom can be most concisely stated as the beam spot location on MC3.

The measured sensing matrix is shown in Table 6-1 in units of Watts per radian. This measurement agrees reasonably well with simulations in the higher order mode simulation package *Finesse*[30][5].

The high frequency power fluctuations out of the input mode cleaner due to angular fluctuations of the mirrors were found to be low enough that the angular control loops are engaged with very low unity gain frequencies. The control scheme has the mirrors following the input beam by feeding the WFS signals to the cavity mirrors with a 500 mHz unity gain frequency. In addition, a 10 mHz servo running on top of the WFS loops adjusts the pointing of the input beam to keep the beam centered on the QPD behind MC2.

6.4 Losses and Thermal Effects

A significant effort was made to characterize the losses in the IMC since lost light inside of the cavity can lead to a number of deleterious effects. Since the power buildup in the input mode cleaner is fairly high light lost to scattering off of the mirror surfaces can send photons out of the beam which have a chance to scatter back into the beam

and end up on a photodetector. These scattered photons cause spurious noise since they often interact with parts of the detector which are not isolated from the ground motion, such as the vacuum walls. Light lost to absorption in the input mode cleaner causes heat to build up in the mirrors which deforms the surface of the mirrors as discussed in Section 3.4. This deformation of the surface of the mirrors can lead to thermal effects which reduce efficiency of the coupling between the mode cleaner and the main interferometer.

6.4.1 Cavity Pole

The cavity pole for the input mode cleaner, given by equation (6-5), is sensitive to the round-trip losses of the IMC. In the limit of low losses we can model the loss to a reduced reflectivity of one of the mirrors. Doing so transforms equation (6-5) to

$$\Omega_0 = \frac{c}{L} \frac{1 - r^2 \sqrt{1 - \frac{\ell}{r^2}}}{r^2 \sqrt{1 - \frac{\ell}{r^2}}}, \quad (6-9)$$

where ℓ is the round trip power loss inside of the IMC, attributed solely to a loss in reflectivity of one of the mirrors. Hence, if we can measure the cavity pole of the input mode cleaner we can get a measure of the round-trip losses. Inverting (6-9) gives the losses expressed as a function of the cavity pole

$$\ell = \frac{r^4 (c + \Omega_0 L)^2 - c^2}{r^2 (c + \Omega_0 L)^2}. \quad (6-10)$$

Notice that one difficulty with measuring the losses in this manner is that the losses compete directly with the reflectivity of the input mode cleaner mirrors which are relatively high at 6000 ppm.

In Section 3.2.1 we saw that the laser frequency noise is filtered by the cavity pole in transmission of the input mode cleaner. A similar calculation shows that amplitude noise on the laser is filtered identically. This fact was exploited as a measurement of the cavity pole. A swept sinusoidal drive was applied to the error point of the intensity stabilization servo inside of the PSL enclosure. A photodiode was placed in a pickoff of




Figure 6-7: The measured IMC cavity pole is shown together with a simple single pole fit. The fit has a pole frequency of $8,686 \pm 108$ Hz which gives a finesse of 522. These values are within the error bars on the measured mirror reflectivities eliminating the possibility of excessive losses.

the beam inside of the PSL, and another matching photodiode was placed in a pickoff of the light transmitted from the IMC. Note that since the amplification electronics of most photodiodes have some color (frequency dependence) it is important for this measurement to use well matched photodiodes. Alternatively, the relative transfer function of the photodiodes can be measured ahead of time and calibrated out of the measurement. A network analyzer was then used to measure the transfer function of the mode cleaner to laser intensity noise from 0.5 kHz to 100 kHz. The results together with a fit to a simple single pole transfer function (3-44) are shown in Figure 6-7.

The measured cavity pole extracted from the fit shown in Figure 6-7 with one sigma fitting error is

$$f_0 = \frac{\Omega_0}{2\pi} = 8,686 \pm 108 \text{ Hz.} \quad (6-11)$$

This is very close to the design value[7] of 8717 Hz. Putting this into equation 6–10 gives the measured round trip losses in the IMC to be

$$\ell = 164 \pm 147 \text{ ppm.} \quad (6-12)$$

Clearly this method does not tightly constrain the losses; it is more effectively applied to higher finesse cavities than the input mode cleaner. Additionally, this calculation of the losses does not include the uncertainty in the reflectivity of the input mode cleaner mirrors since this was not specified by the vendor. An uncertainty in the reflectivity of 20 ppm (a typical value) increases the error in the round-trip losses to 168 ppm which is larger than the measured losses. In the next few sections we will try to get a more direct measure of the losses to better constrain this number.

6.4.2 Scattering

The dominant loss in the input mode cleaner, and indeed most optical cavities, is due to light being scattered out of the beam at the mirror surfaces. Light scattering in high quality dielectric mirrors is caused by a number of factors. Point defects on the mirror surface or inside the dielectric coating, surface micro-roughness, and large scale figure errors all contribute to scattering of photons from the optical beam. The large beam sizes in the IMC mean that the figure errors and surface roughness¹ are the dominant source of scatter. There are different formalisms for dealing with different types of surface scattering, but scattering from dielectric laser mirrors falls into the smooth-surface regime which is governed by the Rayleigh-Rice theory[79]. In this regime the fluctuations of the surface height of the mirror are significantly less than the wavelength of the incident light.

¹ Figure error and surface roughness both quantify deviations in the surface height of the mirror from the nominal. Figure error has a length scale comparable to the beam size while surface roughness has much shorter length scales.

If we work in the smooth surface limit, in which the amplitude of the vertical roughness on the surface of the mirror is much less than the wavelength of the light, then the angular distribution of the scattered light is determined solely from the statistical properties of the surface height fluctuations.[19] In particular, the angular distribution of the scattered light can be determined from a simple mapping from the two-dimensional power spectral density (PSD) of the surface fluctuations. In this limit each spatial wavelength can be thought of as a diffraction grating contributing to the power at the first order diffraction angle of such a grating, given by

$$\sin \theta_s = \sin \theta_i \pm \frac{\lambda}{d}, \quad (6-13)$$

where λ is the wavelength of the radiation, d is the spatial wavelength, and θ_i and θ_s are the incident and scattered angles respectively.

In this limit the bi-directional reflectance distribution function (BRDF) can be calculated from the 2 dimensional power spectrum by[19]

$$BRDF = \frac{1}{l} \frac{dl}{d\Omega} = 4k^4 \cos \theta_i \cos^2 \theta_s Q W(p, q), \quad (6-14)$$

where $k = \frac{2\pi}{\lambda}$ is the wavenumber of the light, Ω is a unit of solid angle θ_i is the incidence angle of the light, θ_s is the scatter angle, and $W(p, q)$ is the two-dimensional power spectrum given by

$$W(p, q) = \frac{1}{A} \left| \frac{1}{2\pi} \int dx \int dy e^{i(px+qy)} Z(x, y) \right|^2. \quad (6-15)$$

Here $Z(x, y)$ is the height of the surface as a function of spatial coordinates x and y . Q is the angle dependent polarization reflectance and has a rather complicated form which can be found in the appendix of Church et al.[19]. It contains information about how light of a given incident polarization ends up at a given scattered angle and polarization. The scattering in a given direction is proportional to $W(p, q)$ evaluated at the wavenumbers

$$p = k(\sin \theta_s \cos \phi_s - \sin \theta_i) \quad \& \quad q = k(\sin \theta_s \sin \phi_s). \quad (6-16)$$

Table 6-2: The scattering results for the MC1 and MC2 mirrors.

Optic	θ_i (deg)	θ_s (deg)	BRDF ($\frac{10^{-6}}{\text{sr}}$)	d (μm)	$\sqrt[4]{k^4 W}$ (1)	TIS (ppm)
MC1	45	-22	0.5 ± 0.1	1.0	0.021	1.6 ± 0.3
MC2	1	61	8.9 ± 1.9	1.2	0.044	28.0 ± 6.0

Note: θ_i is the incident angle, θ_s is the scattered angle, and BRDF is the measured BRDF. d is the spatial wavelength which would lead to scattering at that angle and $\sqrt[4]{k^4 W}$ gives a measure of the required amplitude at that spatial wavelength to lead to the observed scatter. Finally, TIS gives the total integrated scatter if we assume that the measured BRDF is constant (a somewhat dubious assumption).

Once one has the BRDF either by calculating from the surface PSD or from measurement, the total integrated scatter (TIS) can be calculated from[54]

$$TIS = \int_0^{2\pi} \int_0^{\pi/2} BRDF \cos \theta_s \sin \theta_s d\theta_s d\phi_s. \quad (6-17)$$

The TIS of a perfectly Lambertian scatterer is unity since it scatters all light from the beam with uniform intensity in all directions which leads to a BRDF for a Lambertian scatterer of $\frac{1}{\pi}$.

In order to measure scattering in the input mode cleaner, digital cameras (GigE cameras), were used together with different lens setups to image the surface of the three optics. Because the IMC is in vacuum, the available views of the optics are limited, and the ability to measure the entire BRDF is therefore limited. Between the four views of the three optics there are views available roughly equally spaced between 4° and 60° off of the optic normal. The four views are not comparable strictly by scattering angle though since MC1 and MC3 are at a large angle of incidence while MC2 is not.

Most of the scattering measurements were made by David Feldbaum and Eleanor King, but are reported here for completeness. For technical reasons two of the four views were not available in time for publication in this dissertation, but the results for the other two are shown in Table 6-2. The columns of the table show the incident

and scattering angles as well as the measured BRDF at those angles. The last three columns make some important assumptions whose validity will be discussed momentarily. The column labeled d gives the size of the diffraction grating which would lead to scattering at the given angles, and is calculated from equation (6–13). The column labeled $\sqrt[4]{k^4 W}$ is calculated from the BRDF using equation (6–14) and gives an estimate of the amplitude (normalized to the wavelength of the incident light) of the surface fluctuations which would be needed to give the observed BRDF. Finally, the column labeled TIS gives the total light scattered by the optic if one assumes that the measured BRDF value is constant, i.e. a Lambertian scatterer.

The assumption that the measured scatter is due to surface fluctuations is somewhat dubious. The camera images used to measure the scatter off of the optics clearly show point scatterers which could either be defects in the reflective coating or dust particles sitting on the surface. In the case of MC2 these point defects contribute as much as half of the total scattered power. In addition, the assumption of a constant BRDF which goes into the calculation of the TIS is likely invalid since this is not typical of high quality optics[54] which usually show increasing BRDF at angles within 20° of the specular angle.

6.4.3 Absorption

The absorption in the input mode cleaner was measured in two different complementary ways. The first probes the round-trip Gouy phase as a function of input power and infers from this the amount of thermal lensing at each of the mirrors. From this the total absorption can be inferred with a model of how the mirror surface deforms with absorbed power. The second method tracks the shift of the fundamental body mode of each of the three mirrors of the input mode cleaner individually. The body mode tracking has the advantage of being able to distinguish between absorption at the individual optics but has the disadvantage of being difficult to infer the absolute absorption. The




Figure 6-8: The shape of the first eight resonances of the IMC mirrors together with the eigenfrequency. The shapes and frequencies were calculated using a Comsol finite element simulation. The color shows the deformation along the optical axis (z coordinate).

Gouy phase measurement has the disadvantage of only being sensitive to the total absorption while being relatively easy to infer absolute absorption from.

Lets discuss the body mode measurement first since it can be used to inform the Gouy phase measurement. The majority of the work for this measurement was done by W. Korth from Caltech, but we will include it here for completeness. Determination of the eigenfrequencies of a solid elastic cylinder can be achieved analytically[35][71][39], but the results typically involve complicated series expansions which are not illustrative to stare at on paper. The important physical parameters are the shear modulus, bulk modulus, and Poisson's ratio, all of which can have some temperature dependence. In addition, the thermal expansion of the optic itself shifts the frequencies as the boundaries shift even without temperature dependence of the material properties. For this reason, it is a difficult problem to extract the absolute amount of absorption from this calculation. Since the three optics are nearly identical, with the exception of the curvature on the face of MC2, this measurement makes for an excellent way to measure the relative absorption of the three optics.




Figure 6-9: The frequency shift of each of the three input mode cleaner optics together with the input power. The slope of the lines shows that MC1 and MC2 absorb ~ 2.5 times as much power as MC3.

Figure 6-8 shows the shape and frequency of the lowest eight eigenmodes of one of the IMC optics. To measure the relative absorption the lowest order mode, at 17.5 kHz, was tracked in all three optics while sending in 120 W of input power for about 70 minutes, and the data is shown in Figure 6-9. Notice that the slope of the frequency shift is three times as much for MC1 and MC2 as it is for MC3. From this we infer that the amount of power absorbed by MC1 and MC2 is three times higher than the amount absorbed by MC3. This fact can be used to inform the Gouy phase shift absorption measurement.

The technique used to measure the round-trip Gouy phase of the IMC will be discussed at length in the next chapter, so we will only discuss how the absorption in the IMC can be inferred from such a measurement in this section. We begin by deriving a simple expression for the shift in the radius of curvature of an optic which is heated by absorption from a Gaussian beam. According to Winkler et al.[96] the sagitta s of an optic with a prescribed radius of curvature R over a beam size ω is given approximately

by

$$s \simeq \frac{\omega^2}{2R}. \quad (6-18)$$

The change in this sagitta induced by absorption of power from a Gaussian beam is derived to be

$$\delta s = \frac{\alpha}{4\pi\kappa} P_a, \quad (6-19)$$

where α is the coefficient of thermal expansion, κ is the thermal conductivity, and P_a is the *absorbed* power. Typical values of these parameters for fused silica are $\alpha = 0.55 \cdot 10^{-6} \text{ K}^{-1}$ and $\kappa = 1.38 \frac{\text{W}}{\text{mK}}$. Combining these two equations gives

$$\delta p = \frac{\alpha}{2\pi\omega^2\kappa} P_a \quad (6-20)$$

where δp is the shift in the radial power ($1/R$).

We can convert this to a shift in Gouy phase using the techniques developed in Section 3.1.4. Namely if we define $M(R)$ to be the matrix for reflection off of a mirror with radius of curvature R and $T(d)$ to be the ray matrix for translation by a distance d (see appendix 10.2.4), then the round trip in the mode cleaner starting at MC1 is given by

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = M\left(\frac{1}{\delta p_1}\right) T(l_{31}) M\left(\frac{1}{\delta p_3}\right) T(l_{23}) M\left(\frac{1}{\frac{1}{R} + \delta p_2}\right) T(l_{12}). \quad (6-21)$$

The Gouy phase can be calculated from this round trip matrix as we showed in Section 3.1.4 using

$$\eta = 2 \cos^{-1} \left(\text{sign}(B) \sqrt{\frac{A + D + 2}{4}} \right). \quad (6-22)$$

If we set $\delta p_1 = \delta p_2 = \delta p_3 = 0$ we find that the round-trip Gouy phase of the cold input mode cleaner is given by the simple expression

$$\eta_{\text{cold}} = 2 \cos^{-1} \left(\sqrt{1 - \frac{L}{2R}} \right) = 2 \cos^{-1}(\sqrt{g}), \quad (6-23)$$

where L is the *round-trip* length and R is the radius of curvature of MC2. This should look very familiar since it is the Gouy phase of an equivalent two mirror cavity with a flat mirror placed at the waist of the IMC. If we assume that the absorption is the same for all three mirrors, $\delta p_1 = \delta p_2 = \delta p_3 = \delta p$, and linearize about a we obtain for the shift in the Gouy phase with absorption

$$\Delta\eta = \eta_{hot} - \eta_{cold} = \frac{\alpha P_a L^2 (2g\omega_2^2 + \omega_{13}^2) + 2l_{31}^2 (1-g)\omega_2^2}{4\pi\kappa \sqrt{g(1-g)} L\omega_2^2\omega_{13}^2}, \quad (6-24)$$

where $L = 32.9485$ m is the total mode cleaner length, $l_{31} = 0.465$ m is the distance between MC3 and MC1, $\omega_{13} = 2.13$ mm is the beam size on MC1/MC3, and $\omega_2 = 3.38$ mm is the beam size on MC2. It is important to remember when using this expression that the absorbed power P_a contains the cavity buildup, the mirror absorption coefficient, and the input power to the cavity.

We could have used the body mode tracking experiment to inform this measurement, but the body mode tracking measurement requires significantly higher power levels to measure than the Gouy phase measurement. For this reason the body mode tracking was only performed once during high power testing while the Gouy phase measurement was performed numerous times. As we will see in just a moment the absorption data from the Gouy phase measurement varies from measurement to measurement so we will make the assumption of equal absorption so that each measurement can be compared to the others.

Table 6-3 shows the measurements and inferred absorption in the IMC over the course of 9 months. It appeared at the time that the absorption was increasing steadily, but there is a difficulty with such an interpretation. Namely, the beam spot location on the optic can drift around over time, and if the absorbers are small particles scattered about on the optic, then the average absorption could be constant and the apparent increase could be caused by beam spot drift over time. This is one of the reasons that the mode tracking data was not used to infer the absorption at times the mode tracking

Table 6-3: The data and inferred absorption from numerous repetitions of the Gouy phase absorption measurements in the IMC.

Date	Pow. (W)	f_{01} (Hz)	Δf_{01} (Hz)	Abs. (ppm/mir.)
1/17/2013	3.11	$29,266,891 \pm 60$	$6,230 \pm 69$	1.69 ± 0.02
	30.5	$29,273,121 \pm 35$		
7/23/2013	0.517	$29,268,352 \pm 65$	245 ± 85	3.51 ± 1.22
	1.034	$29,268,597 \pm 56$		
8/6/2013	0.203	$29,267,831 \pm 14$	358 ± 15	3.28 ± 0.14
	1.011	$29,268,189 \pm 6.4$		
9/30/2013	1.074	$29,266,056 \pm 39$	$1,175 \pm 54$	4.35 ± 0.20
	3.076	$29,267,230 \pm 37$		
6/17/2014	1.79	$29,877,736 \pm 46$	$2,249 \pm 79$	1.97 ± 0.07
	10.24	$29,875,487 \pm 64$		

Note: The Pow. and f_{01} columns show the input power level and location of the TEM_{01} peak while the δf_{01} and Abs. columns show the shift in this peak between power levels and the inferred absorption. Details of measuring the location of the TEM_{01} peak are given in the next chapter.

measurement. If we were to use the mode tracking data to inform the model used to calculate the absorption from the Gouy phase shift, we would find that the *total* amount of inferred absorption differs only at the level of 0.013%.

We can use this information to infer the amount of thermal lensing in the IMC. We can calculate the q parameter of the resonant mode in the IMC using the techniques discussed in Section 3.1.4 and the round trip matrix of (6-21). This can be propagated mathematically to the interferometer's input to ask how it overlaps with the spatial mode there. The overlap between two Gaussian beams with q parameters q_1 and q_2 can be calculated using

$$O(q_1, q_2) = 4 \frac{\Im(q_1)\Im(q_2)}{|q_1^* - q_2|^2}. \quad (6-25)$$

This is the overlap between the field strengths and includes the relative phase. We are more interested in the power lost due to this overlap which is given by $O(q_1, q_2)^2$. Doing this calculation for the input mode cleaner shows that even for 6 ppm of absorption per mirror and 200 W of input power the thermal lensing amounts to a 0.31% reduction in power coupling into the interferometer.

This minimal amount of thermal lensing was confirmed during high power testing of the input mode cleaner using the same type of digital cameras (GigE cameras) as were used for the scattering measurements. During these tests the power was increased in 30 W steps up to 120 W over the course of several hours and an image of the beam in transmission of the IMC was recorded every minute. Fitting of these camera images to a Gaussian profile showed no appreciable change in the beam size within the accuracy of the measurement which was about 5% of the beam size[62].

6.5 Length and Frequency Noise

As with all lock-in experiments, the measured feedback signals to the length and frequency paths are a measure of the sensitivity of the input mode cleaner (IMC) to length and frequency fluctuations. It is important to understand the source of this noise in the cavity measurement because while some noise sources are quieted by the feedback system, others, e.g. sensing noises, are impressed by the control system and show up as noise at the output of the IMC which then gets injected into the interferometer.

In this section we will discuss noise in the length and frequency control loops of the input mode cleaner. We will do so by starting with the noise budget for the input mode cleaner which was developed during initial commissioning of the IMC. We'll discuss the different known noise sources and see how they contribute to the overall sensing noise of the instrument. Finally, we'll close with a discussion of the unknown noise in the IMC between $\sim 3 - 80$ Hz.

6.5.1 Noise Budget

The noise budget for the input mode cleaner is shown in Figure 6-10. The red and orange curves show the feedback signals to the length and frequency respectively, and the dashed gray line shows the coherent sum of these two. The sum of these two lines represents the sensed length noise of the input mode cleaner or the frequency noise of the incident laser light. This does not necessarily mean that all of the noise is actually length or frequency noise since other sources of noise arising from the sensing and control can be dominant if the feedback system is not carefully designed. The goal of the noise budget is therefore to show that we understand where the measured noise in the IMC comes from so that we can tell if it is sensing noise or true length/frequency noise.

In order to understand how a particular noise source shows up in the IMC sensing chain we need to know the transfer function between each noise source and the control signals. Since it is not always possible to measure the transfer function between any two points, the heart of the IMC noise budget is a linear time-invariant (LTE) control model built in the commercially available software package Simulink. This control model is the same one used to produce Figure 6-6. This control loop uses a very realistic model of the aLIGO triple suspensions developed by Mark Barton[9].

At low frequencies the dominant sources of noise in the IMC are seismic and damping sensor noise. The seismic noise is measured using accelerometers which are mounted inside of the in-vacuum isolated table. This motion is then propagated through the control model which is dominated by the triple suspension transfer function at those frequencies. The damping sensor noise arises because the Advanced LIGO suspensions employ active damping instead of passive eddy current damping. The sensors used to sense the motion of the suspension are shadow sensors[15] which is to say that they pass a flag which is attached to one of the stages of the suspension between an LED and a photodiode. As the optic moves more or less light is blocked

Input Mode Cleaner Noise Budget




Figure 6-10: The noise budget of the Advanced LIGO input mode cleaner. The measured noise in the length path is shown in red, the measured noise of the frequency path is shown in orange, and the sum total of these two is shown in dashed gray. The sum total of the understood noises is shown in dashed black, and the individual terms are discussed in the text of Section 6.5.1. Notice that there is a significant deficit in the noise between $\sim 3 - 80$ Hz which is the topic of Section 6.5.2

from reaching the photodiode and the voltage change is interpreted as suspension motion. The noise of these sensors is approximately $1 \frac{\text{nm}}{\sqrt{\text{Hz}}}$ at 0.1 Hz and decreases as $f^{-1/2}$ at higher frequencies. This noise shows up in the damping signal which is applied to the suspensions and is a dominant source of noise only at the suspensions resonances. The damping sensor noise in Figure 6-10 is calculated by propagating this noise through the control model of the IMC.

At high frequencies the measured noise is dominated by frequency noise of the incident laser beam caused by acoustic vibrations of the laser table on which the pre-stabilized laser (PSL) is built (injection bench vibrations). These vibrations create

frequency noise because vibrations of the injection bench cause the resonant modes of the mechanical structures which support the optics to ring up. This causes the mirror surface to vibrate which modulates the path length that the laser travels between the source and the IMC. This path length modulation shows up as phase modulation of the light beam which is related to frequency modulation through

$$f(t) = \frac{1}{2\pi} \frac{d}{dt} \phi(t) \quad (6-26)$$

where $\phi(t)$ is the instantaneous phase and $f(t)$ is the instantaneous frequency. This noise is projected to the noise budget by measuring the motion of the mirror mounted at the top of the periscope on the PSL table with an accelerometer and converting this measured motion to frequency noise using equation (6-26). Note that the calculated noise does not exactly match all of the peaks in the measured noise spectrum, but this is due to the fact that we only measure the motion at one location on the PSL table. Empirical tests, i.e. banging on the table while watching the noise, show that most of these peaks are due to some kind of vibrations on the PSL injection bench.

The rest of the terms in the noise budget do not contribute to the measured noise, but are still important because some of them will set the sensing noise limit of the IMC. The sensing noise is the limiting factor in the ability of the IMC to suppress the frequency noise of the laser at higher frequencies where the seismic noise is sufficiently rolled off. Briefly, the radiation pressure noise is calculated by measuring the intensity noise inside of the PSL and calculating the radiation pressure caused by these fluctuations (including the IMC buildup). This calculated noise is then propagated through the controls model described above. The VCO noise is the expected, not measured, noise of the voltage controlled oscillator (VCO), which is the actuator used to control the frequency of the laser via an acousto-optic modulator[48], propagated through the controls model. The measured dark noise is a simple measurement of the output of the electronics feedback chain of the input mode cleaner with no light on the

photodiode, propagated through the input mode cleaner. Finally, the shot noise at the very bottom of the plot is a simple calculation of the expected white spectrum shot noise on the photodiode propagated through the controls model.

6.5.2 Noise Deficit

The missing noise between 5 Hz and 80 Hz was the subject of many investigations, but it is still not fully understood. The level of the noise in this region averages around that shown in Figure 6-10, but it fluctuates on many day timescales by as much as an order of magnitude but never coming down to the level of the understood noise. The result of the investigations were that there were many possible ways of increasing the noise, but no ways of decreasing it were found. In this section we will describe these investigations and their conclusions or lack thereof.

The first test that was tried was to increase the amount of pointing noise going into the input mode cleaner. As shown in Figure 6-11 injecting white noise into the piezo-actuated mirror which points into the IMC does increase the noise in the deficit region. Since the bandwidth of the wavefront sensors control loops is ~ 500 mHz it was not possible for the control system to be injecting this noise into the cavity. The suspected mechanism for this coupling is spatial non-uniformities in the locking photodiode of the IMC which has been seen in the past as one of the limiting factors for laser power stabilization[49][81]. These non-uniformities in the response of the photodiode couple the pointing fluctuations to voltage noise in the feedback system which show up as measured length noise. A back-of-the-envelope calculation using typical values seen in the past put this coupling mechanism below the witnessed coupling, but it was never conclusively confirmed or denied.

An increase in noise in the deficit region was also seen during high power testing of the input mode cleaner[61]. The increase in the noise was not linear in the power however; at 30 W of input power the noise at 10 Hz was $2 \cdot 10^{-12} \frac{\text{m}}{\sqrt{\text{Hz}}}$, this increased to $1 \cdot 10^{-11} \frac{\text{m}}{\sqrt{\text{Hz}}}$ at 60 W, $1.5 \cdot 10^{-11} \frac{\text{m}}{\sqrt{\text{Hz}}}$ at 90 W, and $2 \cdot 10^{-11} \frac{\text{m}}{\sqrt{\text{Hz}}}$ at 120 W. Although this was




Figure 6-11: The effect of driving white noise into both channels of the piezo-actuated mirror on the measured IMC length noise. The piezo-actuated mirror is mounted at the top of the PSL periscope and is able to actuate on the beam in both pitch and yaw. Driving white noise into both channels of the mirror showed a marked increase in the noise in the deficit region, but the coupling mechanism was unclear.

repeatable in the sense that returning to 30 W reduced the noise back to the original level, the limited window available for high power testing did not allow for checking of the day-to-day repeatability. The measured intensity noise at the time was well below the level required for radiation pressure fluctuations to be the cause of the noise in the deficit region, but the nonlinear, saturation-like increase in the noise with increasing power is further evidence that this is not the mechanism. The effect of intensity noise was investigated nonetheless by driving the error point of the PSL's intensity stabilization servo which created power fluctuations going into the IMC on the order of 10%. No coherence was found however with the length or frequency control signals of the IMC and no significant increase in the noise was witnessed.

In addition to these studies, there was a period of time during which the noise in the deficit region would show peaks corresponding to harmonics of ~ 1 Hz. These peaks would show up during times of particularly high seismic motion, and the camera

in the PSL looking at the reflection off of the reference cavity would show the motion of the fundamental frequency. This was traced to the fact that the locking photodiode for the reference cavity had been placed such that the reflected power went back towards the reference cavity, creating a low finesse, uncontrolled optical cavity between the reference cavity and the locking PD. Although this finding is not that important on its own, it points toward one more possible source of noise in the deficit region of the IMC. Scattered light seems unlikely however since it characteristically shows up as a zoo of resonances (similar to the IMC noise above 100 Hz) as it picks up the mechanical motion of the objects which scatter it back into the beam.

One final possibility which was not investigated very thoroughly is temperature fluctuations of the reference cavity inside of the PSL. We can estimate the magnitude of the temperature changes necessary to produce the noise witnessed in the deficit region of the IMC with a rather simple calculation. Length fluctuations of the reference cavity would show up as equivalent length fluctuations of the IMC by the relation

$$\delta \ell_I(f) = \frac{L_I}{L_R} \delta \ell_R(f), \quad (6-27)$$

where I denotes the IMC parameters and R denotes the reference cavity parameters. The reference cavity is made of a solid fused silica spacer with optically contacted mirrors on both ends. Temperature fluctuations of the reference cavity will show up as length fluctuations through the relation $\delta \ell_R(f) = \alpha \delta T_R(f)$, where $\alpha \simeq 0.5 \cdot 10^{-6} \text{ K}^{-1}$. The reference cavity is mounted in a vacuum chamber which provides a low pass filter for the surrounding temperature. The cavity temperature fluctuations are related to the external temperature fluctuations by

$$\delta T_R(f) = \left| \frac{f_0}{f_0 + if} \right| \delta T_E(f), \quad (6-28)$$

where the pole frequency was measured to be $f_0 = 216 \mu\text{Hz}$ [68]. Putting all of this together and inverting shows that the necessary temperature fluctuations external to the

reference cavity required to look like a given IMC length fluctuation is

$$\delta T_E(f) = \frac{L_R}{L_I \alpha} \left| \frac{f_0 + if}{f_0} \right| \delta \ell_I(f). \quad (6-29)$$

Plugging in $L_I = 32.9485$ m, $L_R = 0.406$ m, and $\delta \ell_I(10 \text{ Hz}) = 2 \cdot 10^{-13} \frac{\text{m}}{\sqrt{\text{Hz}}}$ shows that the necessary temperature fluctuations to cause the noise at 10 Hz is $\delta T_E(10 \text{ Hz}) = 230 \frac{\mu\text{K}}{\sqrt{\text{Hz}}}$. Experience suggests that un-driven temperature fluctuations of this level at 10 Hz are unrealistic.

6.6 Conclusion

In conclusion, the IMC is ready for gravitational wave detection; the control system is robust and lock acquisition is rapid and the scattering and absorption, though slightly higher than expected, are still low enough to not degrade the performance of the IMC. The measured noise of the IMC is largely understood, and propagation of this through the control model indicates that the frequency noise sent to the interferometer is within the requirements. The noise in the deficit region is not yet understood, but a number of hints have been found which should help lead to its resolution if it is found to cause a problem for operation of the aLIGO interferometers.

Part IV

EIGENSPECTRA MEASUREMENTS IN RESONANT OPTICAL INTERFEROMETERS

CHAPTER 7 AN INTRODUCTION TO EIGENSPECTRA MEASUREMENTS IN RESONANT OPTICAL INTERFEROMETERS

Part III of this dissertation contained a description of the experiments performed on the input optics of the Advanced LIGO interferometers. In this chapter and the following ones we will switch gears to a discussion of a general class of measurements which can be very useful for characterizing resonant optical interferometers. The work of this part of the dissertation, unlike that of Part III, is almost entirely my own. I designed and built the experiments, collected the data, and interpreted the results. The one notable exception is the initial identification of the peak locations of the input mode cleaner eigenspectra measurement (i.e. the model shown in Figure 8-2) in which Koji Arai of the California Institute of Technology and Paul Fulda of the University of Florida were both instrumental.

Measurement of the eigenspectra of resonant optical interferometers involves using a separate probe beam with a well defined relationship to the main resonant beam to trace out the resonance structure of an interferometer. This is analogous to looking at the Fourier transform of the sound produced by a musical instrument. Carrying this analogy a little farther, lets try to understand what information can be garnered from a measurement of the eigenspectra of a resonator, be it a musical instrument or an optical resonator.

7.1 Can One Hear the Shape of a Drum?

We are interested in this section in addressing what information is carried in the eigenspectra of a resonator. In our case the resonator is an optical interferometer, but the general question has been an area of active mathematical research since the early 1900s. The field was given new vigor and a more interesting title with the publishing of Mark Kac's article "Can one hear the shape of a drum?" in 1966[42]. A drum is similar to an optical field in that both obey the wave equation and, upon factoring off an oscillatory time component, both obey the Helmholtz equation (see Section 2.1). Spending a few

moments investigating the answer to this question will therefore give us some insight into what information is available to us about optical resonators through knowledge of their resonant spectra.

One of the first interesting results which came from the consideration of the Jeans radiation problem is that one can hear the area of the drum[42], or, in the case of the blackbody theory, the volume of the container. In addition, Kac shows that it is also possible to hear the circumference of a drum, though the proof relies on assumptions specific to a bounded two dimensional surface.

In 2010 Zelditch and Hezari[37] went a step further and showed that it is sufficient for the drum to be bounded by a surface which is analytic in order for it to have a unique eigenspectrum. Before Zelditch however a number of different non-analytic shapes had already been found whose eigenspectra were identical in theory. This was shown to be true in a physics experiment where Sridhar and Kudrolli[87] measured the spectral lines of microwave cavities up to the 54th eigenvalue and found them to be equal.

These theorems, even the one which is proven for three dimensions, do not carry over immediately to laser resonators unfortunately. This is because the boundary conditions of a laser resonator are Neumann instead of Dirichlet at the mirrors and the volume is infinite since the resonator is not bounded except at the mirrors. Being good physicists, lets disregard these issues and assume that, since the surfaces of the mirrors can be described by an analytic function on the scales of interest to us, the eigenspectrum of an optical resonator fully contains the information about the shape of the mirrors.

Even if this assumption is true these measurements run into practical problems which prevent the experiment from fully determining the surface structure of the mirrors. Namely, the power distribution of the higher order modes of a resonator gets further from the optical axis with increasing mode number. Hence, there is some point at which the mode sizes of the resonator get large enough that the power lost around the edges




Figure 7-1: The losses from the various higher order modes at a mirror with radius equal to 3ω . Notice that the clipping reaches the level of a few percent for mode orders greater than 6 and hits the level of 20 to 40 percent for modes of order 12 and greater. This imposes a practical limit on our ability to probe the eigenspectra of optical cavities since large losses prevent us from being able to ring up the resonances and measure their location.

of the mirrors prevents us from identifying the eigenfrequencies of its resonances.

Figure 7-1 shows the losses versus mode number for a mirror with a radius of 3ω . For the zeroth order mode the losses are less than a part per million, but by the seventh order modes they have reached the level of a few percent, and by the twelfth order modes they start nearing 30%. Modes with losses this high are difficult to identify because they are difficult to ring up and because the peaks are so broad that the center frequency becomes blurry. In addition, generating higher order modes from an initially Gaussian beam becomes more difficult as the mode numbers increase as illustrated in the left hand matrix of Figure 7-3.

With the knowledge that detection of resonances at very high orders will be difficult, we can make a simple argument for the size of spatial structures which we will be able to detect on the mirror surface with eigenspectra measurements. If a mirror has a phase map described by the function $\alpha f(x, y)$, then the beam which reflects off of that mirror




Figure 7-2: The separation between the lowest order zeros of the Hermite polynomial term in a Hermite-Gauss beam. This separation gives us an idea of the spatial deformations of the mirror surface which we can detect with the eigenspectra probe.

will pick up a spatially varying phase given by

$$E_r = E_i e^{-i\frac{2\pi}{\lambda}\alpha f(x,y)}. \quad (7-1)$$

If α is small then we can expand this term to first order giving

$$E_r \simeq E_i - \frac{i2\pi}{\lambda}\alpha f(x,y) E_i. \quad (7-2)$$

So, we see that if the function is given by a multiple of Hermite polynomials

$$f(x,y) = H_n \left(\frac{\sqrt{2}x}{\omega} \right) H_m \left(\frac{\sqrt{2}y}{\omega} \right), \quad (7-3)$$

then the dominant coupling at the mirror is to the nm Hermite-Gauss mode. We may therefore estimate the sensitivity of the eigenspectra measurements to spatial deformations at the surface of the mirror by examining the spatial separation of the zeros of the Hermite polynomial. This is precisely what is plotted in Figure 7-2 from which we

can see that the smallest spatial resolution that we can reasonably hope to achieve with the eigenspectra measurements is on the order of $\frac{\omega}{3}$.

7.2 Creating and Detecting Higher Order Modes

In order to probe the eigenspectra of resonant optical cavities, it is necessary to create and detect the higher order modes (HOMs) of a laser beam (discussed in Section 2.1). It is important to create these HOMs in order to have a probe for each of the resonant frequencies which we want to measure, and it is obviously important that we be able to detect them so that we can measure the eigenspectra. The process of creation and detection works similarly in practice, but the mathematics of how and why both cases work are slightly different so we will discuss both.

In order to create higher order modes on an initially Gaussian beam we typically clip the beam in some way. The measurements described in the following sections use a dental toothscraper with its point stuck into the center of the beam to create higher order modes at the input to both the IMC and the PRC experiments. In this section we will model the toothscraper as a triangular wedge of the beam which is blocked. We can ask how this scatters the beam into the higher order modes by calculating the overlap integral between our beam and the Hermite-Gauss (HG) eigenfunctions

$$\langle \Psi | \psi_{nm} \rangle = \int \Psi^* \psi_{nm} dA, \quad (7-4)$$

where ψ_{nm} is the normalized beam of order n, m , given in the Hermite-Gauss basis for a beam with no ellipticity as

$$\psi_{nm}(x, y) = \sqrt{\frac{2}{\pi}} \sqrt{\frac{1}{\omega_0^2 m! n! 2^{n+m} \left(1 + \frac{z^2}{z_R^2}\right)}} H_n \left(\frac{\sqrt{2}x}{\omega_0 \sqrt{1 + \frac{z^2}{z_R^2}}} \right) H_m \left(\frac{\sqrt{2}y}{\omega_0 \sqrt{1 + \frac{z^2}{z_R^2}}} \right) \exp \left(-\frac{(x^2 + y^2)(\lambda z_R^2 + i\pi\omega_0^2 z)}{\lambda\omega_0^2(z^2 + z_R^2)} \right) \quad (7-5)$$

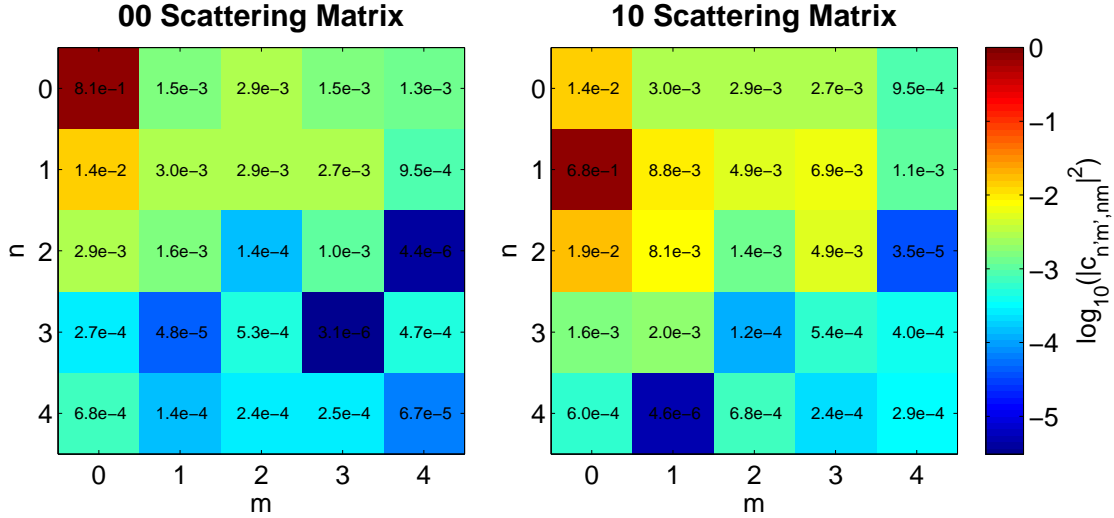


Figure 7-3: The scattering matrices for pie shaped clipping with a wedge angle of 0.2π radians. Each matrix shows the power in each of the modes after this clipping. The 00 scattering matrix shows these values with a 00 mode at the input, and the 10 scattering matrix uses a 10 mode at the input. The 00 matrix is important for creating higher order modes, and the multiple (elementwise) of the two matrices is important for detecting the 10 mode as discussed in the text.

as discussed in Section 2.1. The HG eigenfunctions provide an excellent vector space on which to analyze spatial beam characteristics since they are orthogonal

$$\langle \psi_{n'm'} | \psi_{nm} \rangle = \delta_{n'n} \delta_{m'm}. \quad (7-6)$$

In order to analyze how a particular beam shows up in the higher order modes we will define a coefficient which tells us how one of the HG modes partially transforms into the others when it is clipped;

$$c_{n'm',nm} = \int_{Trans} \psi_{n'm'}^* \psi_{nm} dA. \quad (7-7)$$

This coefficient can take values other than zero and one because the integral is taken only over the portion of the beam which is not blocked. Notice that $c_{n'm',nm}$ is real if the beam parameters (ω_0 , z_R , & z) of ψ_{nm} and $\psi_{n'm'}$ are the same.

As mentioned previously we will model the clipping of the beam as a blocked triangular wedge of size 0.2π radian. The scattering elements, given by (7-4),

are calculated by converting (7-7) to cylindrical coordinates and taking the angular integration limits to be $(0, 1.8\pi)$. Figure 7-3 shows the scattering matrices for a beam which is initially in the 00 and 01 modes for this particular type of clipping.

For the purpose of generating HOMs the important matrix is the 00 scattering matrix on the left of Figure 7-3 which shows how an initially 00 beam is scattered into the HOMs. Detection of the HOMs is slightly different. In general, detection for eigenspectra requires detection of a beat note between the 00 mode of the carrier and the $n_0 m_0$ mode of the sideband where $n_0 m_0 \neq 00$. Any two non-equal higher order modes are however orthogonal, as expressed in eq. (7-6), so the overlap integral between them is zero. In order create a detectable signal we use the same method as was used to create HOMs; clipping of the beam. The clipping scatters the 00 mode of the carrier into the full spectrum of higher order modes (00 scattering matrix in Figure 7-3), as well as scattering the $n_0 m_0$ mode of the sideband into the full spectrum of HOMs (10 scattering matrix in Figure 7-3 is an example). The detected beatnote in the end consists of the sum of all of these different HOMs of both the carrier and the sideband interfering together. It is not enough, therefore, to simply consider interference between only the 00 portions of the beams when calculating the signal strength.

Figure 7-4 shows the sum of the RF power detected in all of these modes up to order 4 for different levels of cylindrical (pie shaped) clipping. The vertical axis shows the detected RF power for both carrier and sideband beams of unit power; perfectly overlapped single mode beams would give 2. The horizontal axis shows the amount of total power clipped from the beam, and the different sidebands which we would like to detect are shown in different colors/linestyles. The right hand plot runs the full range from 0% to 100% clipping while the left hand plot shows a blowup of the shaded region. In the experiments of the following chapters clipping at the 10% level was chosen as it is simultaneously sensitive to all of the modes considered.




Figure 7-4: The power of the RF beat note is shown for different levels of cylindrical clipping with a carrier and sideband of unit power. Modes up to order 4 are considered to calculate the scattering matrices between the 00 mode and the n_0, m_0 mode (as in Figure 7-3) which are then multiplied elementwise and summed. The plot on the left hand side is a zoomed view of the shaded area in the right hand plot.

7.3 Fitting the Measurements

Extracting information from the measured eigenspectra requires fitting the individual peaks in some way that extracts the interesting information without being overly sensitive to spurious nearby features or noise. In all of the eigenspectra measurements, the structure of the resonances is identical to the standard Pound-Drever-Hall[24] (PDH) error signal commonly used to lock resonant optical interferometers. The I phase portion of this error signal was found to provide a robust fitting function from which we can extract all of the parameters of interest. The general procedure used to fit each individual peak, to be described in more detail shortly, is to first rotate the signal into the I phase and then fit it to the I phase portion of a PDH signal with weighting around the center of the signal.

The fitting function is given by the imaginary part (i.e. the I phase) of

$$F = b + d \frac{a - 1}{a - e^{-i 2\pi \frac{\epsilon_0}{c} (f - f_0)}}, \quad (7-8)$$

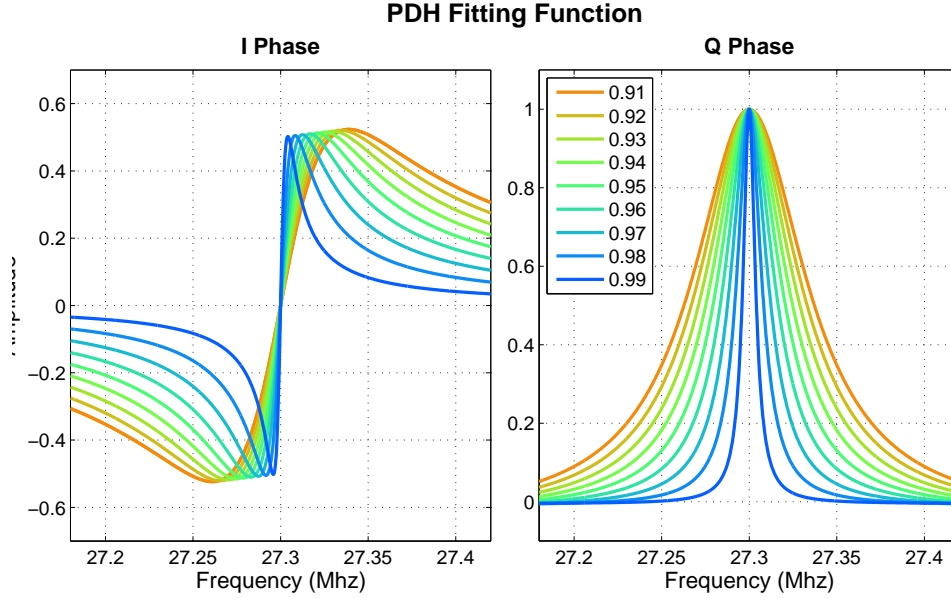


Figure 7-5: The imaginary part (I phase) and real part (Q phase) of the PDH signal used to fit the data from the eigenspectra measurements. The function is given by equation (7-8), where the parameters used in this plot are $b = 0$, $d = 1$, $f_0 = 27.3$ MHz, and a is varied over the values shown in the legend.

where a , b , d , and f_0 are the fitting parameters. Some examples of the imaginary (I phase) and real (Q Phase) parts of this function are shown in Figure 7-5. The two parameters with physical significance are a which is made up of the mirror reflectivities and losses, and f_0 which is the center frequency of the resonance. For a two mirror cavity with input and output reflectivities $r_1 = \sqrt{R_1}$ and $r_2 = \sqrt{R_2}$ and losses between the two mirrors of L (where L is the loss in *power* not amplitude) the coefficient a is given by

$$a = (1 - L)r_1r_2 \quad \rightarrow \quad L = 1 - \frac{a}{r_1r_2}. \quad (7-9)$$

This is used to calculate the losses from the fitting parameters in the later sections.

In order to fit the data to this function some pre-processing is necessary. First the phase shift due to a delay in the electronics is fitted to a measurement with a large frequency span and used to subtract this phase shift from all of the other measurements which typically have narrower frequency spans. The phase shift due to a delay is given

by

$$e^{-i2\pi\tau f}, \quad (7-10)$$

where τ is the delay in the measurement and is typically near $\tau \simeq 0.16 \mu\text{s}$. In order to fit this delay the measurement must be taken over many different peaks so that the phase shift at each of the peaks does not influence the fit.

Once the phase shift due to the delay is taken out the data must be rotated into the I phase which is accomplished by multiplying by a static phase. The signal is rotated into the I phase algorithmically by using a procedure similar to a least squares fitting routine. There are two parameters that are varied, the rotation phase and the center point of the peak. The algorithm folds the data about the center point and looks for anti-symmetry in the imaginary part and symmetry in the real part by calculating the sum of the square of the difference between the two parts. In the end the parameters which are picked are the rotation phase and symmetry point which minimize this sum, exactly like a least squares problem. The center point is used as an initial guess for f_0 when fitting to the I phase signal.

Finally, the I phase portion of this data is fit to the imaginary part of equation (7-8). Because this is a non-linear fit with four parameters the typical fitting routines found in e.g. Matlab which use a gradient search of the χ^2 parameter do not do a good job if the initial guesses are not very near the final location. This is because noisy data can have a lot of very small local minima in the least squares parameter as the fitting coefficients are varied. For this reason the fitting algorithm starts with an intelligent guess at the parameters which is passed to a blanket search over a reasonable set of values for all of the parameters. The end result of this blanket search is a 4 dimensional tensor filled with the least squares values of all combinations of the trial parameters. The parameters which give the minimum of this least squares tensor are passed to Matlab's *Isqcurvefit* function for final fitting. The errors for each of these fitted parameters are calculated by

the Jacobian estimate and residual values which are passed out from this function in the standard way.

CHAPTER 8 EIGENSPECTRA MEASUREMENTS OF THE INPUT MODE CLEANER

In this chapter we will discuss the usefulness of these eigenspectra measurements in characterizing the Advanced LIGO (aLIGO) input mode cleaner (IMC). The IMC was described in detail in Chapter 6 so this chapter will focus on the design and setup of the experiment and the measurements performed.

8.1 Experimental Setup

The layout of the IMC eigenspectra probe is shown in Figure 8-1. The aLIGO electro-optic modulator (EOM), discussed in Chapter 4, is used to add phase sidebands to the beam coming from the pre-stabilized laser (PSL). The EOM is driven by the radio frequency (RF) output of the network analyzer through an RF amplifier which boosts the drive strength to 24 dBm. The beam is clipped with a dental toothscraper after passing through the EOM and before being injected into the vacuum system.

Inside of the vacuum system the beam is routed to the IMC and injected through MC1. The triangular IMC cavity is made up of the three mirrors MC1, MC2, and MC3, and the transmitted beam is transmitted through the MC3 mirror. The first steering mirror after the IMC, IM1, has a transmissivity of 500 ppm, and this beam is brought out of the vacuum system to the IOT2L table.¹

On the IOT2L table another dental toothscraper is inserted into the beam before it lands on a radio frequency photodiode (RFPD). The RFPD uses a 1 mm diode, and its amplification circuitry has a frequency response which is flat up to ~ 80 MHz. The signal of this RFPD is sent to the input of the network analyzer, and the data is recorded to disk.

The probe beam for this measurement is a pair of phase sidebands which are added to the beam by modulating at the 45.5 MHz input of the aLIGO EOM. Using

¹ The IOT2L table is the same one to which the IMC reflected beam is brought.

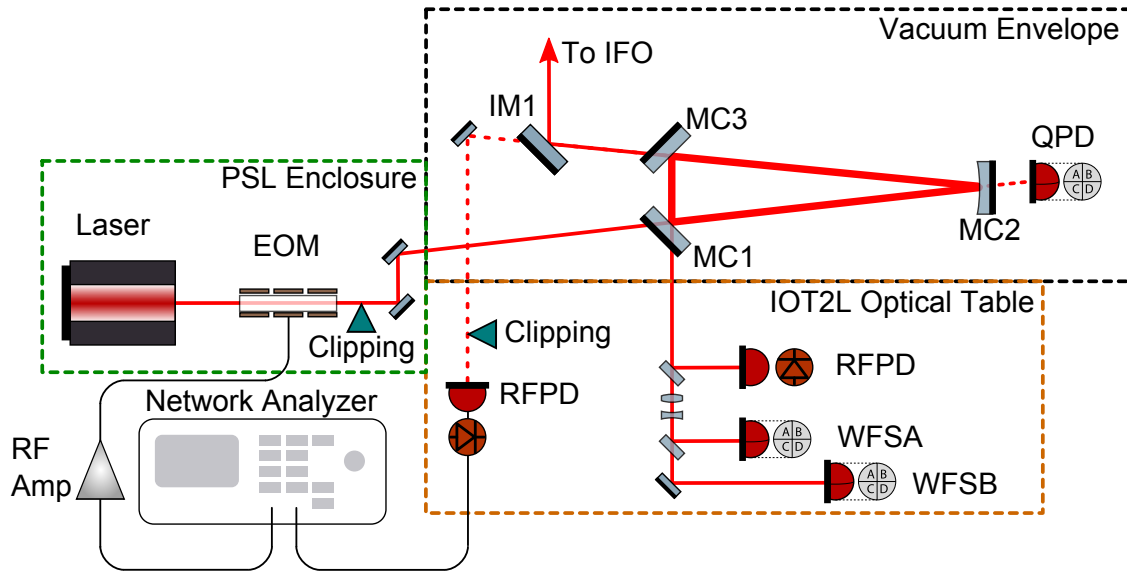


Figure 8-1: The experimental setup of the IMC eigenspectra measurement. The probe beam in this experiment is generated by adding phase sidebands to the main beam by borrowing the 45.5 MHz port of the electro-optic modulator. The signal is demodulated in transmission of the IMC using the pickoff behind IM1 which is brought out of the vacuum system to an optical table called IOT2L.

phase sidebands has the advantage that the probe beam is automatically phase locked and aligned to the main beam.² This type of probe beam adds a small amount of complication to the analysis, however, because each particular cavity resonance will show up twice when sweeping over a full free spectral range (FSR) of the cavity (see Figure 8-2); once when the lower sideband passes through the resonance and again when the upper sideband does so. An additional limitation particular to phase sidebands arises because the 180° phase difference between the two sidebands causes the total signal to cancel out at the zeroth order resonance of an ideally resonant cavity.³ This is

² This is equally true of amplitude sidebands which were used to measure the cavity pole in Section 6.4.1.

³ Although, as we will see shortly, this also allows for an accurate probe of the location of the FSR.

why it was necessary to use amplitude modulation (AM) to measure the cavity pole in Section 6.4.1.

The 45.5 MHz input of the EOM is used to add phase sidebands to the beam because, as discussed in Section 4.2, the response of the resonant circuit attenuates the drive above the resonance. This is also the reason that the cavity scans were typically taken between the FSRs at 27.3 MHz and 36.4 MHz since it is high enough to be well into the shot noise limited regime and low enough for the EOM circuitry to have a relatively flat response. Additionally, 24.0 MHz must be avoided in order not to destabilize the IMC and cause it to lose lock since 24.0 MHz is the RF frequency of the sidebands used to sense the IMC length.

8.1.1 Calculation of the Signal

Lets calculate the signal that shows up at the RFPD in transmission of the IMC as we map out one of the HOM resonances. We will define the total power of the TEM₀₀ Gaussian beam incident upon the first toothscrapper to be P_0 . The power in each of the sidebands will then be $P_s = \frac{\beta^2}{4} P_0$, where β is the modulation depth, and the power in the carrier will be $P_c \simeq P_0$ in the limit of small β . The field immediately after the toothscrapper will be

$$E_0 = E_c e^{-i\omega_0 t} \sum_{nm} c_{nm,00} \psi_{nm} + E_s e^{-i(\omega_0 + \Omega)t} \sum_{nm} c_{nm,00} \psi_{nm} - E_s e^{-i(\omega_0 - \Omega)t} \sum_{nm} c_{nm,00} \psi_{nm}. \quad (8-1)$$

Here ψ_{nm} is as defined in (7-5), ω_0 is the carrier frequency, Ω is the RF frequency of the sidebands, and $c_{n'm',nm}$ is the mode overlap coefficient defined in equation (7-7).

Lets call the particular sideband of interest $n_0 m_0$. The cavity will reject all but the 00 mode of the carrier and all but the $n_0 m_0$ mode of *one* of the two sidebands. Taking into account the transmissivity of the pickoff optics, the field immediately before the second toothscrapper is

$$E_{t1} = E_c e^{-i\omega_0 t} t_{PO} c_{00,00} \psi_{00} \pm E_s e^{-i(\omega_0 \pm \Omega)t} t_{PO} t(\pm\Omega) c_{n_0 m_0,00} \psi_{n_0 m_0}, \quad (8-2)$$

where t_{PO} is the field transmissivity of the pickoff optics, $t(\pm\Omega)$ is the field transmissivity of the IMC to the sidebands (i.e. the thing we want to measure), and the plus or minus sign is chosen depending on whether the resonance is probed with the upper or lower sideband. Note that we have assumed that the transmissivity of the IMC to the 00 mode of the carrier is unity.

The second toothscraper will scatter both the carrier and the sideband into the full litany of HOMs;

$$E_{t2} = E_c e^{-i\omega_0 t} t_{PO} c_{00,00} \sum_{nm} c_{nm,00} \psi_{nm} \pm E_s e^{-i(\omega_0 \pm \Omega)t} t_{PO} t(\pm\Omega) c_{n_0 m_0, 00} \sum_{nm} c_{nm, n_0 m_0} \psi_{nm}. \quad (8-3)$$

Finally, the power detected at the RFPD will be

$$\begin{aligned} \langle P \rangle = & |E_c|^2 |t_{PO}|^2 |c_{00,00}|^2 \sum_{n' m', nm} c_{n' m', 00}^* \psi_{n' m'}^* c_{nm, 00} \psi_{nm} \\ & + |E_s|^2 |t_{PO}|^2 |c_{n_0 m_0, 00}|^2 |t(\pm\Omega)|^2 \sum_{n' m', nm} c_{n' m', n_0 m_0}^* \psi_{n' m'}^* c_{nm, n_0 m_0} \psi_{nm} \\ & \pm E_c^* E_s |t_{PO}|^2 c_{00,00}^* c_{n_0 m_0, 00} t(\pm\Omega) e^{-i(\pm\Omega)t} \sum_{n' m', nm} c_{n' m', 00}^* \psi_{n' m'}^* c_{nm, n_0 m_0} \psi_{nm} \\ & \pm E_c E_s^* |t_{PO}|^2 c_{00,00} c_{n_0 m_0, 00}^* t^*(\pm\Omega) e^{i(\pm\Omega)t} \sum_{n' m', nm} c_{n' m', 00} \psi_{n' m'} c_{nm, n_0 m_0}^* \psi_{nm}^*. \quad (8-4) \end{aligned}$$

The photodetector will integrate these terms over all space, and the orthogonality of the HG modes, eq. (7-6), implies that only those terms under the summations which satisfy $n = n'$ and $m = m'$ will survive. The first two lines will simply be the total DC power in the carrier and sideband respectively; we will drop them since the RF signal is the one of interest. We can simplify the RF part by recalling that the overlap coefficients are real and by assuming, without loss of generality, that the carrier and sideband were in phase

at the first toothscraper. With these simplifications the RF part becomes

$$\langle P_\Omega \rangle = \pm 2E_c E_s T_{PO} c_{00,00} c_{n_0 m_0,00} \left(\sum_{nm} c_{nm,00} c_{nm,n_0 m_0} \right) \left\{ \Re[t(\pm\Omega)] \cos(\Omega t) + \Im[t(\pm\Omega)] \sin(\Omega t) \right\} \quad (8-5)$$

So, we can see that demodulation of this signal maps out both quadratures of the transmission coefficient as the RF frequency sweeps across the resonance.

We can use this to estimate the power in the signal detected at the photodiode to make sure that it is above the electronics noise of our equipment. The magnitude of the RF signal is

$$|\langle P_\Omega \rangle| = P_0 \beta T_{PO} c_{00,00} c_{n_0 m_0} \left(\sum_{nm} c_{mn,00} c_{nm,n_0 m_0} \right). \quad (8-6)$$

Using equation (4-1) and the parameters given in Section 4.2 we can estimate a modulation depth of $\beta = 0.02$ for 24 dBm drive. As stated above, the IM1 pickoff fraction is $T_{PO} = 500 \cdot 10^{-6}$. Referring to the 00 scattering matrix in Figure 7-3 we can see that the two coefficients outside of the sum have values on the order of $c_{00,00} = 0.9$ and $c_{n_0 m_0,00} \simeq 0.05$. Finally, referring to Figure 7-4 we can see that for clipping at the 10 % level we can expect the sum in the above equation to have a value on the order of

$$\sum_{nm} c_{mn,00} c_{nm,n_0 m_0} \simeq 0.05.$$

Putting all of this together we find that the RF power in our signal is

$$|\langle P_\Omega \rangle| \simeq 23 \cdot 10^{-9} P_0. \quad (8-7)$$

The photodiode used for these measurements has a responsivity of $\eta = 0.74 \frac{\text{A}}{\text{W}}$ and a transimpedance of $R = 1 \text{ k}\Omega$ which gives an electrical signal with an amplitude of 170 μV with $P_0 = 10 \text{ W}$ of power into the IMC which is well above the network analyzer noise floor of $17 \frac{\text{nV}}{\sqrt{\text{Hz}}}$.

Finally, as one last check, we can calculate the shot noise equivalent power due to the DC light on the photodiode. It is given by the Schottky formula as

$$S_p(\omega) = \sqrt{\frac{2e}{\eta} \langle P_{DC} \rangle} = 50 \frac{\text{pW}}{\sqrt{\text{Hz}}}, \quad (8-8)$$

where e is the electron charge and η is the PD responsivity. This is again significantly below the signal level, giving an SNR of greater than 10^3 for 1 second integration time with 10 W input to the IMC.

8.1.2 Example Scan

Figure 8-2 shows an example scan across a full free spectral range of the IMC. Overlaid on top of the measured data is the prediction from a model built in the modeling package Finesse[31][30] by Paul Fulda and Koji Arai. The height of the peaks in the model were chosen simply to match the measurement and therefore should be considered arbitrary. This measurement shows that modes up to order 4 are detectable with this measurement technique. The model did not include the contribution from the P polarization which shows up as a broad peak in the center of the FSR.

8.2 Frequency/Length Offsets and Measurements of the Length

When probing one of the free spectral ranges of a resonant optical cavity with phase sidebands one should get no signal. This is because, as mentioned previously, the phase sidebands come with opposite signs, and when they are transmitted identically through the optical cavity they cancel each other. This provides a useful probe of offsets in the control system of the input mode cleaner.

Consider what will happen if the control system of the IMC is holding the laser slightly off resonance from the cavity's length due to e.g. an electronic offset in the feedback servo. For concreteness, let us say that the laser frequency is slightly higher than the ideal resonance. This means that the carrier is sitting off center on




Figure 8-2: An example HOM sweep of the IMC across a full FSR together with a Finesse model. Each peak is labeled with its Hermite-Gauss mode number for easy reference.

its transmission curve⁴. As we increase the RF frequency of the phase modulation the two phase sidebands move away from the carrier, one decreasing in frequency and the other increasing in frequency, until they hit the next FSR. The upper sideband will hit its transmission peak slightly before the FSR though, and the lower sideband will hit its transmission peak slightly after the FSR. This leads to a double-lobed signal which still goes to zero at the FSR, but has increased RF power with opposite phase on either side of the FSR. Note that this effect is identical to the one which converts PM sidebands into AM sidebands discussed in Section 4.4.

Figure 8-3 shows the results of using this mechanism to tune the offset in the IMC control system. The red line shows a sweep across one of the FSRs of the IMC with an electronic offset in the servo, and the blue line shows the null signal at the FSR

⁴ Examples of simple cavity transmission curves are shown in Figure 3-2




Figure 8-3: The effect of tuning the control loop offset on the eigenspectra measurement of the free spectral range. An offset in the control loops causes a double lobed signal at the FSR when probing it with phase sidebands. Reducing the locking point offset causes the signal to disappear completely at the FSR.

Table 8-1: Two measurements of the IMC length with the eigenspectra probe.

Date	ν_0 (Hz)	ℓ (m)	$\Delta\ell$ (mm)
Design	9,099,471	32.947	N/A
21 Sep., 2012	$9,099,055 \pm 100$	$32.9485 \pm 4 \times 10^{-4}$	1.5 ± 0.4
8 Nov., 2013	$9,099,385 \pm 30$	$32.9473 \pm 1 \times 10^{-4}$	0.3 ± 0.1

Note: ν_0 is the measured free spectral range, ℓ is the round-trip length, and $\Delta\ell$ is the difference between the measured length and the design length. The design parameters are taken from the input optics final design document: Arain et al. “Advanced LIGO Input Optics Final Design Document.” LIGO DCC T0900386 (2011).

when the offset is tuned out. Notice that the red line shows the double lobe pattern just described.

The narrow dip and rapid phase transition mapped out when the sideband frequency is swept across the FSR with a microscopic offset allows for a very precise identification of the macroscopic length of the IMC. The FSR is related to the macroscopic

length through

$$\ell = \frac{c}{\nu_0}, \quad (8-9)$$

where c is the speed of light, ν_0 is the free spectral range, and ℓ is the *round-trip* length of the cavity. This measurement was performed on two occasions about 1 year apart, and the results are shown in Table 8-1. The ~ 1 mm shift in the length of the IMC which occurred between the two measurements is not well understood. The noise of the reference oscillator was not taken into account in the error analysis, and it is possible that this drifted over such long timescales.

8.3 Input Polarization Tuning

The broad peak in the middle of the two s polarization FSR peaks in Figure 8-2 is due to the residual p polarization left on the input beam to the IMC. This peak shows up in the middle of the s polarization FSR peaks because of the fact that the IMC has an odd number of mirrors. In cavities with an odd number of mirrors, the two different polarizations differ in the round-trip phase that they pick up by π . This fact makes the IMC an excellent polarizer.

This provides a useful way of tuning the input polarization to the IMC. To correct the polarization, a half wave plate was placed in the input beam inside of the PSL enclosure and adjusted to minimize the p polarization peak. Figure 8-4 shows the results of this adjustment. The method was found to be sensitive at the 0.05° level of adjustment in the HWP which corresponds to 0.1° of polarization rotation in the input beam. Surprisingly, this polarization adjustment was found to significantly reduce the noise in the IMC wavefront sensors. Our hypothesis was that this was connected to RFAM in the p polarization beam which was reflected off of the cavity, but this was never confirmed.

8.4 Absorption Measurements

The measurements in the previous sections are useful for tuning up the IMC, but one of the key results of applying the eigenspectra method to the IMC was to measure the absorption. The model used to infer the absorption from the shift in the frequency of




Figure 8-4: The p polarization peak of the IMC eigenspectra measurement and the effect of adjusting the input polarization. This peak is used as a fiducial to adjust the input polarization to the input mode cleaner. The blue trace shows the p polarization peak before tuning, and the orange line shows the same peak after adjustment of a half wave plate placed in the input beam to the IMC.

the 01 mode was described in Section 6.4.3, and the various measurements over time were presented in Table 6-3. In this section we will go into a little more detail about the measurements; focusing particularly on the shift in all of the higher order modes and one possible way in which the measurement can be contaminated. Namely, by a power dependent shift of the beam spot position on the optics.

One of the key assumptions that went into the absorption model described in Section 6.4.3 was that the surface deformation of the heated optic is well described by the Winkler approximation. This approximation describes the change as a simple parabolic change in the height of the surface of the optic. This is not strictly true since the change in the surface of the optic is localized on the beam spot position. Figure 8-5 shows, amongst other things, the surface deformation calculated from a finite element

model⁵ together with the surface deformation calculated with the Winkler approximation. The two clearly differ, though the difference is not large until one goes far away from the beam.

One of key predictions of the Winkler based model is that the shift induced by thermal effects is a pure shift in the Gouy phase of the cavity. In Section 2.1.2 we saw that the phase picked up during the propagation of higher order Gaussian modes beyond that of a plane wave is equal to

$$\psi_{nm} = (n + m + 1)\psi_{00}, \quad (8-10)$$

where ψ_{00} is the Gouy phase picked up by the fundamental mode. This means that, since we are measuring relative to the 00 mode, the shift in the frequency locations of the various HOM resonances should be given simply by

$$\Delta f_{nm} = (n + m)\Delta f_{01}, \quad (8-11)$$

where Δf_{01} is the shift of the 01 mode.

One way to test this hypothesis is to track many of the HOM locations as the power is changed in the IMC. This was done over the course of 6 hours during which the power was ramped from ~ 2 W to ~ 10 W five times, measuring the HOM locations for ~ 30 minutes at each point. The results are summarized in Table 8-2. The final column in that table shows the difference between the measured frequency shift and the expected frequency shift given by equation (8-11) where Δf_{01} is taken to be the mean of the shift in the 01 and 10 modes. Except for the anomalously low shift of the 03/21 mode, the shifts are all within 30% of what is expected.

⁵ Finite element models of the surface deformation agree very well with the Hello-Vinet analytic solution[36] discussed in Section 3.4.

Table 8-2: The results of tracking the eigenfrequencies of the IMC HOMs up to order 4 as the power is stepped up and down.

Mode	f @ 1.79 W (Hz)	f @ 10.24 W (Hz)	Δf (Hz)	$\frac{\Delta f_m - \Delta f_e}{\Delta f_e}$
01	29,877,736 \pm 46	29,875,487 \pm 64	2,249 \pm 79	0.17
10	29,269,212 \pm 23	29,270,797 \pm 36	1,585 \pm 42	-0.17
02/20	31,238,784 \pm 51	31,241,509 \pm 74	2,725 \pm 90	-0.29
11	27,907,460 \pm 64	27,904,237 \pm 68	3,223 \pm 93	-0.16
03/21	28,663,192 \pm 48	28,663,248 \pm 78	56 \pm 92	-0.99
30/12	30,485,200 \pm 114	30,490,484 \pm 130	5,285 \pm 173	-0.08
04/22/40	28,527,843 \pm 299	28,519,444 \pm 226	8,399 \pm 375	0.10
13/31	30,630,681 \pm 119	30,636,285 \pm 115	5,605 \pm 165	-0.27

Note: The results are largely in line with expectations from the Winkler approximation model except for the anomalously low shift of the 03/21 mode. The middle three columns show the measured frequencies and shifts, and the final column shows the normalized difference between the measured shift and the expected shift given by (8-11) where Δf_{01} is taken to be the mean of the 01 and 10 modes.

The reason for the anomalously low shift in the 03/21 mode is not well understood, but there are a number of possibilities. This particular peak produced the most non-standard signal during the measurement, meaning that the shape of the peak did not look very much like a typical Pound-Drever-Hall signal. Regardless of the cause, it is possible that the best fit to this peak of a PDH signal is dominated by features which do not shift with increased power. Another possible explanation, although less likely, is that the shift in the spot positions on the mirrors during this measurement (to be discussed post haste) moved the beam to a location in which the frequency shift induced by the motion canceled out the frequency shift induced by the thermal effects.

Figure 8-5 shows the relative scales of the deformations at the surface of the MC2 mirror. Two cross-sections of the surface roughness of the mirror are shown together with two different calculations of the thermal deformation and a plot of the power in the

Comparison Between Surface Roughness and Thermal Bump

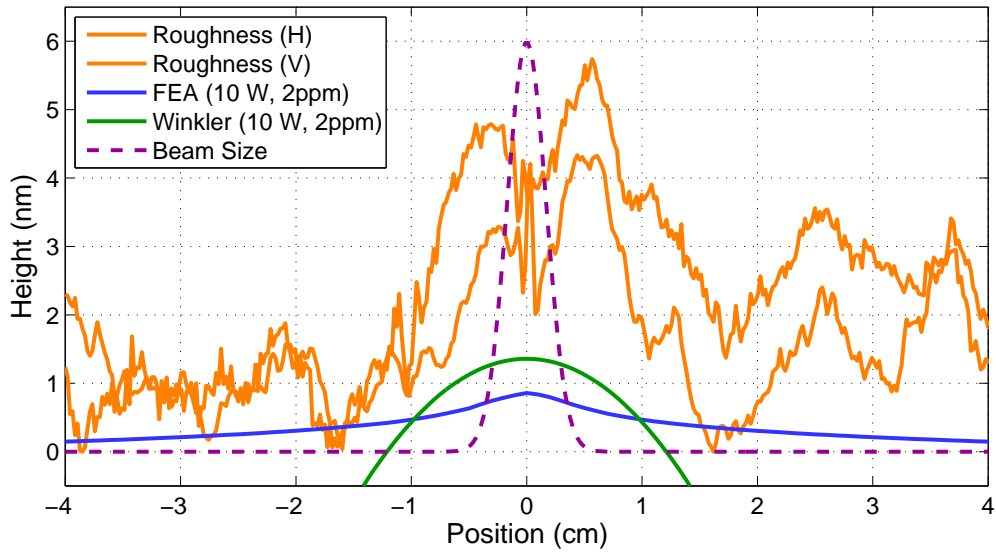


Figure 8-5: The height of various deformations on the surface of MC2. The two lines labeled surface deformation show the two cross sections of the measured surface uniformity of the MC2 optic. The two lines labeled FEA and Winkler show the surface deformation induced by heating at the optic caused by 10 W of power into the IMC and 2 ppm of uniform absorption. The FEA line is calculated with a finite element analysis tool and the Winkler line is calculated from the Winkler approximation described in Section 6.4.3 and has a 0.5 nm offset for clarity. The dashed line shows the beam size on that optic for reference, and its height has no absolute meaning

beam for scale. The point of this plot is to show that the surface roughness of the optic is comparable in scale to the thermal deformations which we are trying to measure with the eigenspectra probe. If the beam remains in a static position on the optic while the power is ramped up and down, then the surface roughness will not contribute to the power dependent shift in the frequencies of the various HOM peaks. However, if there is a uniform shift in the beam spot position on the optic as the power is changed, then this surface roughness can contribute significantly. A Fourier transform based optical simulation was put together by Hiro Yamamoto[97] which suggests that shifting the beam position on the optic by amounts on the order of the beam size is capable of producing eigenfrequency shifts up to 1 kHz which is comparable with the shifts expected from the thermal bump at 10 W of input power.

Table 8-3: The motion of the beam spot location on the three IMC mirrors during the measurements summarized in Table 8-2.

Optic	DOF	1.79 W (ω)	10.24 W (ω)	Δ (ω)
MC1	P	-0.26	-0.19	0.07
	Y	0.91	1.00	0.09
MC2	P	-0.08	-0.08	0.00
	Y	-0.62	-0.62	0.00
MC3	P	-0.19	-0.22	-0.01
	Y	-0.05	-0.06	-0.01

Note: All numbers are in units of the beam size at the optics.

During the measurements summarized in Table 8-2 the beam spot location was tracked throughout the six hour measurement. This was accomplished by dithering all of the optics in pitch and yaw at frequencies spaced roughly 2 Hz apart between 12 Hz and 20 Hz. The IMC control signal which is fed back to the laser frequency to keep it locked to the cavity was band-passed around these frequencies, demodulated against the dither signals, and heavily low-passed. This demodulated signal tracks the spot location on each of the mirrors in pitch and yaw because the coupling coefficient between angle and length is proportional to the offset from the center and undergoes a sign change when it crosses the center. The motion was found to be very uniform as the power was increased and decreased which points to the cause being some power dependence in the WFS control loops. The amount of spot shift on each of the mirrors is summarized in Table 8-3. During the measurement, none of the beams shifted by more than one tenth of the beam size at the optics.

CHAPTER 9 EIGENSPECTRA MEASUREMENTS OF THE POWER RECYCLED MICHELSON INTERFEROMETER

Eigenspectra measurements were also very useful in the early stages of Advanced LIGO (aLIGO) to characterize the power recycling cavity. At the Livingston Observatory the central portion of the interferometer was installed first, and the end test masses (ETMs) followed a few months later. During this intermediate time period the PRC eigenspectra measurement was set up to characterize the power recycling cavity (PRC) with the interferometer resonant in the power recycled Michelson configuration (PRMI). The PRMI state uses the Michelson formed by the input test masses (ITMs) as the rear mirror of a resonant cavity formed with the power recycling mirror.

Since the light of the RF sidebands in the full aLIGO configuration does not get into the arm cavities, this provides an excellent probe of the modal structure seen by the sideband light. When the full interferometer is locked, however; the modal structure of the carrier will see an effective cavity formed by the power recycling mirror and the ETMs. I.E. the carrier will effectively not see the spatial structure in the bulk of the ITMs. This means that the mode probed in the PRMI configuration will be different than the one seen by the carrier in a fully locked interferometer.

9.1 Experimental Setup

The optical and electronic layout of the PRC probe experiment is shown in Figure 9-1. The red beam shows the forward traveling portion of the main interferometer beam which comes in from the PSL, is transmitted through the IMC (MC1-3), reflects off of four steering mirrors (IM1-4), and is transmitted through the Faraday isolator before entering the power recycling cavity (PRC) by transmitting through the power recycling mirror (PRM). The reflected beam, shown in green, is deflected by the Faraday isolator and sent to the HAM 1 vacuum chamber (not shown) where a portion of the light is

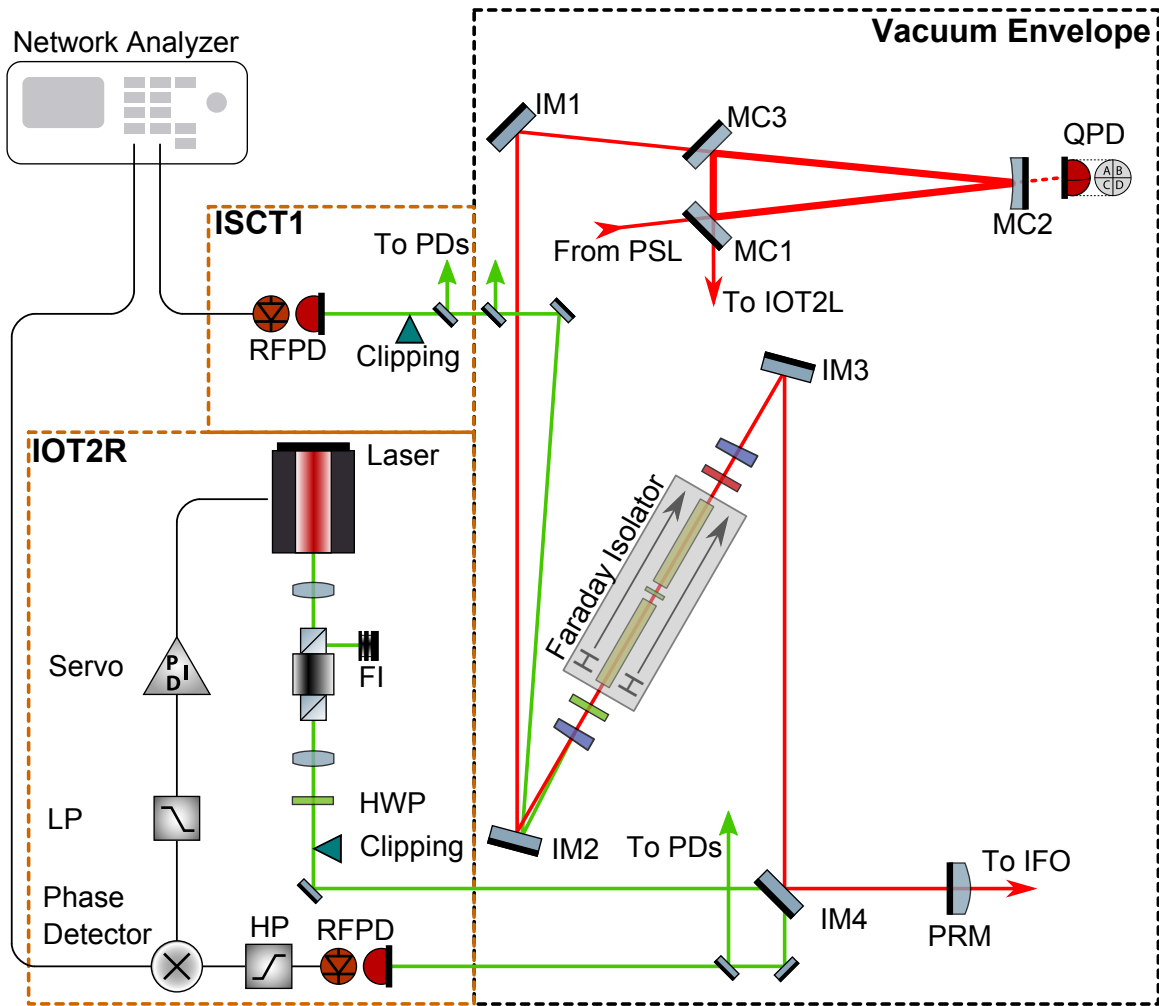


Figure 9-1: The layout of the PRC probe experiment. The main interferometer beam is shown in red while the reflected beam and relevant pickoff beams are shown in green.

picked up and sent to the ISCT1 table¹. On the ISCT1 table the light is further split between two RF photodiodes and two wavefront sensors, only one of which is shown in Figure 9-1. One of the RF photodiodes on this table is a narrow band photodiode tuned to the modulation frequencies of the interferometer, but the other is one of the aLIGO broadband RFPDs[28] (BBPD), and it is this photodiode which is used for readout of

¹ The portion which remains in vacuum is sent via suspended mirrors to in-vacuum RF photodiodes and wavefront sensors which are used for low noise readout of the length and angular degrees of freedom of the interferometer.

the PRC eigenspectra². The beam is partially clipped with a dental toothscraper or a partially closed iris immediately before the BBPD in order to be sensitive to the higher order modes as discussed in Section 7.2. All totaled the transmissivity of all of the optics combined from the Faraday isolator pickoff to the BBPD is $T_{PO} = 6250$ ppm (one 5% pickoff and three 50/50 beamsplitters).

The IM4 steering mirror immediately before the power recycling cavity is slightly transmissive, $T_{IM4} = 2400$ ppm, and therefore transmits a pickoff of both the forward and backward traveling beams which are both shown in green in the layout. The forward going beam has additional beam splitters behind it which pick off 10% of the light for a quadrant detector and then 90% of the light for the in-vacuum intensity stabilization PDs (summarized as one BS in Figure 9-1) before sending the remaining light out of the vacuum system to the IOT2R table. The backward going beam is brought directly to the IOT2R table without any in-vacuum attenuation.

This backward going path is used as the injection path for the auxiliary laser, and the injection path includes a Faraday isolator to dump the light coming out of the interferometer. A set of lenses and a half wave plate (HWP) are also included in the injection path to match the spatial mode and polarization of the auxiliary laser to that of the interferometer. Additionally, a dental toothscraper is inserted into the beam immediately before entering the vacuum system in order to generate HOMs on the auxiliary beam. The majority of the power injected from the auxiliary laser reflects off of IM4 and travels along the same path as the pickoff of the forward going main beam back to the IOT2R table. These two beams are aligned onto an RFPD for beat

² The broadband photodiode is used in the Advanced LIGO locking scheme to detect the beat notes between the sidebands and their second harmonics. These signals have the advantage of being relatively insensitive to the resonance condition of the carrier in the recycling cavities which changes dramatically when the arms are brought onto resonance.

note detection. This setup has the advantage of using IM4 as the beamsplitter which combines the auxiliary beam and the main beam. This is advantageous because any phase noise picked up by the auxiliary beam before entering the vacuum system due to e.g. vibrations of the optical table and mirrors is suppressed by the phase locked loop. Additionally, any phase noise picked up by the two beams on their way out of the vacuum system is not impressed since it is common to both beams.

9.1.1 The Phase-Locked Loop

The feedback system uses a common control scheme known as a phase-locked loop (PLL) to hold the frequency of the auxiliary laser at a specific RF frequency away from the main laser. This offset RF frequency is set by the RF output of the network analyzer. To accomplish this the difference beat note between the two lasers is read out with an RFPD whose bandwidth is 300 MHz. This RF signal is fed to a phase detector which is fed the RF output of the network analyzer at the other input. The output of the phase detector is simply the multiple of the two RF signals which, for two sinusoidal signals, is given by

$$\begin{aligned} V_{PD} &= V_1 V_2 \sin(\omega_1 t) \sin(\omega_2 t + \phi) \\ &= \frac{V_1 V_2}{2} (\cos((\omega_2 - \omega_1)t + \phi) - \cos((\omega_2 + \omega_1)t + \phi)). \end{aligned} \quad (9-1)$$

The low pass filtering after the phase detector filters out the sum signal and leaves the signal which shows up at the difference frequency. When the two frequencies are the same, this error signal simplifies to

$$\begin{aligned} V_{PD} &= \frac{V_1 V_2}{2} \cos(\phi) \\ &\simeq -\frac{V_1 V_2}{2} \left(\phi - \frac{\pi}{2} \right), \end{aligned} \quad (9-2)$$

where the second line holds for ϕ near $\pi/2$. So we see that we have a linear error signal in the phase difference of the two. This allows us to hold the auxiliary laser at a well defined offset from the frequency of the main laser.

The feedback mechanism actuates on the auxiliary laser's frequency. This is accomplished with two different mechanisms, both of which actuate on the effective length of the resonant cavity inside of the laser. The high bandwidth, low range actuator uses a piezo-electric crystal to press on the laser cavity which changes its true length as well as its effective length by shifting the index of refraction of the laser crystal via the elasto-optic effect. The low bandwidth, high range actuator uses a small thermo-electric cooler attached to the laser crystal which shifts the actual resonator length through thermal expansion of the resonator and also shifts its effective length through the thermally dependent shift in its index of refraction, dn/dT . The fast actuator has an actuation coefficient of $0.3 \frac{\text{MHz}}{\text{V}}$ while the slow actuator has an actuation coefficient of $1.7 \frac{\text{GHz}}{\text{V}}$. The important point here is that the feedback mechanism actuates on the laser frequency.

The relationship between phase and frequency is given by

$$\phi(t) = \int \omega(t) dt \quad \Rightarrow \quad \phi(s) = \frac{\omega(s)}{s}, \quad (9-3)$$

where the second equation is expressed in the Laplace domain. This means that the control system has a natural $1/f$ shape in the loop so that a stable control loop can be attained simply by applying the appropriate gain and ensuring that the delays in the system don't take away too much phase margin.

9.1.2 Feedback Electronics

In practice, we need more gain at low frequencies since the linear region of the error signal is fairly narrow. For this reason the feedback servo shown in Figure 9-2 was designed to boost the low frequency gain and still give a unity gain frequency of ~ 50 kHz. The fast path is shown at the top of the schematic and starts with a pair of inputs which get summed together and are followed by a monitor output. The extra input and monitor were included to facilitate with measurement of the open loop transfer function of the locked control loop. This is followed by a variable gain stage which allows

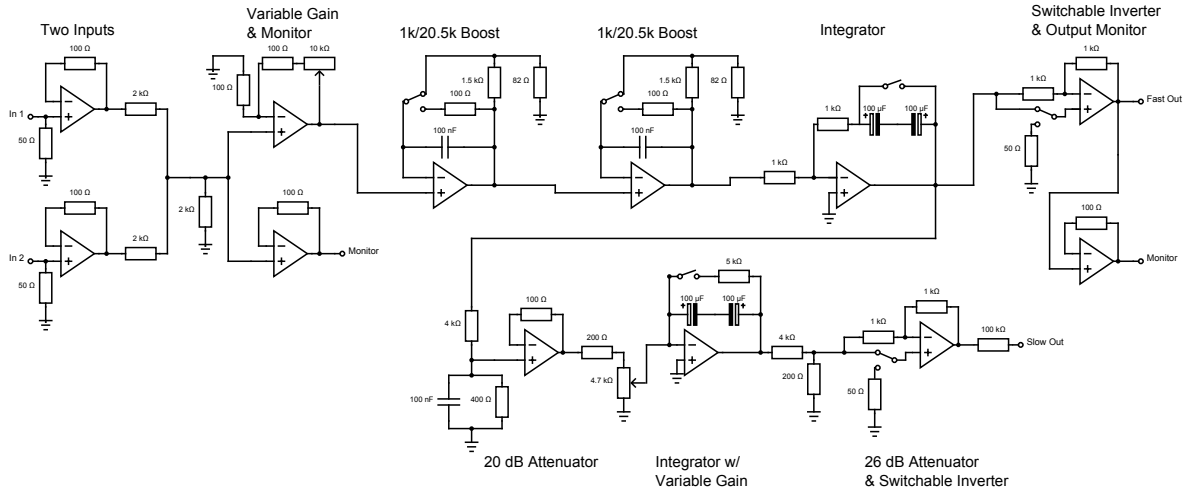


Figure 9-2: The feedback controller used for the phase-locked loop of the PRC probe. The goal was to boost the gain at low frequencies to assist with lock acquisition and provide a very attenuated pickoff to actuate on the slow feedback of the laser which has an incredibly high internal gain.

the gain to be varied from 6 dB to 40 dB. Following that are a pair of boost stages which have a gain of 0 dB when switched off. When they are switched on the boost stages have a pole at 1 kHz and a zero at 20.5 kHz with a gain of 1 at high frequencies and a gain of 40 dB at low frequencies. Next is an integrator with a pole at DC and a zero at 3 Hz. Finally, the fast loop ends with a switchable inverter with a gain of 0 dB and a monitor output for watching the control signal.

The slow path is picked off immediately before the switchable inverter in the fast path. The goal of the slow path was to provide large DC gain so that the control signal from the fast path could be offloaded to the slow path, but with enough attenuation that the crossover frequency would be low even with the massive amount of relative actuation gain in the slow path. The slow path therefore starts with 20 dB of attenuation and follows with an integrator which has a zero at 70 mHz when switched off and a pole at DC when switched on. The final stage of the slow path has a switchable inverter in order to get the correct sign in the feedback loop.

9.1.3 Calculation of the Signal

We can now calculate the signal we will measure at the RF photodiode in reflection of the PRMI. The auxiliary beam is clipped on the table before being injected into vacuum and joining the main beam. We will write it as

$$E_a = E_s e^{-i(\omega_0 + \Omega)t} \sum_{nm} c_{nm,00} \psi_{nm}, \quad (9-4)$$

where $P_s = |E_s|^2$ is the power in the auxiliary beam and the rest of the notation is the same as that used in Section 8.1.1. After transmitting through IM4 this beam will combine with the main beam such that the field incident on the PRC is given by

$$E_0 = E_c e^{-i\omega_0 t} + E_s t_{IM} e^{-i(\omega_0 + \Omega)t} \sum_{nm} c_{nm,00} \psi_{nm}, \quad (9-5)$$

where $P_c = |E_c|^2$ is the power in the main beam entering the interferometer.

Lets pause for one moment to observe one interesting fact about a beam with a single sideband such as ours. A single sideband beam can be written as

$$\begin{aligned} E &= E_c e^{-i\omega_0 t} + E_s e^{-i(\omega_0 + \Omega)t} \\ &= E_c e^{-i\omega_0 t} + \left(\frac{E_s}{2} - \frac{E_s}{2} \right) e^{-i(\omega_0 - \Omega)t} + \left(\frac{E_s}{2} + \frac{E_s}{2} \right) e^{-i(\omega_0 + \Omega)t}. \end{aligned} \quad (9-6)$$

This is a beam which has both phase modulation and amplitude modulation sidebands, both with a modulation depth of E_s . Hence, the PRC eigenspectra probe experiment exhibits properties of both phase and amplitude sideband measurements. As pointed out in the last chapter, this will matter most when we measure the resonance structure at the FSR where the PM sidebands can get converted to AM if there is an offset in the control loops.

For most of the PRMI experiments the interferometer was locked to the RF sidebands so that the carrier light is completely reflected off of the cavity. Likewise, all of the HOMs of the sideband will also be reflected except for the one which we are measuring which will trace out the reflectivity coefficient. Taking into account the

finite pickoff fraction on the way to ISCT1, the field before being clipped by the second toothscraper will be

$$E_r = E_c t_{PO} e^{-i\omega_0 t} + E_s t_{IM} t_{PO} e^{-i(\omega_0 + \Omega)t} \left(r(\Omega) c_{n_0 m_0, 00} \psi_{n_0 m_0} + \sum_{nm \neq n_0 m_0} c_{nm, 00} \psi_{nm} \right), \quad (9-7)$$

where $T_{PO} = |t_{PO}|^2$ is the pickoff fraction. The toothscraper on the ISCT1 table will spray all of the modes in both the carrier and the sideband into the full spectrum of HOMs, giving

$$\begin{aligned} E_{r2} = & E_c t_{PO} e^{-i\omega_0 t} \sum_{nm} c_{nm, 00} \psi_{nm} + E_s t_{IM} t_{PO} e^{-i(\omega_0 + \Omega)t} r(\Omega) c_{n_0 m_0, 00} \sum_{nm} c_{nm, n_0 m_0} \psi_{nm} \\ & + E_s t_{IM} t_{PO} e^{-i(\omega_0 + \Omega)t} \sum_{nm \neq n_0 m_0} \sum_{n' m'} c_{nm, 00} c_{n' m', 00} \psi_{n' m'} \end{aligned} \quad (9-8)$$

We are now in a position to calculate the power in this beam. As usual, we will get DC terms, RF terms at Ω (the ones of interest), and RF terms at 2Ω . Because of the complexity of the full resulting expression, we will only show the relevant terms which show up at 1Ω . These are

$$\begin{aligned} \langle P \rangle_{\Omega} = & 2E_c E_s T_{PO} \sqrt{T_{IM}} c_{n_0 m_0, 00} \left(\sum_{nm} c_{nm, 00} c_{nm, n_0 m_0} \right) (\Re[r(\Omega)] \cos(\Omega t) + \Im[r(\Omega)] \sin(\Omega t)) \\ & + 2E_c E_s T_{PO} \sqrt{T_{IM}} \left(\sum_{nm \neq n_0 m_0} \sum_{n' m'} c_{n' m', 00} c_{nm, 00} c_{n' m', nm} \right) \cos(\Omega t), \end{aligned} \quad (9-9)$$

after spatial integration by the photodetector. We can see that, just as with the IMC measurement in Chapter 8, this measurement maps out the reflectivity coefficient of each particular higher order mode. In this case there is an additional term due to the beating of all of the HOMs which are reflected off of the cavity with the carrier.

We can estimate the signal strength from this equation. Since the signal has the RF contribution from all of the other higher order modes which are reflected off of the cavity, the problem in this case is one of dynamic range instead of signal strength on resonance. As a worst case we will assume that the large double sum in the above expression is equal to 1 which means that the full beat between all of the higher order

modes and the carrier is detected. In this case, the background RF power at the photodiode will be

$$\langle P \rangle_{BG} = 2\sqrt{P_c}\sqrt{P_s}T_{PO}\sqrt{T_{IM}} = 475 \mu\text{W}, \quad (9-10)$$

for $P_c = 3 \text{ W}$ and $P_s = 200 \text{ mW}$. The change in the signal when the HOM comes onto resonance will be given by

$$\langle P \rangle_{Sig} = 2\sqrt{P_c}\sqrt{P_s}T_{PO}\sqrt{T_{IM}}c_{n_0m_0,00} \sum_{nm} c_{nm,00} c_{nm,n_0m_0} (1 - |r(\Omega)|) = 1.1 \mu\text{W}, \quad (9-11)$$

where we have used a 0.05 for both the sum and the HOM coefficient and used the measured cavity visibility of 90%. This gives a necessary dynamic range of less than 10^3 which is well within the limits of the network analyzer.

9.2 Non-Ideal Clipping Effects

The discussion in Section 7.2 on the detection of higher order modes by clipping would lead one to believe that the dependence of the signal on clipping should be one of magnitude only. This is, however, not what was seen in the PRC probe experiment. In this experiment there was a qualitative difference in the appearance of the resonance structure of the various modes as well as a difference in the apparent location and width of these peaks.

Figure 9-3 shows the measurement of the zeroth and first order peaks with various different clipping configurations. Each plot contains one curve with no clipping, two curves with clipping by a dental toothscraper inserted into the beam, and two curves clipped by an iris placed symmetrically on the beam. In the case of the toothscraper and iris, the beam was clipped at two different power levels, one blocking 15 % of the power and the other blocking 50 % of the power. The curves in both cases vary qualitatively with different clipping even though the analysis in Section 7.2 suggests that they should only vary in magnitude.

This effect is believed to come from the fact that the alignment and mode of the auxiliary beam and the main beam are not perfectly matched as was assumed in the

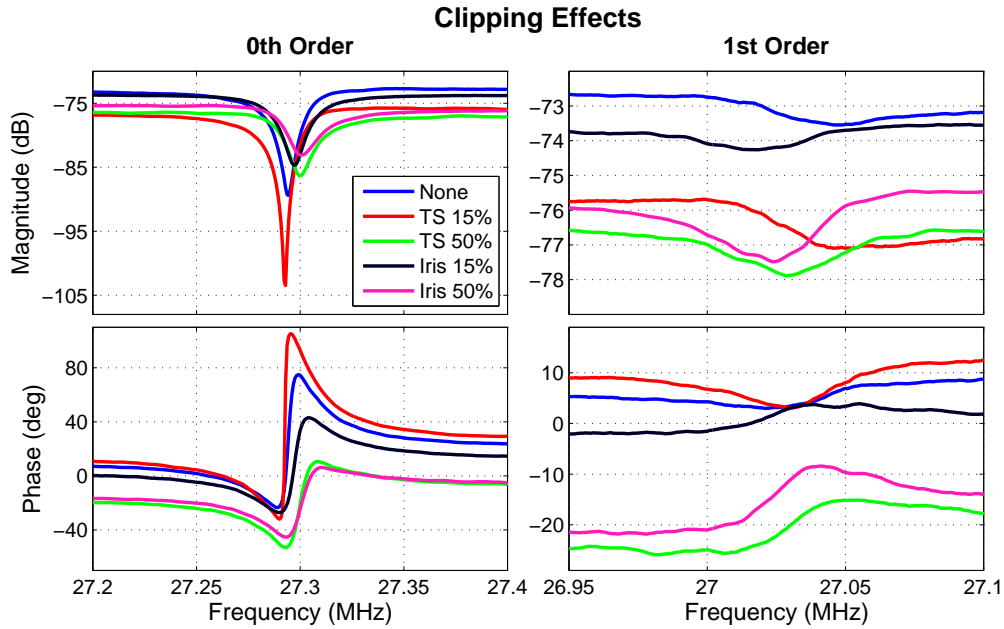


Figure 9-3: The measured signals at the zeroth and first order peaks for a number of different types of clipping. In the legend, TS refers to clipping with the dental toothscraper and Iris refers to clipping with an iris centered on the beam, and the percentage tells how much power is blocked. The average of the fitting results from this data is used to calculate the length and Gouy phase of the PRC cavity.

calculations of the previous sections. Recall that part of the results in Section 7.2 relied on the assumption that the coefficients, $C_{nm,n'm'}$, were real which was shown to be true for two beams with the same beam parameters. In the case of two beams with differing beam parameters these coefficients can be complex since the HOMs of the two beams have phase shifts relative to each other. In this case clipping more or less of the beam could introduce relative phase effects in the detection scheme.

9.3 PRC Cavity Parameters

This clipping issue introduces some ambiguity into measurements of the PRC parameters. For this reason, the parameters discussed in this section are calculated by first fitting curves with various different types of clipping (some of which are shown in Figure 9-3) and taking the average of the results as the actual value. The error is calculated by calculating the standard error of this mean which is given by

$$\sigma_{\mu} = \frac{\sigma}{\sqrt{N}}, \quad (9-12)$$

where σ_μ is the standard error of the mean, σ is the standard deviation of the sample population and $N = 7$ is the number of curves.

The fitting to each individual peak is done as described in Section 7.3 where the signal is rotated into the I phase and fit to a PDH-like signal. In particular, the calculation of the losses uses the known reflectivities of the power recycling mirror (PRM) and the input test mass (ITM)[51]. For the fits to the zeroth order mode the losses come out to be a little under 1% meaning that some other loss besides the mirror transmissivities is present in the power recycling cavity. There is some loss to the dark port due to contrast defect at the beamsplitter and some loss due to small but finite transmissivities of the other mirrors in the cavity, but not enough to give the measured value.

The free spectral range of the cavity is calculated from the measured location of the zeroth order mode by

$$FSR = \frac{f_0}{10.5}. \quad (9-13)$$

the factor of $\frac{1}{2}$ comes in because the cavity is locked on the sideband³ while the auxiliary laser is locked to the carrier. From this we can calculate the macroscopic *round-trip* length of the cavity from equation (8-9).

The other parameter of interest is the Gouy phase of the cavity. This parameter is significant because it is a measure of how stable the cavity is with respect to angular fluctuations. A larger round-trip Gouy phase implies that the power stored in the cavity is less susceptible to fluctuations in the alignment of the mirrors and vice versa. On the flip side, a more stable cavity has weaker control signals for automatically maintaining the alignment and vice versa. For these reasons, the design value of the PRC Gouy phase was chosen to be 29.9° in the cold state and 50° in the hot state when the thermal lens

³ The PRC cavity is intentionally designed to be anti-resonant for the sideband when the carrier is resonant and vice-versa so that they will both be simultaneously resonant when the carrier resonates in the arms.

Table 9-1: The cavity parameters measured with the PRC probe.

Parameter (Unit)	Value	Design	Difference
Fitted Values			
f_0 (Hz)	$27,296,998 \pm 86$	27,298,413	$-1,415 \pm 86$
f_1 (Hz)	$27,030,551 \pm 1246$	27,082,478	$-51,927 \pm 1246$
L_0 (%)	0.85 ± 0.01		
L_1 (%)	3.78 ± 0.14		
Derived Values			
Δf (Hz)	$266,447 \pm 1246$	215,935	$50,512 \pm 1246$
FSR (Hz)	$2,599,714 \pm 8$	2,599,849[7]	-135 ± 8
l_0 (m)	115.320 ± 0.001	115.313	0.007 ± 0.001
η (deg)	36.90 ± 0.17	29.9	7.00 ± 0.17

Note: The values and errors are calculated from the average values of curves with many different clipping styles, some examples of which are shown in Figure 9-3. The upper part of the table has the key fitted parameters, and the lower part of the table has the important parameters derived from these.

in the ITMs makes the cavity more stable. The round-trip Gouy phase of the PRC can be calculated from the difference between the frequencies of the zeroth order and the first order modes by

$$\eta = \frac{2\pi l_0 \Delta f}{c}, \quad (9-14)$$

where l_0 is the round-trip length just discussed, c is the speed of light, and Δf is the frequency difference between the zeroth and first order modes.

The measured parameters are shown in Table 9-1. The summary is that the cavity is very close to its proper length, but due to the sensitivity of the cavity mode to errors in the radii of curvature of the mirrors, especially PR2 and PR3, the round-trip Gouy phase is larger than expected in the cold state. This is not a significant problem because the actual design value of the PRC Gouy phase when operating in science mode is 50° , but

the plan had been to design the mode matching to be optimal at 18 W of input power to the interferometer. At 18 W a 25 μm thermal lens should develop in the substrate of the ITMs which will would have brought the cavity Gouy phase from 29.9° to 50° . The high cold state Gouy phase only means that the interferometer will now be optimally mode matched at ~ 12 W of input power and that more will be demanded of the TCS system.

9.4 A Thermal Error Signal

The design operating power of the Advanced LIGO interferometers is nominally 125 W of input power which builds up in the arm cavities to ~ 800 kW. Even with very low absorption optical components this significant power level causes non-negligible amounts of power to be deposited in the optics. The plan for Advanced LIGO is to use a number of different techniques to try and maintain the interferometer's state at its low power operating point.

9.4.1 TCS Sensors and Actuators

The TCS system will have the ability to actuate on the thermal state of the test masses in two ways. The ring heaters radiatively heat the test masses on the barrel near the anti-reflective (AR) face of the optic. This has two effects; it induces a thermal lens in the optic substrate through the photo-thermal effect which has the opposite sign to that of the lens induced by the main laser beam, and it deforms the radius of curvature of the highly reflective (HR) surface of the mirror which also has the opposite sign to the change induced by the main laser beam. In addition to the ring heaters, the TCS system includes a long wavelength laser (CO_2) which can be used to create a lens inside of the recycling cavity very near to the test masses. Depending on the pattern of CO_2 laser projected onto the compensation plate, the lens induced can be negative or positive. So, the general topology of the TCS actuation is that the ring heaters correct the radius of curvature of the mirrors and as a side benefit also remove some of the thermal lens in the substrate. The CO_2 laser is then used to compensate for the leftover thermal lens after correction by the ring heater.

The sensors for the TCS system are few however. The main sensors for the TCS system are the Hartmann sensors which are capable of imaging the total thermal state of each ITM separately. This includes the changes in the substrate and at the HR surface of the mirror, all wrapped up together. The other sensors for the TCS system are the power levels in the various parts of the interferometer because the amount of power coupled into the interferometer can change drastically as thermal effects start to become large. The difficulty with the Hartmann sensors is that they are not very strongly tied to the interferometer in that they can measure the thermal state but give no information about which state is the best. Examining the power levels in the various parts of the interferometer also has problems. Most notably, the ideal operating point is at the maximum where the derivative goes to zero so they can not be used as a proper error signal and can be affected by many changes in the state of the interferometer other than the thermal state.

In the rest of this section, we will discuss a way in which the PRC probe laser can be used to create an error signal for the TCS system which is strongly tied to the interferometer and dependent almost solely on the thermal state. The idea is to use the location of the first order mode resonance to track the round-trip Gouy phase in the PRC as a tracer of the thermal state.

9.4.2 The Error Signal

The location of the first order mode is proportional to the round-trip Gouy phase of the cavity which is highly sensitive to the curvature of the mirrors and lenses which make up the cavity. This fact was shown and exploited already in Section 8.4 to measure the absorption of the input mode cleaner. The goal in this section is to show how this signal develops as the thermal state of the power recycling cavity is intentionally altered with the central heating from the TCS CO₂ lasers.

The ITMs of the Livingston interferometer have a non-negligible cold state lens in the substrate which is due to inhomogeneities in the relatively thick mirrors⁴. This substrate inhomogeneity leads to an effective radius of curvature as seen through the substrate of 1939.9 m for ITMX and 1912.8 m for ITMY. This means that adding some positive lens to the ITMY compensation plate with the CO₂ laser helps to match the ITMs to each other which reduces the contrast defect and the losses in the power recycling cavity.

Figure 9-4 shows the evolution of the I phase component of the first and second order modes as this process happens. The legend shows the amount of time passed in minutes since 400 mW of CO₂ power is applied to the ITMY compensation plate. Notice that the error signal shows a shift in frequency as well as a reduction in losses. The error signal of the first order mode is what we will lock the auxiliary laser to in order to track the thermal state of the interferometer.

The shift in the frequency of this mode will be our tracer of the Gouy phase of the cavity and therefore the thermal state. The reduction in losses will make the optical gain of the error signal larger as ITMY warms up. This could cause havoc in a high bandwidth loop with low margins at the UGF. Since the thermal drifts are on very long timescales and the first order mode location is otherwise very stable, we will use a simple $\frac{1}{f}$ loop with a UGF around 1 Hz. In this state the control loop will be unconditionally stable up to the point where delays in the loop take away phase margin at the UGF which should not occur unless the gain increases by many orders of magnitude.

Figure 9-5 shows the result of fitting the data in Figure 9-4 as well as the data from the zeroth order resonance with the fitting procedures described in Section 7.3. As expected, the frequency of the first order and second order modes both shift towards more negative frequencies indicating that the power recycling cavity is getting more

⁴ All of the test masses have a 34 cm diameter and are 20 cm thick

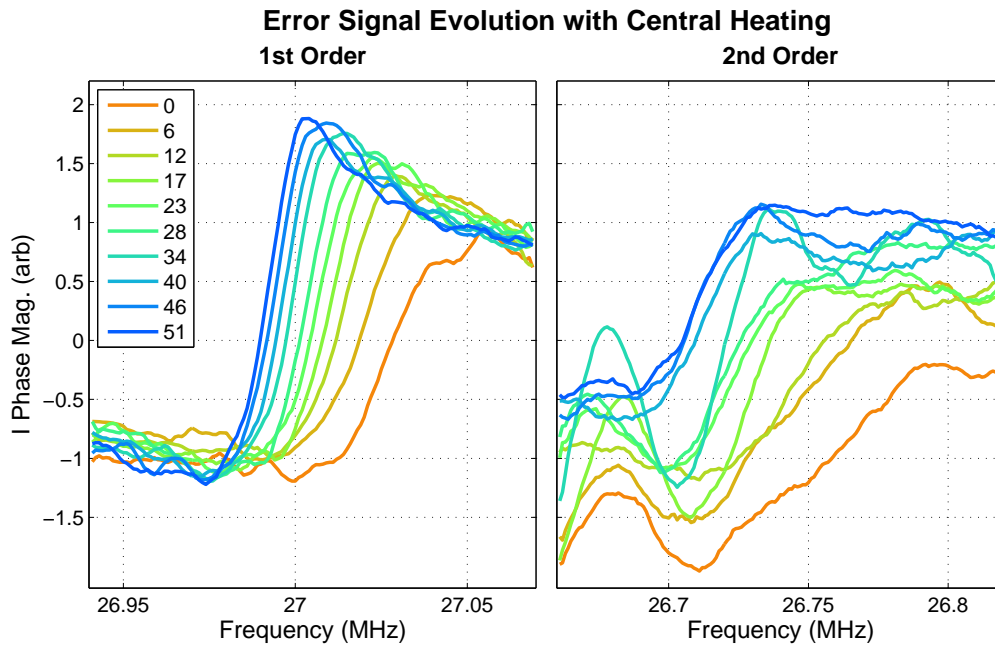


Figure 9-4: The evolution of the I phase component of the first and second order modes as central heating is applied to ITMY. The legend shows the number of minutes since the 400 mW beam was applied. Tracking of the first order mode by locking the auxiliary laser to this error signal provides an excellent tracer of the thermal state of the PRC. In both modes we can see a visible shift in frequency and reduction in losses as the ITMY radius of curvature is matched to that of ITMX.

stable. Also as expected the zeroth order mode does not shift since the FSR should not change with the thermal state of the interferometer. Somewhat surprisingly, the second order mode does not shift twice as much as the first order mode as one would expect from a simple parabolic change at the ITMs. This non-ideal shifting of the HOMs was also seen in the input mode cleaner and indicates the presence of non-parabolic changes in the power recycling cavity.

9.4.3 Continuous Tracking of the Thermal State

Continuous tracking of the first order mode requires only a small change to the setup of the PRC probe shown in Figure 9-1. Namely, the output of the network analyzer is replaced with a voltage controlled oscillator (VCO) so that the offset frequency between the main laser and the auxiliary laser can be controlled by applying a voltage from a DAC connected to the aLIGO digital control system. The readout of the RFPD




Figure 9-5: The frequency and loss of the lowest order modes as 400 mW of CO₂ power is applied to ITMY. This is slightly more than the amount of power needed to match the radii of curvature to ITMX. The actual data for the first and second order modes is shown in Figure 9-4. Notice that the losses in all modes reduce as the contrast defect at the beam splitter gets better and the mode size in the power recycling cavity gets smaller.

on ISCT1 is sent to an RF amplifier⁵ to increase the signal strength and then to a demodulator. The other port of the demodulator is fed a copy of the RF signal from the VCO against which the amplified signal of the RFPD is demodulated. After properly rotating the phase of the demodulation waveform, the output is precisely the I phase error signal shown on the left hand side of Figure 9-4. This demodulated error signal is fed to an ADC of the aLIGO digital control system.

The servo used to control the auxiliary laser is a simple $\frac{1}{f}$ loop which has a pole at zero, i.e. an integrator. The gain of the digital controller is chosen such that the unity gain frequency before heating ITMY is 1 Hz. As mentioned before, this is more than

⁵ The LF output of the Advanced LIGO diplexer amplifier was borrowed for this purpose.

enough bandwidth to follow the thermal changes in the interferometer and has low enough gain to avoid any issues as the optical gain changes with the thermal state. This setup locks the auxiliary laser to the first order mode of the power recycling cavity using the same PDH error signal that is typically used to lock a zeroth order mode to an optical cavity. The readout of the thermal state of the interferometer is the amount of voltage sent to the VCO in order to keep the first order mode of the auxiliary laser resonant in the cavity which can be converted to frequency shift of the first order mode and Gouy phase shift of the cavity via the VCO calibration.

After closing the control loop on the first order mode the same process as is shown in Figure 9-5 was begun. Namely, 400 mW of central heating was applied to ITMY and the PRC was allowed to warm up. After approximately an hour and a half the heating to ITMY was increased by 200 mW and 200 mW of central heating was also applied to ITMX, and this was shut off after ~ 45 min. The experiment was ended by an earthquake slightly before the three hour mark.

Figure 9-6 shows the results of the experiment. Notice that the continuously tracked location of the first order resonance provides an excellent tracer of the thermal state of the cavity. Notice also that the shift during the first hour is very similar to that which is shown in Figure 9-5 as expected. The buildup inside of the power recycling cavity and the contrast defect also go in the direction expected from matching the radii of curvature of the two test masses. There was a brief lock loss of the PRMI at 1.9 hours which caused the probe servo to also lose lock. Additionally, there was a lock loss of the probe servo for an unknown reason at 2.2 hours.

9.4.4 Conclusion

The thermal error signal developed in this section could be a useful tool for sensing the thermal state of the interferometer. For the power recycling cavity there are effectively two important thermal degrees of freedom; the common lensing by the ITMs and the differential lensing by the ITMs. Figure 9-6 shows that this error signal




Figure 9-6: The results of tracking the thermal state of the PRC by locking the auxiliary laser to the resonance of the first order mode. Notice that the buildup increases and the contrast defect decreases as the effective ROC of ITMY is brought to match that of ITMX. At 1.7 hrs common mode heating is applied to both ITMX and ITMY. Throughout the experiment the location of the first order mode clearly follows the thermal state of the interferometer both in common and differential changes between the two ITMs. There was a brief lock loss of the PRMI at 1.9 hours and of the first order probe servo at 2.2 hours.

is sensitive to both, but being able to control both requires two sensors. Hence, this error signal could be used in conjunction with the Hartmann sensors to control both the differential and common thermal degrees of freedom in the power recycling cavity.

One of the largest problems with thermal issues in the aLIGO interferometers is due to instabilities developing in the angular control loops which keep the interferometer aligned. Controlling the round-trip Gouy phase of the cavities should alleviate many of these thermal problems with the angular control loops which depend heavily on the cavity Gouy phase. Similar systems could also be used to detect the thermal states of

the arms and signal recycling cavity, though they would require that another laser be injected into each of the cavities.

CHAPTER 10 CONCLUSION AND OUTLOOK

The outlook for the Advanced LIGO (aLIGO) detectors and the field of gravitational wave astronomy looks very promising at the moment. As of this writing the Livingston (LLO) detector is fully installed and the Hanford detector will be so within the next few months. The Livingston detector has been brought to and held at its operating point for time periods of up to a few hours by a nearly deterministic, automated process.¹ And the process of noise hunting in the LLO detector has begun, and is progressing rapidly.

Figure 10-1 shows the fruits of this rapid commissioning progress. The sensitivity of the interferometer is already better at all frequencies except between ~ 30 Hz and ~ 300 Hz than it was during the sixth science run of initial LIGO (called eLIGO). Notice in particular, the marked improvement in low frequency noise due to the improved seismic isolation systems of the aLIGO detectors. The limiting noise in all frequency bands is still the subject of active investigation, but the prospect of reaching 100 Mpc for NS/NS inspirals seems just around the corner.

10.1 The Input Optics

That the input optics (IO) have been contributing to this marked success is visible in the fact that they have been virtually forgotten about by the commissioners working on noise hunting of the full interferometer. This means that none of the components of the input optics have been found to be contributing significantly to the measured noise thus far. As the noise hunting progresses it is likely that some unexpected problems will be found, but the fact that the IO are operating so well thus far is a testament to an excellent design. The details of the key components of the input optics were discussed in Part III

¹ In contrast, the process of locking the initial LIGO interferometers involved a number of steps during which stochastic seismic motion brought parts onto resonance and a brief period of relinquished control as the carrier was brought onto resonance in the arms.




Figure 10-1: Comparison between the current noise of the aLIGO interferometer at LLO and the best eLIGO sensitivity. Notice the marked improvement at low frequencies due to the improved seismic isolation. The limiting noise sources in the other regions are the subject of active investigations. Note that this is a preliminary curve and the calibration may not be completely accurate.

of this dissertation, but we will briefly recap the identified areas where improvement may be necessary and describe some strategies for those improvements.

The upper sideband (at 45.5 MHz) of the electro-optic modulator is able to reach modulation depths of 0.4 which looks to be sufficient for readout and control of the interferometer at the moment. If a deeper modulation depth is necessary the resonant circuitry will need to be tested with the significant drive levels required to reach these deep modulation depths. In addition, the residual amplitude modulation was found to be low enough at times, but temperature drift can cause it to drift above the 10^{-4} AM/PM ratio required[47]. The fix for this is temperature stabilization of the crystal in the EOM which requires relatively minor modifications of the EOM housing.

Out-of-vacuum high power testing of the Faraday isolator (FI) indicates that it will meet all of the requirements at power levels up to the full aLIGO input power of

165 W except for the marginally out of spec thermal lensing. The in-vacuum isolation performance at low power was measured and found to be within spec, but it more than 10 dB lower than was measured out-of-vacuum. In addition, the thermal performance of the FI has yet to be measured in-vacuum due to the limited availability of high input power. Thermal effects are exaggerated in vacuum due to the lack of convective cooling so it is possible that some issues could arise. Thermal lensing issues can be fixed by replacement of the DKDP element with a thinner crystal; an operation which requires very little down time except for the vacuum intrusion. Thermal induced isolation ratio issues can be largely dealt with by adjustment of the motorized half wave plate (HWP) which will trade off isolation ratio at the cost of dumping more power in transmission of the FI.

The input mode cleaner (IMC) is operating very well, and appears to be ready to reach the full aLIGO sensitivity. The control system is robust and lock acquisition is rapid. The scattering and absorption are slightly higher than expected, but are not high enough to cause problems with the IMC itself. The scattered light could become an issue as the full interferometers reach lower and lower noise, but the remoteness from the dark port and significant level of baffling in the IO chain leave the expectation that scattering at other points in the interferometer will be a headache long before the scattering in the IMC. The frequency noise of the IMC is well understood except for the region between ~ 5 Hz and ~ 80 Hz, and modeling suggests that the understood noises will not be a problem for the full aLIGO sensitivity. The noise in the deficit region may turn out to create frequency noise which is too high for the full sensitivity, and hunting of this noise could take some serious effort. If this is the case some clues as to its source are provided in Section [6.5.2](#).

10.2 Eigenspectra Measurements

Part [IV](#) of this dissertation discussed a relatively new class of measurements on resonant optical interferometers which involve using a probe beam to characterize their

eigenspectra. This measurement class was applied as a characterization technique of the input mode cleaner and the power recycling cavity (PRC) in which it was used to measure the typical parameters which characterize a resonant optical cavity (length, Gouy phase, etc.). In addition to this simple characterization the eigenspectra probe technique was used to measure thermal effects in both the IMC and PRC by measuring the shift in the first order modes (a tracer of the cavity Gouy phase) with applied heating. In the case of the IMC this shift was due to absorption at the surface of the mirrors while in the PRC this shift was due to application of CO₂ heating at the ITMs.

There are a number of other ways in which this class of measurements might be applied to the Advanced LIGO interferometers both for characterization during commissioning and sensing/control during science mode. The vision put forth is one in which the interferometer is probed with lasers from all entry points giving simultaneous but independent sensitivity to all of the resonant cavities.

10.2.1 Arm Cavity Absorption

The technique described in Section 8.4 which is used to measure the absorption in the input mode cleaner could be straightforwardly applied to the arms. In order to do so a probe laser would need to be setup which was capable of delivering enough power to the arms to see its resonances. This might be able to be accomplished with a probe laser injected at the dark port where a Faraday isolator allows injection of a probe beam with very low loss or by probing the arm cavities from the end stations. Probing from the end stations might have a dynamic range issue if the probe frequency is near 1064 nm since the ETM reflectivity is so high at that wavelength. The measurement could be done with the green lasers ($\lambda = 532$ nm) used for arm length stabilization, but the

resolving power of the measurement scales with the finesse of the cavity and the current state of the ETM transmissivity² might not be sensitive enough.

10.2.2 Probing the Full Interferometer in the Gravitational Wave Band

During initial LIGO there was some evidence of higher order mode generation in the interferometer causing readout noise beyond the simple shot noise contribution[8]. These effects are presumably greatly reduced by the output mode cleaner, but it would nevertheless be useful to have an understanding of the low frequency resonance structure of the full interferometer. Here low frequency means in and just above the gravitational wave band; 0 Hz – 500 kHz.

The full Advanced LIGO gravitational wave interferometers consist of four coupled cavities as discussed in Section 1.3. The coupled cavities force the probe laser to be shifted from the carrier only by the frequency of interest since coupled cavities do not have a well defined free spectral range³. This could prove to be a very useful measurement since calculation of the higher order mode spectrum for such a complex instrument is difficult with simulations and essentially impossible on paper. Measurement of the low frequency resonance structure should, however, be fairly straightforward.

Probing this low frequency resonance structure can be accomplished with a setup very similar to the one described in Chapter 9 with one major modification. The probe laser used in that setup can not go down to DC due to the phase locking topology. By adding an AOM or changing the phase locking topology to make the offset frequency tunable in the low frequency band the rest of the setup could stay the same. By reading

² The currently installed ETMs have a reflectivity on the order of 50 % for 532 nm light which gives the arm cavities a finesse on the order of ~ 10 for this wavelength. There are some plans to correct this in the near future though.

³ Recall that the PRC measurement probed the eigenspectra by measuring around the tenth FSR and the IMC by measuring around the fifth.

out the eigenspectra at the different interferometer ports (arm transmission, dark port, reflection, etc) a full map of the low frequency resonance structure in each of the cavities could be made. In addition, problems which are thought to be related to higher order modes can be confirmed by using the probe laser to intentionally ring up or cool the resonances.

10.2.3 Thermal State Sensing

The TCS error signal described in Section 9.4 could be extended to the arm cavities. By injecting probe lasers which could see the two arms independently, either from the ends or from the dark port, and locking them to the first order mode of these cavities, an error signal which is capable of reading out the Gouy phase of the two cavities is attained. These error signals could be used in conjunction with the TCS actuation systems to maintain the interferometer in a given thermal state by controlling the Gouy phase structure. Controlling the thermal state in this way has the advantage of fixing the higher order mode resonance structure of the interferometer as well as the optical gain of the angular control signals.

10.2.4 Parametric Instability Cooling

Parametric instabilities are caused in the arm cavities of gravitational wave detectors when one of the mechanical modes of the test masses has substantial spatial overlap with a higher order mode whose offset frequency from the carrier is near the resonant frequency of the mechanical mode (within the Q of the HOM). [12][11][99][41] These instabilities are typically out of the sensitive band of gravitational wave detectors, but can be low enough to saturate low noise readout electronics.

If a particular acoustic mode/HOM pair is identified as problematic during the Advanced LIGO commissioning period and can not be dealt with by other means (such as tuning the test mass ROCs with the ring heaters), then it is possible to damp the mode with optical feedback from a probe laser. The idea is to lock a probe laser to the higher order mode which is coupling to the test mass acoustic mode. Amplitude

modulation of the probe laser can then be used as narrow band feedback to the acoustic mode of the test mass. In order not to inject extra light, the injection laser can have all of its power put into the higher order mode of interest through use of a phase plate and mode cleaner[34].

APPENDIX: RAY MATRICES

This appendix has a list of 4x4 matrices, the properties of which are discussed in Section 2.3. The standard 2x2 ABCD matrices discussed in Section 2.2 can be plucked out of the 4x4 matrices by taking the first two columns and rows. The table is essentially a reprint of the table in the appendix of Shaomin's review article on matrix methods in misaligned optical systems[82].

Recall from Section 2.3.2 that misalignment coordinates of a system with a single center of curvature are ambiguous excepting that they must satisfy

$$\frac{\rho}{R} + \phi = C, \quad (\text{A-1})$$

where ρ , ϕ are the displacement and tilt of the axis of the ABCD system with respect to the optical axis, R is the radius of curvature, and C is an arbitrary constant. This does not affect the form of the 4x4 matrices, but must be kept in mind in order to interpret them correctly.

Table A-1: A table of 4x4 ray matrices.

Name	Description	Matrix
Translation	Translation by a distance d .	$\begin{bmatrix} 1 & d & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
Plane Interface	A plane interface from index n_i to index n_f .	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & n_i/n_f & 0 & \phi(1 - n_i/n_f) \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
Spherical Interface	A spherical interface of radius of curvature R from index n_i to index n_f .	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -\left(1 - \frac{n_i}{n_f}\right)/R & \frac{n_i}{n_f} & \rho\left(1 - \frac{n_i}{n_f}\right)/R & \phi\left(1 - \frac{n_i}{n_f}\right) \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
Thin Lens	A thin lens of focal length f .	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -\frac{1}{f} & 1 & \frac{\rho}{f} & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
Spherical Mirror	A spherical mirror of radius R .	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -\frac{2}{R} & 1 & \frac{2\rho}{R} & -2\phi \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$

Note: The coordinates ρ and ϕ are used to denote the displacement and tilt, respectively, of the axis of the ABCD system with respect to the optical axis of the system.

REFERENCES

- [1] Complex beam parameter and ABCD law for non-gaussian and nonspherical light beams. 31.
- [2] R. Abbott, R. Adhikari, G. Allen, S. Cowley, E. Daw, D. DeBra, J. Giaime, G. Hammond, M. Hammond, and C. Hardham. Seismic isolation for advanced LIGO. *Classical and Quantum Gravity*, 19(7):1591, 2002. Available from: <http://iopscience.iop.org/0264-9381/19/7/349>.
- [3] F. Acernese, P. Amico, N. Arnaud, D. Babusci, R. Barill, F. Barone, L. Barsotti, M. Barsuglia, F. Beauville, and M. A. Bizouard. Results of the virgo central interferometer commissioning. *Classical and Quantum Gravity*, 21(5):S395, 2004. Available from: <http://iopscience.iop.org/0264-9381/21/5/003>.
- [4] D. Z. Anderson. Alignment of resonant optical cavities. *Applied Optics*, 23(17):2944–2949, 1984. Available from: <http://www.opticsinfobase.org/abstract.cfm?id=27768>.
- [5] K. Arai, C. Bond, D. Brown, L. Carbone, S. Doravari, A. Freise, P. Fulda, and K. Kokeyama. Finesse simulation for the alignment control signals of the aligo input mode cleaner. *LIGO DCC*, (T1300074), 2013. Available from: <https://dcc.ligo.org/LIGO-T1300074>.
- [6] Koji Arai. On the accumulated round-trip gouy phase shift for a general optical cavity, March 2013. Available from: <https://dcc.ligo.org/LIGO-T1300189>.
- [7] M. Arain, R. Martin, G. Mueller, D. Reitze, D. Tanner, and L. Williams. Advanced ligo input optics final design document. *LIGO DCC*, (T0900386), 2011. Available from: <https://dcc.ligo.org/T0900386-v7>.
- [8] S. W. Ballmer. *LIGO interferometer operating at design sensitivity with application to gravitational radiometry*. PhD thesis, Massachusetts Institute of Technology, 2006. Available from: <http://18.7.29.232/handle/1721.1/36396>.
- [9] M. Barton. Models of the advanced ligo suspensions in mathematica. *LIGO DCC*, (T020205), 2010. Available from: <https://dcc.ligo.org/LIGO-T020205/public>.
- [10] P. Baues. The connection of geometrical optics with the propagation of gaussian beams and the theory of optical resonators. *Opto-electronics*, 1(2):103–118, 1969. Available from: <http://link.springer.com/article/10.1007/BF01418102>.
- [11] V. B. Braginsky, S. E. Strigin, and S. P. Vyatchanin. Analysis of parametric oscillatory instability in power recycled LIGO interferometer. *Physics Letters A*, 305(3):11–124, 2002. Available from: <http://www.sciencedirect.com/science/article/pii/S0375960102013579>.

- [12] V. B. Braginsky, S. P. Vyatchanin, and S. E. Strigin. Parametric oscillatory instability in fabry-perot interferometer. *Physics Letters A*, 287:331–338, September 2001.
- [13] H. B. Callen and T. A. Welton. Irreversibility and generalized noise. *Physical Review*, 83(1):34–40, 1951. Available from: <http://lptms.u-psud.fr/membres/trizac/Ens/M2MQPL/CallenWelton.pdf>.
- [14] J. B. Camp, H. Yamamoto, S. E. Whitcomb, and D. E. McClelland. Analysis of light noise sources in a recycled michelson interferometer with Fabry–Perot arms. *JOSA A*, 17(1):120–128, 2000. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=JOSAA-17-1-120>.
- [15] L. Carbone, S. M. Aston, R. M. Cutler, A. Freise, J. Greenhalgh, J. Heefner, D. Hoyland, N. A. Lockerbie, D. Lodhia, N. A. Robertson, C. C. Speake, K. A. Strain, and A. Vecchio. Sensors and actuators for the advanced LIGO mirror suspensions. *Classical and Quantum Gravity*, 29(11):115005, June 2012. Available from: <http://stacks.iop.org/0264-9381/29/i=11/a=115005?key=crossref.012439014e45d280b1722ea3e59e0d5b>.
- [16] C. M. Caves. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Physical Review Letters*, 45(2):75, 1980. Available from: <http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.45.75>.
- [17] C. M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, 23(8):1693, 1981. Available from: <http://journals.aps.org/prd/abstract/10.1103/PhysRevD.23.1693>.
- [18] Y. Chen. Macroscopic quantum mechanics: theory and experimental concepts of optomechanics. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 46(10):104001, May 2013. Available from: <http://stacks.iop.org/0953-4075/46/i=10/a=104001?key=crossref.f6fb280b25ca177dd9c80e4daf9cd4e0>.
- [19] E. L. Church, H. A. Jenkinson, and J. M. Zavada. Relationship between surface scattering and microtopographic features. *Optical Engineering*, 18(2):182125, 1979. Available from: <http://opticalengineering.spiedigitallibrary.org/article.aspx?articleid=1221802>.
- [20] E. Constantin and N. H. Pavel. Green function of the laplacian for the neumann problem in r_+^n . *Libertas Mathematica*, 30:57–69, 2010. Available from: <http://system.lm-ns.org/index.php/lm/issue/view/57>.
- [21] R. DeRosa, J. C. Driggers, D. Atkinson, H. Miao, V. Frolov, M. Landry, J. A. Giaime, and R. X. Adhikari. Global feed-forward vibration isolation in a km scale interferometer. *Classical and Quantum Gravity*, 29(21):215008, November 2012. Available from: <http://stacks.iop.org/0264-9381/29/i=21/a=215008?key=crossref.bef82431b88497444ddb51b3ee94f2aa>.

- [22] K. L. Dooley, M. A. Arain, D. Feldbaum, V. V. Frolov, M. Heintze, D. Hoak, E. A. Khazanov, A. Lucianetti, R. M. Martin, and G. Mueller. Characterization of thermal effects in the enhanced LIGO input optics. *arXiv preprint arXiv:1112.1737*, 2011. Available from: <http://arxiv.org/abs/1112.1737>.
- [23] K. L. Dooley, M. A. Arain, D. Feldbaum, V. V. Frolov, M. Heintze, D. Hoak, E. A. Khazanov, A. Lucianetti, R. M. Martin, and G. Mueller. Thermal effects in the input optics of the enhanced laser interferometer gravitational-wave observatory interferometers. *Review of Scientific Instruments*, 83(3), March 2012. Available from: <http://scitation.aip.org/content/aip/journal/rsi/83/3/10.1063/1.3695405>.
- [24] R. W. P. Drever, John L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(2):97–105, 1983. Available from: <http://link.springer.com/article/10.1007/BF00702605>.
- [25] J. C. Driggers, J. Harms, and R. X. Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Physical Review D*, 86(10), November 2012. Available from: <http://link.aps.org/doi/10.1103/PhysRevD.86.102001>.
- [26] A. S. Eddington. The total eclipse of 1919 May 29 and the influence of gravitation on light. *The Observatory*, 42:119–122, March 1919.
- [27] M. F. Erden and H. M. Ozaktas. Accumulated gouy phase shift in gaussian beam propagation through first-order optical systems. *JOSA A*, 14(9):2190–2194, 1997. Available from: <http://www.opticsinfobase.org/abstract.cfm?id=1903>.
- [28] M. Evans and K. Arai. aligo broadband photodetector. *LIGO DCC*, (T1100467), 2013. Available from: <https://dcc.ligo.org/T1100467/public>.
- [29] G. Fowles. *Introduction to modern optics*. Dover Publications, New York, 1989.
- [30] A. Freise. Finesse: Frequency domain interferometer simulation software. 2013. Available from: <http://www.gwoptics.org/finesse/>.
- [31] A. Freise, D. Brown, and C. Bond. Finesse: Frequency domain interferometer simulation software. *arXiv:1306.2973 [physics.comp-ph]*, 2013. Available from: <http://arxiv.org/abs/1306.2973>.
- [32] P. Fritschel. Interferometer sensing and control design requirements. *LIGO DCC*, (T070236). Available from: <https://dcc.ligo.org/LIGO-T070236>.
- [33] P. Fritschel, N. Mavalvala, D. Shoemaker, D. Sigg, M. Zucker, and G. Gonzalez. Alignment of an interferometric gravitational wave detector. *Applied Optics*, 37(28):6734–6747, 1998. Available from: <http://www.opticsinfobase.org/abstract.cfm?id=61206>.

- [34] P. Fulda. *Precision interferometry in a new shape : higher-order Laguerre-Gauss modes for gravitational wave detection*. PhD thesis, University of Birmingham, 2013.
- [35] V.T. Grinchenko and G.L. Komissarova. Analysis of frequency spectrum and vibrational modes of long cylinders. *Soviet Applied Mechanics*, 16(1):1–5, 1980. Available from: <http://dx.doi.org/10.1007/BF00884604>.
- [36] P. Hello and J. Y. Vinet. Analytical models of thermal aberrations in massive mirrors heated by high power laser beams. *Journal de Physique*, 51(12):1267–1282, 1990. Available from: http://jphys.journaldephysique.org/articles/jphys/abs/1990/12/jphys_1990__51_12_1267_0/jphys_1990__51_12_1267_0.html.
- [37] H. Hezari and S. Zelditch. Inverse spectral problem for analytic $(z/2z)^n$ -symmetric domains in r^n . *arXiv:0902.1373 [math.AP]*, 2010. Available from: <http://arxiv.org/abs/0902.1373>.
- [38] R. A. Hulse and J.H. Taylor. Discovery of a pulsar in a binary system. *Astrophysical Journal*, 195:L51–L53, January 1975.
- [39] J. R. Hutchinson. Axisymmetric vibrations of a free finite-length rod. *The Journal of the Acoustical Society of America*, 51(1B):233–240, 2005. Available from: <http://scitation.aip.org/content/asa/journal/jasa/51/1B/10.1121/1.1912835>.
- [40] J. D. Jackson. *Classical electrodynamics*. Wiley, New York, 1999.
- [41] L. Ju, S. Gras, C. Zhao, J. Degallaix, and D.G. Blair. Multiple modes contributions to parametric instabilities in advanced laser interferometer gravitational wave detectors. *Physics Letters A*, 354(5-6):360–365, June 2006. Available from: <http://linkinghub.elsevier.com/retrieve/pii/S0375960106001952>.
- [42] Mark Kac. Can one hear the shape of a drum? *American Mathematical Monthly*, 73(4):1–23, April 1966.
- [43] E. Khazanov, N.F. Andreev, A. Mal'shakov, O. Palashov, A.K. Poteomkin, A. Sergeev, A.A. Shaykin, V. Zelenogorsky, I.A. Ivanov, R. Amin, G. Mueller, D.B. Tanner, and D.H. Reitze. Compensation of thermally induced modal distortions in faraday isolators. *IEEE Journal of Quantum Electronics*, 40(10):1500–1510, October 2004. Available from: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1337031>.
- [44] E. A. Khazanov, O. V. Kulagin, S. Yoshida, D. B. Tanner, and D. H. Reitze. Investigation of self-induced depolarization of laser radiation in terbium gallium garnet. *Quantum Electronics, IEEE Journal of*, 35(8):1116–1122, 1999. Available from: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=777210.

- [45] H. Kogelnik and T. Li. Laser beams and resonators. *Proceedings of the IEEE*, 54(10):1312–1329, 1966. Available from: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1447049.
- [46] Hg Kogelnik. Imaging of optical modes – resonators with internal lenses. *The Bell System Technical Journal*, 44(3):455–494, March 1965. Available from: <https://archive.org/details/bstj44-3-455>.
- [47] Keiko Kokeyama, Kiwamu Izumi, William Z. Korth, Nicolas Smith-Lefebvre, Koji Arai, and Rana X. Adhikari. Residual amplitude modulation in interferometric gravitational wave detectors. *Journal of the Optical Society of America A*, 31(1):81, January 2014. Available from: <http://www.opticsinfobase.org/abstract.cfm?URI=josaa-31-1-81>.
- [48] P. Kwee, C. Bogan, K. Danzmann, M. Frede, H. Kim, P. King, J. Pld, O. Puncken, R. L. Savage, and F. Seifert. Stabilized high-power laser system for the gravitational wave detector advanced LIGO. *Optics express*, 20(10):10617–10634, 2012. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=oe-20-10-10617>.
- [49] Patrick Kwee, Benno Willke, and Karsten Danzmann. Shot-noise-limited laser power stabilization with a high-power photodiode array. *Optics letters*, 34(19):2912–2914, 2009. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=ol-34-19-2912>.
- [50] Liufeng Li, Fang Liu, Chun Wang, and Lisheng Chen. Measurement and control of residual amplitude modulation in optical phase modulation. *Review of Scientific Instruments*, 83(4):043111, 2012. Available from: <http://scitation.aip.org/content/aip/journal/rsi/83/4/10.1063/1.4704084>.
- [51] LIGO. Ligo optics. 2014. Available from: <https://galaxy.ligo.caltech.edu/optics/>.
- [52] R. Loudon. *The quantum theory of light*. Oxford University Press, Oxford New York, 2000.
- [53] T. T. Lyons, M.W. Regehr, and F. J. Raab. Shot noise in gravitational wave detectors with fabry-perot arms. *Applied Optics*, 39(36):6761–6770, December 2000. Available from: <http://www.opticsinfobase.org/ao/abstract.cfm?uri=ao-39-36-6761>.
- [54] F. Magana-Sandoval, R. X. Adhikari, V. V. Frolov, J. Harms, J. Lee, S. Sankar, P. R. Saulson, and J. R. Smith. Large-angle scattered light measurements for quantum-noise filter cavity design studies. *JOSA A*, 29(8):1722–1727, 2012. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=josaa-29-8-1722>.

- [55] R. Martin, L. Williams, Muzammil Arain, Volker Quetschke, Guido Mueller, David Reitze, and David Tanner. Is an in-vacuum faraday isolator really required in the input optics for advanced ligo? Available from: <https://dcc.ligo.org/LIGO-G0900173>.
- [56] B. J. Meers. Recycling in laser-interferometric gravitational-wave detectors. *Physical Review D*, 38(8):2317, 1988. Available from: http://prd.aps.org/abstract/PRD/v38/i8/p2317_1.
- [57] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34(203):333–345, November 1887.
- [58] P. Milonni. *Lasers*. Wiley, New York, 1988.
- [59] E. Morrison, B. J. Meers, D. I. Robertson, and H. Ward. Automatic alignment of optical interferometers. *Applied optics*, 33(22):5041–5049, 1994. Available from: <http://www.opticsinfobase.org/abstract.cfm?id=42086>.
- [60] C. L. Mueller. Expectations for the carm control loops in the advanced ligo detectors. *LIGO DCC*, (T1300409), 2013. Available from: <https://dcc.ligo.org/LIGO-T1300409>.
- [61] C. L. Mueller. Llo alog 5903. *Ligo Livingston aLog*, 2013. Available from: <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=5903>.
- [62] C. L. Mueller. Llo alog 5957. *Ligo Livingston aLog*, 2013. Available from: <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=5957>.
- [63] C. L. Mueller. Llo alog 8608. *Ligo Livingston aLog*, 2013. Available from: <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=8608>.
- [64] C. L. Mueller, D. Feldbaum, and R. Martin. Testing results from the first aligo faraday isolator. *LIGO DCC*, (T1200101), 2013. Available from: <https://dcc.ligo.org/LIGO-T1200101>.
- [65] G. Mueller and R. Amin. Eo-modulators for advanced ligo part i, mar 2003. Available from: <https://dcc.ligo.org/LIGO-T020025/public>.
- [66] Harry Nyquist. Thermal agitation of electric charge in conductors. *Physical review*, 32(1):110–113, 1928.
- [67] O. V. Palashov, D. S. Zheleznov, A. V. Voitovich, V. V. Zelenogorsky, E. E. Kamenetsky, E. A. Khazanov, R. M. Martin, K. L. Dooley, L. Williams, and A. Lucianetti. High-vacuum-compatible high-power faraday isolators for gravitational-wave interferometers. *JOSA B*, 29(7):1784–1792, 2012. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=josab-29-7-1784>.

- [68] J. H. Poeld. Reference cavity temperature control. *LIGO DCC*, (T140015), 2014. Available from: <https://dcc.ligo.org/LIGO-T140015>.
- [69] S. Privitera. Sensitivity achieved by the ligo and virgo gravitational wave detectors during ligo's sixth and virgo's second and third science runs. *LIGO DCC*, (T1100338), 2012. Available from: <https://dcc.ligo.org/LIGO-T1100338>.
- [70] M. Rakhmanov. *Dynamics of laser interferometric gravitational wave detectors*. PhD thesis, California Institute of Technology, 2000. Available from: <http://docuserver.ligo.caltech.edu/docs/public/P/P000002-00.pdf>.
- [71] S. N. Rasband. Resonant vibrations of free cylinders and disks. *Journal of the Acoustic Society of America*, 57(4):899–905, April 1975.
- [72] M. W. Regehr. *Signal extraction and control for an interferometric gravitational wave detector*. PhD thesis, California Institute of Technology, 1994. Available from: <http://www.ligo.caltech.edu/~rana/docs/Theses/Thesis-Regehr.pdf>.
- [73] N. A. Robertson, M. Barton, M. Meyer, J. Romie, C. Torrie, and J. Kissel. Ham small triple suspension final design document. *LIGO DCC*, (T0900435), 2009. Available from: <https://dcc.ligo.org/LIGO-T0900435>.
- [74] N. A. Robertson, G. Cagnoli, D. R. M. Crooks, E. Elliffe, J. E. Faller, P. Fritschel, and Go. Quadruple suspension design for advanced LIGO. *Classical and Quantum gravity*, 19(15). Available from: <http://iopscience.iop.org/0264-9381/19/15/311>.
- [75] B. Saleh. *Fundamentals of photonics*. Wiley-Interscience, Hoboken, N.J, 2007.
- [76] J. Sathian and E. Jaatinen. Reducing residual amplitude modulation in electro-optic phase modulators by erasing photorefractive scatter. *Optics Express*, 21(10):12309, May 2013. Available from: <http://www.opticsinfobase.org/abstract.cfm?URI=oe-21-10-12309>.
- [77] P. N. Schatz and A. J. McCaffery. The faraday effect. *Quarterly Reviews, Chemical Society*, 23(4):552–584, 1969. Available from: <http://pubs.rsc.org/en/content/articlepdf/1969/qr/qr9692300552>.
- [78] Arthur L. Schawlow and Charles H. Townes. Infrared and optical masers. *Physical Review*, 112(6):1940, 1958. Available from: <http://journals.aps.org/pr/abstract/10.1103/PhysRev.112.1940>.
- [79] S. Schroder, A. Duparre, L. Coriand, A. Tunnermann, D. H. Penalver, and J. E. Harvey. Modeling of light scattering in different regimes of surface roughness. *Optics express*, 19(10):9820–9835, 2011. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=oe-19-10-9820>.
- [80] B. Schutz. *A first course in general relativity*. Cambridge University Press, Cambridge New York, 2009.

- [81] F. Seifert, P. Kwee, M. Heurs, B. Willke, and K. Danzmann. Laser power stabilization for second-generation gravitational wave detectors. *Optics letters*, 31(13):2000–2002, 2006. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=ol-31-13-2000>.
- [82] W. Shaomin. Matrix methods in treating decentred optical systems. *Optical and quantum electronics*, 17(1):1–14, 1985. Available from: <http://link.springer.com/article/10.1007/BF00619988>.
- [83] K. Shiraishi, F. Tajima, and S. Kawakami. Compact faraday rotator for an optical isolator using magnets arranged with alternating polarities. *Optics letters*, 11(2):82–84, 1986. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=ol-11-2-82>.
- [84] D. Shoemaker. Advanced ligo anticipated sensitivity curves. *LIGO DCC*, (T0900288), 2010. Available from: <https://dcc.ligo.org/LIGO-T0900288>.
- [85] A. E. Siegman. *Lasers*. University Science Books, Mill Valley, Calif, 1986.
- [86] I. Snetkov, I. Mukhin, O. Palashov, and E. Khazanov. Compensation of thermally induced depolarization in faraday isolators for high average power lasers. *Optics express*, 19(7):6366–6376, 2011. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=oe-19-7-6366>.
- [87] S. Sridhar and A. Kudrolli. Experiments on not hearing the shape of drums. *Physical Review Letters*, 72(14):2175, 1994. Available from: http://prl.aps.org/abstract/PRL/v72/i14/p2175_1.
- [88] J. P. Tache. Derivation of ABCD law for laguerre-gaussian beams JP tach. *Applied optics*, 26(14/15), 1987. Available from: <http://www.opticsinfobase.org/viewmedia.cfm?id=61725&seq=0>.
- [89] D. Tanner, G. Mueller, and D. Reitze. Input optics subsystem design requirements document., September 2009. Available from: <https://dcc.ligo.org/LIGO-T020020>.
- [90] The LIGO Collaboration. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nature Physics*, 7(12):962–965, September 2011. Available from: <http://www.nature.com/doifinder/10.1038/nphys2083>.
- [91] The LIGO Collaboration. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics*, 7(8):613–619, July 2013. Available from: <http://www.nature.com/doifinder/10.1038/nphoton.2013.177>.
- [92] K. S. Thorne. The scientific case for advanced ligo interferometers. *LIGO DCC*, (P000024), 2001. Available from: <https://dcc.ligo.org/LIGO-P000024/public>.

- [93] UF LIGO Group. Upgrading the input optics for high power operation, Jan 2006. Available from: <https://dcc.ligo.org/LIGO-T060267>.
- [94] R. Ward. *Length Sensing and Control of a Prototype Advanced Interferometric Gravitational Wave Interferometer*. PhD thesis, California Institute of Technology, February 2010.
- [95] Edward A. Whittaker, Manfred Gehrtz, and Gary C. Bjorklund. Residual amplitude modulation in laser electro-optic phase modulation. *JOSA B*, 2(8):1320–1326, 1985. Available from: <http://www.opticsinfobase.org/abstract.cfm?uri=josab-2-8-1320>.
- [96] W. Winkler, K. Danzmann, A. Rdiger, and R. Schilling. Heating by optical absorption and the performance of interferometric gravitational-wave detectors. *Physical Review A*, 44(11):7022–7036, 1991. Available from: <http://docuser.v.ligo.caltech.edu/docs/public/P/P910012-00.pdf>.
- [97] H. Yamamoto. A modular fft-based interferometer simulation package. *LIGO DCC*, (G1400572), 2014. Available from: <https://dcc.ligo.org/LIGO-G1400572>.
- [98] A. Yariv. *Optical waves in crystals : propagation and control of laser radiation*. Wiley, New York, 1984.
- [99] C. Zhao, L. Ju, J. Degallaix, S. Gras, and D. Blair. Parametric instabilities and their control in advanced interferometer gravitational-wave detectors. *Physical Review Letters*, 94(12), April 2005. Available from: <http://link.aps.org/doi/10.1103/PhysRevLett.94.121102>.
- [100] D. S. Zheleznov, I. B. Mukhin, O. V. Palashov, E. A. Khazanov, and A. V. Voitovich. Faraday rotators with short magneto-optical elements for 50-kW laser power. *IEEE Journal of Quantum Electronics*, 43(6):451–457, June 2007. Available from: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4200825>.
- [101] D. Zwillinger. *CRC Standard Mathematical Tables and Formulae*. Chapman & Hall/CRC, Boca Raton, 2003.

BIOGRAPHICAL SKETCH

Chris Mueller graduated from Lakeside High School in Hot Springs, Arkansas in 2003, and attended the Georgia Institute of Technology with the anticipated major of Nuclear Engineering. After one year at Georgia Tech, Chris decided to take some time off during which he gained life experience living and working in Atlanta, Georgia. After moving to Orlando, Florida in 2005 Chris began working as a land surveyor while taking night classes at Valencia Community College (now Valencia College). It was during a physics class instructed by Dr. George Coyne at Valencia that Chris decided to get a four year degree in physics. In 2007 Chris transferred to the University of Florida in Gainesville, Florida from which he graduated in 2009 with a Bachelor of Science in physics and a Bachelor of Art in mathematics, both magna cum laude. Chris stayed at the University of Florida to complete his graduate studies in physics beginning in the fall of 2009. In the spring of 2012 Chris moved to Baton Rouge, Louisiana to take part in the commissioning of the newly built Advanced LIGO interferometer located at the LIGO Livingston Observatory. The time period spent at the observatory is when Chris produced most of the results found in this dissertation.