DESIGN AND PERFORMANCE OF HIGH LASER POWER INTERFEROMETERS FOR GRAVITATIONAL-WAVE DETECTION

By

KATHERINE LAIRD DOOLEY

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

2011

For my grandmother

ACKNOWLEDGMENTS

I was one of the few very lucky graduate students who had the privilege to spend my graduate school years at the LIGO Livingston Observatory. I got to learn how the interferometer works first-hand and I had the privilege of getting to know and work with so many of the world's top interferometer experts. Many thanks to my advisor, Dave Reitze, for supporting my desire to reside at the site and for spending a year there himself helping me get my thesis work underway and well-defined for the next several years. Thank you, too, for all of the phone calls over the years.

I wouldn't have made it all this way if it weren't for my teachers, friends, and family who have served as role models throughout my life, making me feel completely comfortable to pursue my interests, including those in math and science. Special thanks to my Uncle Tom, Jamie Lombardi, Ben Hummon, and Mrs. Matts. Most of all, I thank my parents, Janine and Alan, for their support and love all along the way.

My LIGO career would not be if it weren't for the serendipitous timing of chatting with my late neighbor, Noah Goldfish, who learned of my SURF work and continued interest in LIGO. Noah introduced me to Andri Gretarsson who in turn recommended me to the UF group. Thank you, Andri. I am thankful for Guido Mueller, Dave Reitze, and David Tanner's good judgement to recruit me only 2 months before the school year started and despite the fact I had already accepted grad school admission elsewhere. That was the best thing that could have happened for me.

Thank you to all of my fellow graduate students who helped figure things out with me and who made working at the sites all the more fun: Tobin Fricke, Nicolas Smith, Dan Hoak, Anamaria Effler, Ryan DeRosa, Jeff Kissel, and Rupal Amin. Tobin deserves extra recognition for all that he taught me and for his endless patience and encouragement.

Rana Adhikari always made sure I had one too many interesting projects to work on and would bug me on my progress. It was nice to be looked after. Thank you, Rana, and thanks for also teaching me to write good elog entries. Thanks to Sam Waldman, Lisa Barsotti, and Matt

Evans for demonstrating how to be good scientists and for all the time you spent at the sites. Thanks to Rai Weiss for his enthusiasm, for making me make that first outline of my thesis, and for always being an advocate for us students.

Thanks to the LLO staff and the UF LIGO group, especially Tom Evans for letting me borrow his juggling clubs and Guido Mueller for being so reasonable. And David Feldbaum says I must include a thank you for him because I made him search the elog for a spiricon BCS picture one too many times for me. Thank you, David!

Finally, this thesis would not be if it weren't for Valera Frolov, who I thank sincerely for welcoming me as a fellow commissioner and for being a role model on how to approach scientific work. He taught me to think critically and could always be count on whether it be day, night, or weekend, to offer advice and thought-provoking discussions.

This work was supported by the National Science Foundation through grants PHY-0855313 and PHY-0555453. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

A2L	Angle to length
ADC	Analog-to-digital converter
AS	Anti-symmetric
ASC	Angular Sensing and Control
BCS	Beam centering servo
BS	Beam splitter
CWP	Calcite wedge polarizer
DAC	Digital-to-analog converter
DARM	Differential arm
DC	Direct conversion; Direct current
DKDP	Deuterated potassium dihydrogen phosphate
DOF	Degree of freedom
EOM	Electro-optic modulator
ETM	End test mass
FI	Faraday isolator
FPM	Fabry-Perot Michelson
GW	Gravitational wave
HEPI	Hydraulic external pre-isolator
HWP	Half-wave plate
IO	Input Optics
ISC	Interferometer sensing and control
ITM	Input test mass
L2A	Length to angle
LIGO	Laser Interferometer Gravitational-wave Observatory
LHO	LIGO Hanford Observatory
LLO	LIGO Livingston Observatory

LSC	Length Sensing and Control
LVEA	Large vacuum equipment area
MC	Mode cleaner
MICH	Michelson
MMT	Mode-matching telescope
NSPOB	Normalized beam splitter sideband pick-off
OSEM	Optical sensor and electro-magnetic actuator
POB	Beam splitter pick-off
PSL	Pre-stabilized laser
PRC	Power recycling cavity
PRM	Power-recycled Michelson
QPD	Quadrant photodiode
QR	Quartz rotator
RMS	Root mean square
REFL	Reflected beam
RF	Radio frequency
RM	Recycling mirror
ROC	Radius of curvature
RP	Radiation pressure
SNR	Signal to noise ratio
SPOB	Beam splitter sideband pick-off
TFP	Thin film polarizer
TGG	Terbium gallium garnate
TM	Test mass
UGF	Unity gain frequency
VIRGO	Variability of Solar Irradiance and Gravity Oscillations
WFS	Wave-front sensor

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

DESIGN AND PERFORMANCE OF HIGH LASER POWER INTERFEROMETERS FOR GRAVITATIONAL-WAVE DETECTION

By

Katherine Laird Dooley

December 2011

Chair: David Reitze
Major: Physics

A prediction of Einstein's general theory of relativity, gravitational waves (GWs) are perturbations of the flat space-time Minkowski metric that travel at the speed of light. Indirectly measured by Hulse and Taylor in the 1970s through the energy they carried away from a binary pulsar system, gravitational waves have yet to be detected directly. The Laser Interferometer Gravitational-wave Observatory (LIGO) is part of a global network of gravitational-wave detectors that seeks to detect directly gravitational waves and to study their sources.

LIGO operates on the principle of measuring the gravitational wave's physical signature of a strain, or relative displacement of inertial masses. An extremely small effect whose biggest of expected transient signals on Earth is on the order of one part in 10^{23}, gravitational-wave strain can only be measured by detectors so sensitive to displacement as to encounter the effects of quantum physics. To improve their sensitivities and to demonstrate advanced technologies, the LIGO observatories in Hanford, WA and Livingston, LA underwent an upgrade between fall 2007 and summer 2009 called Enhanced LIGO. This study focuses on the experimental challenges of one of the goals of the upgrade: operating at an increased laser power.

I present the design and characterization of two of the interferometer subsystems that are critical for the path towards higher laser power: the Input Optics (IO) and the Angular Sensing and Control (ASC) subsystems. The IO required a new design so its optical components would not be susceptible to high power effects such as thermal lensing or thermal beam drift. The ASC required a new design in order to address static instabilities of the arm cavities caused

by increased radiation pressure. In all, I demonstrate the capability of an interferometric GW detector to operate at several times the highest of laser powers previously used.

CHAPTER 1
THE SEARCH FOR GRAVITATIONAL WAVES

Einstein's predictions of general relativity opened to the scientific community a whole new window of how to look at the universe. Just as scientists had been building detectors to observe directly optical and microwave radiation, they now had the theory in hand to think about building detectors for gravitational radiation. Gravitational waves (GW) are dynamic strains in space-time that travel at the speed of light and are generated by non-axisymmetric acceleration of mass.

Joseph Weber of the University of Maryland and John Wheeler of Princeton University introduced the field of gravitational wave astronomy in the 1960s. Weber built a resonant bar, the first instrument designed to directly observe gravitational waves [1]. Although his bar never made a positive detection, the interest in directly detecting gravitational waves persisted, and new and more sensitive detector designs were conceived.

The most promising of new detector designs for measuring a gravitational wave's distortion of space-time proved to be a laser interferometer. Robert Forward of Hughes Spacecraft built the first bench top prototype in the 1970s [2]. Rai Weiss of M.I.T. and Ron Drever of Caltech with the aid of others developed this concept into what is becoming a worldwide array of large scale interferometers.

The field of ground-based gravitational-wave physics is rapidly approaching a state with a high likelihood of detecting GWs for the first time. Such a detection will not only validate part of Einstein's general theory of relativity, but initiate an era of astrophysical observation of the universe through GWs. A first detection is expected to witness an event such as a binary black hole/neutron star merger. This chapter provides the theoretical framework of gravitational wave generation and presents various ways to detect GWs, including the current status of an effort to do so. I explain the purpose of this dissertation in the context of these current effort.

1.1 The Theory of Gravitational Radiation

Gravitational radiation is a perturbation $|h_{\mu\nu}| \ll 1$ to the flat space-time Minkowski metric $\eta_{\mu\nu} = \text{diag}(-1,1,1,1)$ [3]. The metric describing space-time in the presence of gravitational

radiation is therefore

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \tag{1–1}$$

As in electrodynamics where one has freedom in choosing the vector potential $\vec{A}$ for calculating the magnetic field $\vec{B} = \vec{\nabla} \times \vec{A}$, one also has freedom in general relativity in choosing the form of $h_{\mu\nu}$ for ease of calculation. A convenient and popular choice is called the transverse-traceless (TT) gauge in which

$$h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+(t) & h_\times(t) & 0 \\ 0 & h_\times(t) & -h_+(t) & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{1–2}$$

where the $+$ and $\times$ represent two linearly independent polarizations. Without loss of generality, we consider the h_+ polarization in the example that follows.

For a gravitational wave traveling along the z-axis, the space-time metric is given by:

$$ds^2 = -c^2dt^2 + [1 + h_+(t)]dx^2 + [1 - h_+(t)]dy^2. \tag{1–3}$$

This says the TT coordinate system is stretched along the x axis and and compressed along the y axis by a factor of

$$\sqrt{1 \pm h_+(t)} \approx 1 \pm \frac{1}{2}h_+(t). \tag{1–4}$$

Therefore, for two free masses located a proper distance L from one another along either the x-axis or the y-axis, their separation is magnified by the factor in Eq. 1–4 in the presence of a gravitational wave. Their coordinate separations, however, remain constant. The gravitational wave perturbation is a dimensionless strain:

$$h_+(t) = 2\frac{\Delta L(t)}{L}, \tag{1–5}$$

where $\Delta L(t)$ is the change in separation between the free masses, as illustrated in Fig. 1-1

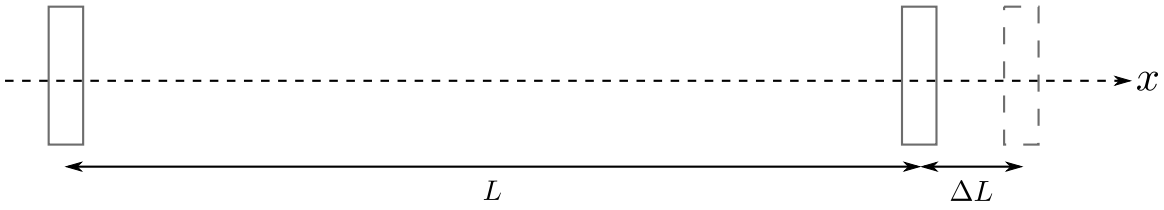

Figure 1-1. Depiction of strain.

1.2 Sources

Any object with an accelerating mass quadrupole moment generates gravitational waves. The typical strain amplitudes, however, are extremely tiny: a binary system of coalescing $1.4M_{\odot}$ neutron stars in the Virgo Cluster (a distance of 15 Mpc) would produce a maximum GW strain on Earth of only 10^{-21} at 800 Hz [4]. The strain is proportional to source mass, M, and velocity, v, and inversely proportional to the distance from the observer, R:

$$h \approx \frac{GMv^2}{Rc^4} \tag{1–6}$$

The quantity G/c^4 is what sets the scale for strain amplitudes because of how small it is: 8.26×10^{-45} $\mathrm{m^{-1}kg^{-1}s^2}$. Consequently, the most promising sources of detectable gravitational waves are nearby, fast-moving, massive astrophysical objects that include

- supernovae [5]
- binary stars (orbiting or coalescing) [6]
- spinning neutron stars [7]
- cosmological/astrophysical background [8]

and can be categorized as producing periodic, burst, or stochastic GWs.

Stably orbiting binary star systems comprised of black holes or neutron stars as well as rapidly spinning non-axisymmetric pulsars are considered periodic sources since they will produce GWs of relatively constant frequency. These reliable sources of GWs require a long integration time to pick out their signal above noise. Supernovae are burst sources since the gravitational collapse will produce a short-lived, unmodeled emission of GWs. Binaries in their final tens of milliseconds of inspiral also fall into this category. Finally, the anisotropies in the

inflation of the universe together with the hum of all distant astrophysical sources will create a stochastic background of radiation. Coherent cross-correlation between multiple detectors is necessary for measuring the constant amplitude, broad-spectrum GW background [9].

Directly detecting gravitational radiation from any such source will reveal information that electromagnetic radiation cannot convey. The frequency of the GW tells about the dynamical timescale of the source. Only through GW radiation, for example, can mass and spin properties of a black hole be revealed. A first detection is expected to witness an event such as a binary black hole/neutron star coalescence [10].

1.3 Methods of Detection

In order to detect directly a gravitational wave, the instrument must be sensitive to strain. Weber's bar and laser interferometers both accomplish this requirement. There is a third method of detection, however, that has already proved successful, although the detection is not direct. Hulse and Taylor observed the rate of change of the orbital period of a binary star system, demonstrating beautifully a precise agreement with the predictions of GR should the rate of change be due to gravitational radiation [11, 12]. Awarded the Nobel Prize for their work, Hulse and Taylor's indirect evidence of GWs has fueled the field to produce a direct detection. Newer methods under active research include pulsar timing [13] and B-mode measurements of the cosmic microwave background polarization. For an approachable overview of the history of the field, including detector design choices and estimated GW strain amplitudes of various sources, refer to Ref. [14].

1.4 State of Ground-based Interferometry

A network of first generation kilometer-scale laser interferometer gravitational-wave detectors completed its integrated 2-year data collection run in 2007, called S5. The instruments were: the American Laser Interferometer Gravitational-wave Observatories (LIGO) [15], one in Livingston, LA with 4 km long arms and two in Hanford, WA with 4 km and 2 km long

arms; the 3 km French-Italian detector VIRGO [16] in Cascina, Italy; and the 600 m German-British detector GEO [17] in Ruthe, Germany. Multiple separated detectors increase detection confidence through signal coincidence and improve source localization through triangulation.

The first generation of LIGO, known as Initial LIGO, achieved its design goal of sensitivity to GWs in the 40 Hz - 7000 Hz band which included a record strain sensitivity of $2 \times 10^{-23}/\sqrt{\text{Hz}}$ at 155 Hz. However, only the closest of sources produce enough GW strain to appear in LIGO's band, and no gravitational wave has yet been found in the S5 data. A second generation of LIGO detectors, Advanced LIGO, has been designed to be at least an order of magnitude more sensitive at several hundred Hz and above and include an impressive increase in bandwidth down to 10 Hz, dramatically increasing the chances of detection. The baseline Advanced LIGO design [18] improves upon Initial LIGO by featuring better seismic isolation, the addition of a signal recycling mirror at the output port, homodyne readout, and an increase in laser power from 10 W to 165 W.

To test some of Advanced LIGO's new technologies so unforeseen difficulties could be addressed and so that a more sensitive data taking run could take place, increasing the chances of detection, an incremental upgrade to the interferometers was carried out after S5 [19]. This project, Enhanced LIGO, culminated with the S6 science run from July 2009 to October 2010. An output mode cleaner was designed, built and installed, and DC readout of the GW signal was implemented [20]. An Advanced LIGO active seismic isolation table was also built, installed, and tested [21, Ch. 5]. In addition, the 10 W Initial LIGO laser was replaced with a 35 W laser [22]. Accompanying the increase in laser power, the test mass Thermal Compensation System [23], the Alignment Sensing and Control, and the Input Optics were modified.

As of the writing of this dissertation (September 2011), construction and installation of Advanced LIGO is underway. The vacuum systems are being retro-fitted to accompany the new layout, and at LLO the 165 W laser has been installed. At both sites, the new seismic isolation platforms and multi-level suspension cages are being mass-produced. By 2012, the first of the suspended mirrors will be installed and testing begun. Simultaneously, VIRGO and GEO

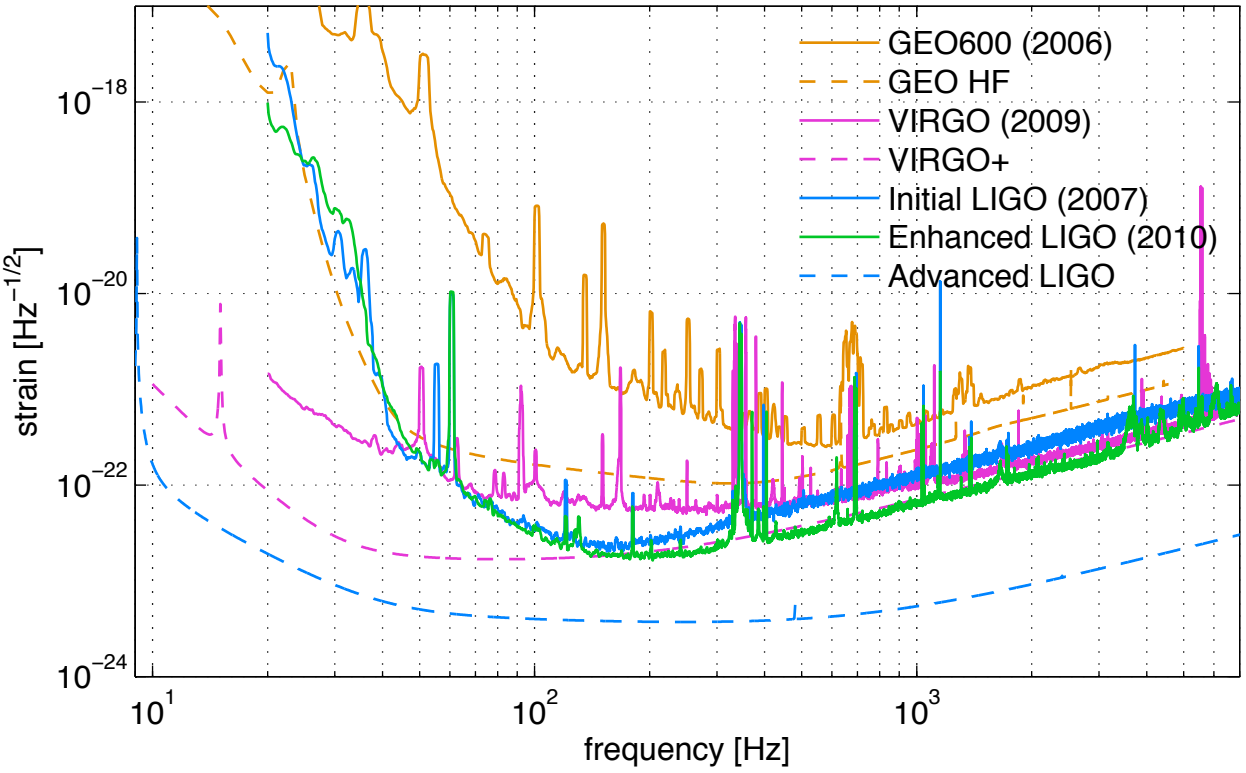

Figure 1-2. Strain sensitivities of LIGO-VIRGO collaboration interferometers. Solid lines show achieved detector noise floors and dashed lines show design noise floors for near-future generation interferometers.

are both undergoing their own upgrades as well [16, 24]. Figure 1-2 shows the achieved and theoretical future noise curves of this network of ground-based GW detectors.

1.5 Motivation for this Work

The purpose of this work is to demonstrate the capability of an interferometric gravitational wave detector to operate at several times the highest of laser powers previously used. From a naïve standpoint, more power is desirable since strain sensitivity improves by $\sqrt{P}$ in the high frequency (> 200 Hz) shot-noise-limited region. However, as detectors become more sensitive at low frequencies (< 40 Hz) in the years to come, radiation pressure noise will become the dominant noise source there, making high laser power operation a design trade-off. Currently, seismic noise limits low frequency sensitivity, so exploring the technical world of increasing the laser power is a fruitful adventure.

More power introduces radiation pressure and thermally induced side effects that must all be addressed for effective interferometer operation. Concerns about the practical difficulties of handling high power effects abounded during Initial LIGO when operating at the design power of 10 W proved more difficult and less straight-forward than expected. To achieve the Advanced LIGO design sensitivity, an ambitious 160 W of input power is needed. Without an understanding of the thermal and radiation pressure problems at 10 W, Advanced LIGO becomes a daunting goal.

The work presented in this dissertation was carried out during Enhanced LIGO to verify and investigate the predicted and unforeseen effects of as much as 25 W of laser power. It also served the purpose of enabling the operation of LIGO at higher powers and record strain sensitivities. I present the design and the measurements I made of the performance of two of the interferometer subsystems that are affected by an increase in laser power: the Input Optics and the Angular Sensing and Control. I show that the thermal and radiation pressure effects on these subsystems are well understood. This work on the Enhanced LIGO detectors informs design choices for Advanced LIGO.

1.5.1 The Input Optics High Power Story

The performance of the Initial LIGO Input Optics degraded with input power as the result of absorbing too much heat. Particular issues that needed to be addressed for any further increase in power included thermal steering of the beam rejected by the interferometer, a decrease in the optical isolation, and thermal lensing that drove the spatial mode of the beam directed at the interferometer away from optimal. We replaced two of the key Input Optics components and modified the others. I describe the design of the improved Input Optics for Enhanced LIGO which includes less absorptive optical components in order to conquer thermal issues at the source and changes to the design architecture that compensate for any residual effects. I also present the set of measurements I made to characterize the Input Optics performance with up to 30 W input power. I show that we can expect the design of the Enhanced LIGO Input Optics to also perform well for Advanced LIGO.

1.5.2 The Angular Sensing and Control High Power Story

Radiation pressure creates torques, a long-known concept, and the optical torque's ability to de-stabilize optical cavities was first recognized in 1991 by Solimeno et al. [25]. However, the theory of radiation pressure's effect on angular mechanical transfer functions was not fully appreciated until 2006, in the paper by Sidles and Sigg [26]. The concern arose that radiation pressure might be the factor limiting Initial LIGO's ability to increase the input power. Eiichi Hirose showed that the optical torque was present and measurable, but that it was *not* limiting at Initial LIGO's power [27].[1] The concern of the optical torque's role in cavity dynamics shifted to Enhanced and Advanced LIGO, which were designed to operate at four times and 20 times the laser power of Initial LIGO, respectively. Lisa Barsotti developed a numerical model of the angular sensing and controls for Enhanced LIGO, specifically including radiation pressure torque. She showed that, in principle, the radiation pressure torque can be controlled without detrimental consequences to the sensitivity of the detector [28]. We implemented Barsotti's theoretical control scheme and I measured its performance with up to 20 W of input power, demonstrating a thorough understanding of the principles at work and providing confidence in the ability to control radiation pressure torques in Advanced LIGO. I also improved upon Hirose's measurement of the optical angular (anti-)spring. In addition, through post-analysis of angular data, I demonstrate the potential of a technique that may be used in Advanced LIGO for reducing the angular control signals.

[1] In fact, after the Enhanced LIGO laser was installed, and before any changes were made to the ASC, we successfully operated the interferometer with 14 W input power.

CHAPTER 2
LASER INTERFEROMETERS FOR GRAVITATIONAL-WAVE DETECTION

We show in this chapter how a laser interferometer can detect gravitational wave strain, and we present the basic design principles of the LIGO detectors. We motivate the desire for higher laser power, and introduce some of the details of the interferometer that are relevant for later chapters. To reduce clutter, I do not specify the polarization of the gravitational wave strain and use simply h for its symbol.

2.1 Measuring Gravitational-wave Strain with Light

Considering a simple Michelson interferometer consisting of a laser, a beam splitter, and two end mirrors each a distance L from the beam splitter, one can understand intuitively why an interferometer can detect gravitational waves. If an appropriately polarized gravitational wave is present, it will stretch one arm and compress the other. For two wave packets leaving the beam splitter at the same time, each heading down a different arm, the roundtrip travel time for the light traveling down the stretched arm is longer than that for the light traveling down the compressed arm. For the stretched arm the roundtrip travel time is:

$$t_{\text{stretched}} = \frac{2L}{c}\left(1+\frac{h}{2}\right), \tag{2–1}$$

and for the compressed arm the roundtrip travel time is:

$$t_{\text{compressed}} = \frac{2L}{c}\left(1-\frac{h}{2}\right). \tag{2–2}$$

A stationary clock at the beam splitter could, in principle, measure the non-zero difference in arrival times, $\Delta t = 2Lh/c$, of the two different wave packets.[1]

[1] It should be noted that h is treated as a constant in Eqs. 2–1 and 2–2. We use the approximation that the gravitational wave wavelength λ_{gw} is much larger than the Michelson arm length L. This means that the temporal variation of $h(t)$ is negligible during the time it takes the photon to make its round trip.

In practice we send a continuous electromagnetic wave into the interferometer. The difference in travel times turns into a difference in phase of the beams returning to the beamsplitter:

$$\Delta\phi_{\text{BS\_MICH}} = \omega\Delta t = \frac{2L}{c}\omega h, \tag{2–3}$$

where ω is the angular frequency of the laser light. Ref. [29] provides a nice discussion about why GW detectors should work from the viewpoint of gauge invariance. We now introduce the modified Michelson interferometer used in LIGO, and in this context continue the discussion of strain measurement and sensitivity.

2.2 Power-recycled Fabry-Perot Michelson Interferometers

The LIGO detector configuration is a power-recycled Fabry-Perot Michelson (FPM) laser interferometer as depicted in Figure 2-1. A beam splitter (BS) directs 1064 nm light from a diode-pumped, power amplified, and intensity and frequency stabilized Nd:YAG laser to the Fabry-Perot arms, which are made of an input test mass mirror (ITM) and an end test mass mirror (ETM). Both arms are of length $L \approx 4$ km and are set to maintain nearly perfect destructive interference of the recombined light at the anti-symmetric (AS) port, where a photodetector is placed to measure any change in power. A power recycling mirror (RM) at the symmetric port directs the constructively-interfered light back into the interferometer.

The Fabry-Perot arms are a modification to the Michelson that increases the change in phase measured at the AS port compared to that for a simple Michelson. Rather than make a single round trip down each arm, the light is trapped by the Fabry-Perot cavity, experiencing many round trips before returning to the beam splitter and interfering with the light from the other arm. The effect is that Eq. 2–3 for the Fabry-Perot Michelson includes a frequency-dependent phase gain factor, $g_\phi(f)$:

$$\Delta\phi_{\text{BS\_FPM}} = \frac{2L}{c}\omega g_\phi(f) h(f). \tag{2–4}$$

For Enhanced LIGO, $g_\phi = 137$ at DC and falls off as $1/f$ after 85 Hz due to the storage time of the light in the arm cavities. The power recycling mirror is a modification to the Michelson

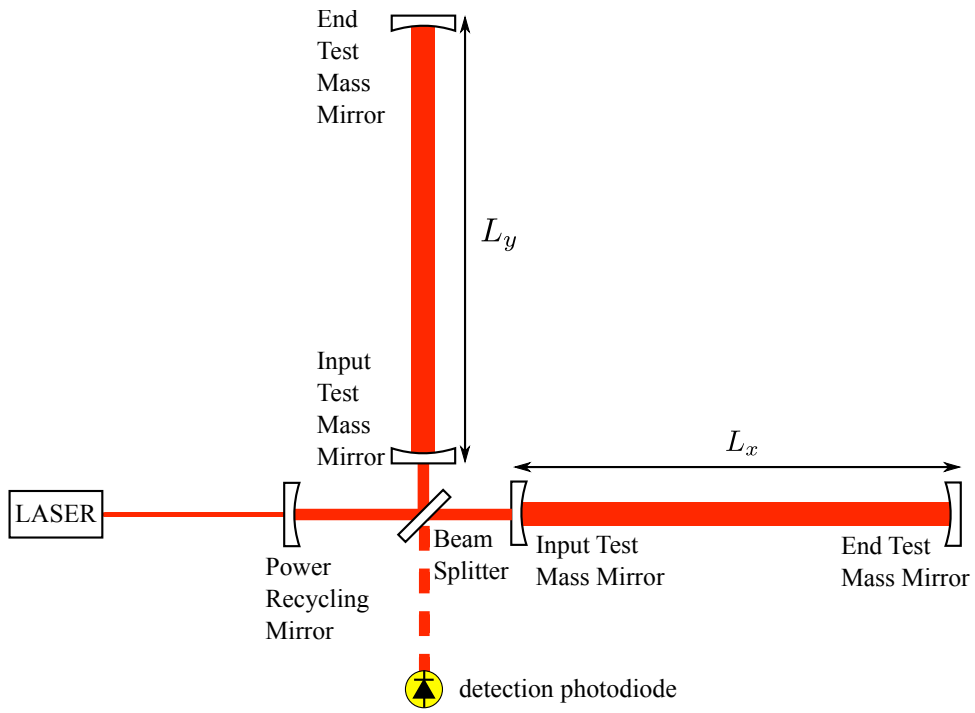

Figure 2-1. Power-recycled Fabry-Perot Michelson laser interferometer.

interferometer that increases the circulating power by a factor of $g_{cr}^2 \approx 40$. Details are found in Appendix A.6.2.

2.2.1 DC Readout

The gravitational wave readout in Enhanced LIGO was not operated precisely at the dark fringe at the AS port. Instead, it used a small offset from the quadratic minimum so that small changes in phase linearly produce power changes as is depicted in Figure 2-2. The offset used was $\phi_0 \approx 6 \times 10^{-5}$ rad. The technique is a form of homodyne detection called DC readout. For details of the implementation of DC readout in Enhanced LIGO, refer to Ref. [30].

With a DC offset, the electric field at the AS port is $E_{AS} = E_{BS} \sin(\phi_0 + \Delta\phi)$. Squaring the electric field and expanding about ϕ_0, we determine the power incident on the photodetector, P_{AS}:

$$P_{AS} = P_{BS} \sin^2(\phi_0 + \Delta\phi) \tag{2–5}$$

$$\approx P_{BS} \sin^2(\phi_0) + 2P_{BS}\phi_0\Delta\phi \tag{2–6}$$

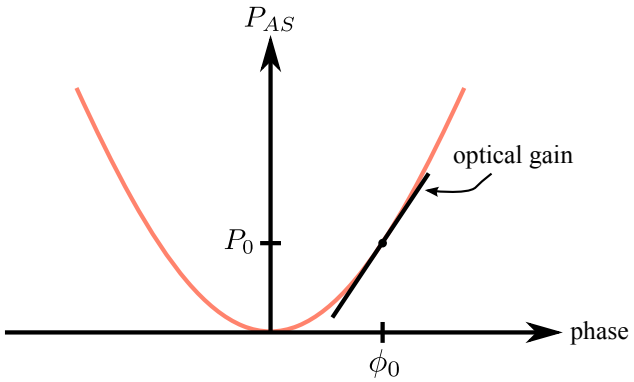

Figure 2-2. The DC readout dark fringe. The AS port is not kept at the dark fringe, but is slightly offset by ϕ_0. Changes in phase at the beam splitter are a linear function of power.

The first term on the right hand side of the expanded P_{AS} is the DC power due to the static offset from the fringe. The second term on the right hand side describes how a change in phase at the beam splitter is converted to a change in power:

$$\frac{dP_{AS}}{d\phi_{BS}} = 2P_{BS}\phi_0 \tag{2–7}$$

This relationship is linear and proportional to the power at the beam splitter. Throughout this dissertation, when we refer to signals falling in a *linear regime*, we mean that they are small enough to be well modeled by a tangent to the actual response, just as for the case described here regarding small phase signals.

2.2.2 DARM

The differential arm length (commonly known as DARM) is of central interest. This is the length degree of freedom affected by gravitational waves. It is defined as

$$\text{DARM} := L_- := L_x - L_y \tag{2–8}$$

where L_x and L_y are the lengths of the x-arm and y-arm, respectively. When there is no gravitational wave, $L_- = 0$, but in the presence of a gravitational wave, the DARM signal is:

$$L_- = Lh \tag{2–9}$$

We see that L is the conversion factor between GW strain and DARM.

2.2.3 DARM and Strain Optical Gain

The DARM optical gain tells how differential displacement is converted to power at the AS port and has units of Watts per meter. Combining Eqs. 2–4, 2–7, and 2–9, the DARM optical gain at DC of the LIGO interferometer with DC readout is:

$$\frac{dP_{AS}}{dL_{-}} = \frac{4}{c} P_{BS} \phi_0 \omega g_{\phi}. \tag{2–10}$$

Likewise, the strain optical gain is:

$$\frac{dP_{AS}}{dh} = \frac{4}{c} P_{BS} \phi_0 L \omega g_{\phi}. \tag{2–11}$$

2.3 Signal Versus Noise

From Eq. 2–11 we note four fundamental ways to increase the strain optical gain for a constant DARM offset and therefore produce more power at the AS port for a given GW strain:

- Increase the arm length.
- Increase the power at the beamsplitter.
- Increase the phase gain of the Fabry-Perot arms.
- Increase the laser frequency.

Although addressing each of these points will improve the strain signal, our ability to detect gravitational wave strain is also dependent on the detector noises which will mask a weak GW signal. No matter how large a signal one might have, it will not be found confidently, or at all, if there is too much noise. For a thorough overview of the interferometer noise analysis and examples of ways to improve strain sensitivity, refer to Ref. [31].

2.3.1 Noise

The sources of noise which contaminate the detector's output can be grouped into two categories:

- displacement noise
- sensing noise

Displacement noises are those that create real motion of the mirrors, while sensing noises are those that arise in the process of measuring the electric field at the detector's output.

The primary displacement noise that plagues terrestrial laser interferometers is motion of the ground, i.e. seismic noise. Thermal motion of the mirrors and their suspensions are another source of displacement noise [32]. The primary sensing noises are electronics ('dark') noise due to thermal noise in resistors and electronic amplifiers, and shot noise which arises from the Poisson statistics of photon arrival at the photodetector. Shot noise appears as a fluctuating power with amplitude spectral density

$$P_{SN} = \sqrt{2Ph_p\nu}, \tag{2–12}$$

where P is the mean power on the photodiode, h_p is Planck's constant, and ν is the frequency of the incident light. Shot noise is spectrally white. The detector electronics are typically designed so that electronics noise is never limiting.

2.3.2 Noise Floor

The detector's noise floor is limited by seismic noise below 40 Hz and by shot noise above 200 Hz. In general we endeavor to push the noise floor down as far as possible so that any underlying GW signals will be revealed. Whether limited by displacement noise or by sensing noise, the noise floor can be lowered by increasing the length of the arms, which acts as the conversion from strain to differential displacement (Eq. 2–9) [2] . Further improvements require considering the noise sources individually.

2.3.2.1 Displacement noise floor

At frequencies where the noise floor is limited by displacement noise, simply increasing the DARM optical gain (Eq. 2–10) will not help. The mirror differential displacements, whether due to gravitational waves or due to ground motion, are converted into power at the AS port in the exact same way. Reduction of displacement noises relies primarily on the development of more sophisticated seismic isolation systems and mirror suspension arrangements. Note that

[2] In the context of LIGO, increasing the arm lengths is not actually an option.

improvements in the displacement noise floor are only linear with respect to arm length, whereas each additional stage of seismic isolation can potentially reduce the displacement noise by $1/f^2$.

2.3.2.2 Sensing noise floor

The noise floor due to sensing noise is improved by increasing the DARM optical gain. In particular, the contribution due to shot noise may be found by dividing the shot noise amplitude spectral density by the strain optical gain:

$$h_{shot} = \sqrt{\frac{h_p}{2P_{BS}\nu}} \frac{c}{4\pi\phi_0 L g_\phi} \tag{2–13}$$

Here we see that the shot noise limit (calibrated in effective strain) drops with increases in the optical gain. Increasing the power in the interferometer improves the shot noise limit because the optical gain increases more quickly ($\propto P_{BS}$, see Eq. 2–10) than the shot noise amplitude (which goes like $\sqrt{P_{BS}}$, assuming the DARM offset is held constant).

2.4 Controlling the Interferometer

The ability of the interferometer to operate as described above requires that the many interferometer cavities be held on resonance. The motion of the mirrors in the absence of control is much too large–on the order of 1 μm, a full wavelength!–to maintain resonance. The motion of the interferometer mirrors must therefore be controlled. A feedback control system is implemented to hold the system sufficiently near the intended operating point (for DARM, within $\sim 10^{-13}$ m) so that the response to residual deviations remains linear. Calibration of the detector must take into account the action of the control system [33].

2.4.1 RF Sidebands

The various length and angular degrees of freedom are sensed through the use of radio-frequency (RF) sidebands on the carrier light. The standard technique for locking optical cavities is the Pound–Drever–Hall method of laser frequency stabilization as explained in Ref. [34]. LIGO uses three sets of RF sidebands at 24.4 MHz, 33.3 MHz, and 61.1 MHz for locking the Fabry-Perot arm cavities, input mode cleaner, and power recycling cavity, respectively. An electro-optic modulater (EOM) phase modulates the 1064 nm carrier light to produces these

reference fields that are critical for interferometer control. Phase modulation is explained in Appendix A.1.

2.4.2 Digital Control in LIGO

Although the interferometer is an analog instrument, it is interfaced through a digital control system.[3] The analog sensor signals are sent through analog-to-digital converters (ADCs), are digitally filtered, and then sent through digital-to-analog converters (DACs) before returning to the interferometer's actuators as control signals. The digital control system allows complex filters to be implemented and tuned from a comfortable control-room environment.

The various LIGO subsystems operate at different sample rates. The length sensing and control (LSC) subsystem, which measures and controls DARM in addition to other length degrees of freedom, operates at 16384 samples/second, while the angular sensing and control (ASC) system, which maintains mirror alignment, operates at 2048 samples/second. In addition to the all-important DARM channel, many other auxiliary data streams are permanently recorded.

2.4.3 Mirror Suspension and Actuation

The primary interferometer optics are suspended in vacuum so that they act like free masses in the horizontal plane at the frequencies in the GW detection band, and so that they are isolated from ground motion. Each mirror is hung from a single wire that loops around the bottom of the barrel of the mirror as shown in Figure 2-3. Stand-offs glued just above the mirror's center of mass on both sides of the barrel mark the final point of contact of the wire with the mirror, and both ends of the wire are clamped to the top of a suspension cage.

Each mirror is equipped with four optical sensor and electro-magnetic (OSEM) actuators for rough sensing and fine control of the mirror position and orientation. Magnets arranged to form the four corners of a square are glued on the mirror's back surface which are enveloped by

[3] There are a select few control systems that remain completely analog, like the laser intensity stabilization servo (ISS). When the frequencies of interest extend beyond several tens of thousands of Hz, the use of computers becomes impractical.

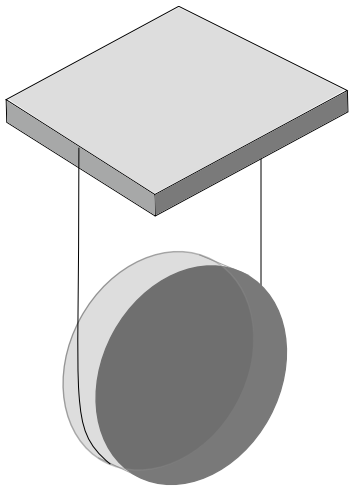

Figure 2-3. Simple sketch of a LIGO suspension.

the OSEM solenoid coil. The currents through each coil may be driven independently. Length control of the cavities, for instance, sends current of the same magnitude through each coil on a given mirror to provide a piston force for changing the mirror's position. OSEM sensing is accomplished through simple shadow sensors.

To avoid thermal noise, the mirror suspensions are designed to minimize dissipation. Damping for the large optics is therefore achieved through electronic servos. Motion of the optics corresponding to a change in cavity length is damped at all times by simple velocity damping servos that use the OSEM sensors and actuators. Angular motion is sensed and controlled at all times via optical levers which provide velocity damping between 0.2 Hz and 2 Hz.[4] The OSEM signals also provide angular feedback, but only when the interferometer is not locked.

2.5 Summary

The modified Michelson interferometer provides a robust foundation on which to build a gravitational wave detector in which fluctuating gravitational wave strains are transduced into measurable optical power fluctuations. For the interferometer to operate properly, the mirrors positions and orientations must be controlled. The noise floor of the interferometer may be

[4] The open loop transfer function of the optical lever servo is described in Appendix C.1.

improved in the shot noise dominated regime by increasing the laser power circulating in the interferometer. Increases in laser power, however, create several challenges. Two such challenges are addressed in this thesis: coping with the higher power in the interferometer's Input Optics (Chapter 3) and dealing with radiation pressure induced angular instabilities caused by high power in the interferometer's arm cavities (Chapters 4, 5, and 6).

CHAPTER 3
INPUT OPTICS DESIGN AND CHARACTERIZATION

3.1 Function of the Input Optics

The Input Optics (IO)[1] is one of the primary subsystems of the Laser Interferometer Gravitational-wave Observatory (LIGO) interferometers. Its purpose is to deliver an aligned, spatially pure, mode-matched beam with phase-modulation sidebands to the power-recycled Fabry-Perot Michelson interferometer. The IO also prevents reflected or backscattered light from reaching the laser and distributes the control sidebands reflected from the interferometer (designated the *reflected port*) to photodiodes for sensing and controlling the length and alignment of the interferometer. In addition, the IO provides an intermediate level of frequency stabilization and must have high overall optical efficiency. It must perform these functions without limiting the strain sensitivity of the LIGO interferometer. Finally, it must operate robustly and continuously over years of operation. The conceptual design is found in Ref. [35].

As shown in Figure 3-1, the Input Optics subsystem consists of four principle components located between the pre-stabilized laser and the power recycling mirror:

- electro-optic modulator (EOM)
- mode cleaner cavity (MC)
- Faraday isolator (FI)
- mode-matching telescope (MMT)

Each element is a common building block of many optical experiments and not unique to LIGO. However, their roles specific to the successful operation of interferometry for gravitational-wave detection are of interest and demand further attention. Here, we briefly review the purpose of each of the IO components; further details about the design requirements are in Ref. [36].

3.1.1 Electro-optic Modulator

The Length Sensing and Control (LSC) and Angular Sensing and Control (ASC) subsystems require phase modulation of the laser light at RF frequencies. This modulation is produced

[1] The Input Optics was originally called the Input-Output Optics (IOO).

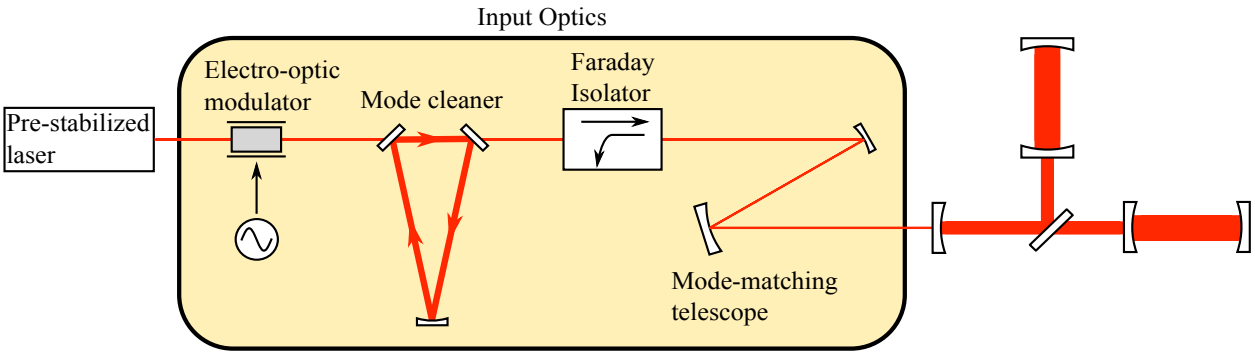

Figure 3-1. Block diagram of the Input Optics subsystem. The IO is located between the pre-stabilized laser and the recycling mirror and consists of four components: electro-optic modulator, mode cleaner, Faraday isolator, and mode-matching telescope. The electro-optic modulator is the only IO component outside of the vacuum system. Diagram is not to scale.

by an EOM, generating sidebands of the laser light which act as references against which interferometer length and angle changes are measured [37]. The sideband light must be either resonant only in the recycling cavity or not resonant in the interferometer at all. The sidebands must be offset from the carrier by integer multiples of the MC free spectral range so that neither MC length fluctuations nor phase modulation of the sidebands (due to phase noise of the RF oscillator) are converted to amplitude modulation.

3.1.2 Mode Cleaner

Stably aligned cavities, limited non-mode-matched (junk) light, and a frequency and amplitude stabilized laser are key features of any ultra sensitive laser interferometer. The mode cleaner, at the heart of the IO, plays a major role to this effect.

A three-mirror triangular ring cavity, the mode cleaner suppresses laser output not in the fundamental TEM_{00} mode, serving two major purposes. It enables the robustness of the ASC since higher order modes would otherwise contaminate the angular sensing signals of the interferometer. Also, all non-TEM_{00} light on the length sensing photodiodes, including those used for the GW readout, contributes shot noise but not signal and therefore diminishes the signal to noise ratio. The mode cleaner is thus largely responsible for achieving an aligned, minimally shot-noise-limited interferometer.

The MC also plays an active role in laser frequency stabilization [37], which is necessary for ensuring that the signal at the anti-symmetric port is due to arm length fluctuations rather than laser frequency fluctuations. In addition, the MC passively suppresses beam jitter at frequencies above 10 Hz.

3.1.3 Faraday Isolator

Faraday isolators are four-port optical devices which utilize the Faraday effect to allow for non-reciprocal polarization switching of laser beams. Any backscatter or reflected light from the interferometer (due to impedance mismatch, mode mismatch, non-resonant sidebands, or signal) needs to be diverted to protect the laser from back propagating light, which can introduce amplitude and phase noise. This diversion of the reflected light is also necessary for extracting length and angular information about the interferometer's cavities. The Faraday isolator fulfils both needs.

3.1.4 Mode-matching Telescope

The lowest-order mode cleaner and arm cavity spatial eigenmodes need to be matched for maximal power buildup in the interferometer. The mode-matching telescope is a set of three suspended concave mirrors between the mode cleaner and interferometer that expand the beam from a radius of 1.6 mm at the mode cleaner waist to a radius of 37 mm at the recycling mirror as shown in Figure 3-2. The MMT should play a passive role by delivering properly shaped light to the interferometer without introducing beam jitter or any significant aberration that can reduce mode coupling.

3.2 Thermal Problems in Initial LIGO

The Initial LIGO interferometers were equipped with a 10 W laser, yet operated with only 7 W input power due to power-related problems with other subsystems. The EOM was located in the 10 W beam and the other components experienced anywhere up to 7 W power. The 7 W operational limit was not due to the failure of the Input Optics; however, many aspects of the IO performance did degrade with power.

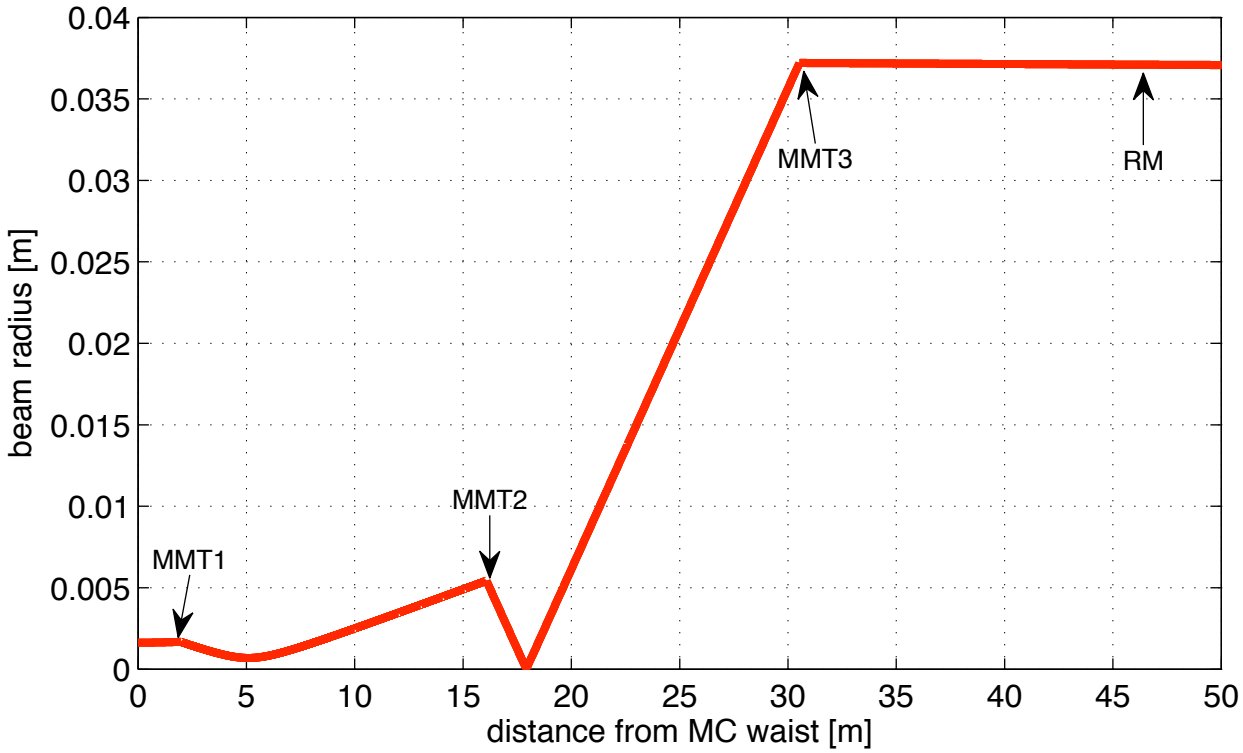

Figure 3-2. Beam profile through the Input Optics. The starting point is the mode cleaner waist and the changes in trajectory are due to the mode-matching telescope mirrors.

One of the primary problems of the Initial LIGO Input Optics [38] was thermal deflection of the back propagating beam due to thermally-induced refractive index gradients in the Faraday isolator. A significant beam drift between the interferometer's locked and unlocked states led to clipping of the reflected beam on the photodiodes used for length and alignment control. Our measurements determined a deflection of approximately 100 μrad/W in the FI. This was mitigated at the time by the design and implementation of an active beam steering servo on the beam coming from the isolator.

There were also known limits to the power the IO could sustain. Thermal lensing in the Faraday isolator optics began to alter significantly the beam mode at powers greater than 10 W, leading to a several percent reduction in mode matching to the interferometer [39]. Additionally, the absorptive FI elements would create thermal birefringence, degrading the optical efficiency and isolation ratio with power [40]. The Initial LIGO New Focus electro-optic modulators had an operational power limit of around 10 W. There was a high risk of damage to the crystals under

the stress of the 0.4 mm radius beam. Also, anisotropic thermal lensing with focal lengths as severe as 3.3 m at 10 W made the EOMs unsuitable for much higher power. Finally, the mode cleaner mirrors exhibited high absorption (as much as 24 ppm per mirror), enough that thermal lensing of the MC optics at Enhanced LIGO powers would induce higher order modal frequency degeneracy and result in a power-dependent mode mismatch into the interferometer [41, 42]. In fact, as input power increased from 1 W to 7 W the mode matching decreased from 90% to 83%.

In addition to the thermal limitations of the Initial LIGO IO, optical efficiency in delivering light from the laser into the interferometer was not optimal. Of the light entering the Input Optics chain, only 60% remained by the time it reached the power recycling mirror. Moreover, because only 90% at best of the light at the recycling mirror was coupled into the arm cavity mode, room was left for improvement in the implementation of the MMT.

3.3 Enhanced LIGO Input Optics Design

The Enhanced LIGO IO design addressed the thermal effects that compromised the performance of the Initial LIGO IO, and accommodated up to four times the power of Initial LIGO. Also, the design was a prototype for handling the 165 W laser planned for Advanced LIGO. Because the adverse thermal properties of the Initial LIGO IO (beam drift, birefringence, and lensing) are all attributable primarily to absorption of laser light by the optical elements, the primary design consideration was finding optics with lower absorption [39]. Both the EOM and the FI were replaced for Enhanced LIGO. Only minor changes were made to the MC and MMT. A detailed layout of the Enhanced LIGO IO is shown in Figure 3-3 and photographs are in Figure 3-4.

3.3.1 Electro-optic Modulator Design

We replaced the commercially-made New Focus 4003 resonant phase modulator of Initial LIGO with an in-house EOM design and construction. Both a new crystal choice and architectural design change allow for superior performance.

The Enhanced LIGO EOM design uses a crystal of rubidium titanyl phosphate (RTP), which has at most 1/10 the absorption coefficient at 1064 nm of the lithium niobate ($LiNbO_3$) crystal

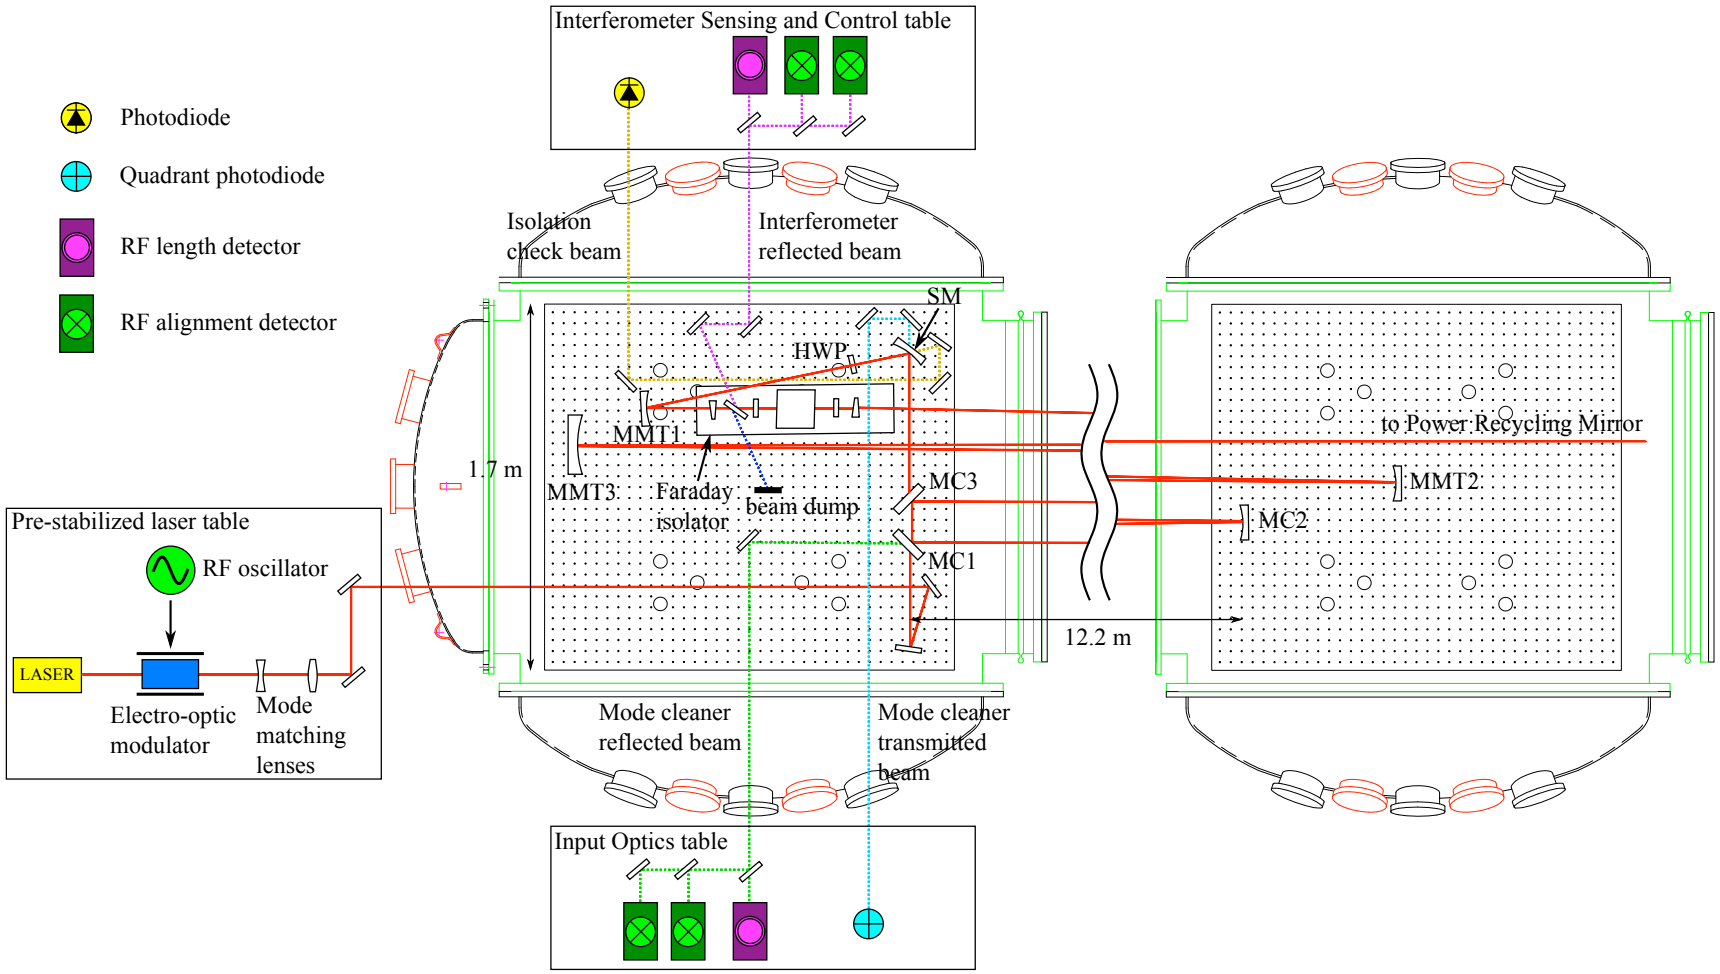

Figure 3-3. Enhanced LIGO Input Optics optical and sensing configuration. The HAM1 (horizontal access module) vacuum chamber is featured in the center, with locations of all major optics superimposed. HAM2 is shown on the right, with its components. These tables are separated by 12 m. The primary beam path, beginning at the pre-stabilized laser and going to the power recycling mirror, is shown in red as a solid line, and auxiliary beams are different colors and dotted. The MMTs, MCs, and steering mirror (SM) are suspended; all other optics are fixed to the seismically isolated table. The laser and sensing and diagnostic photodiodes are on in-air tables.

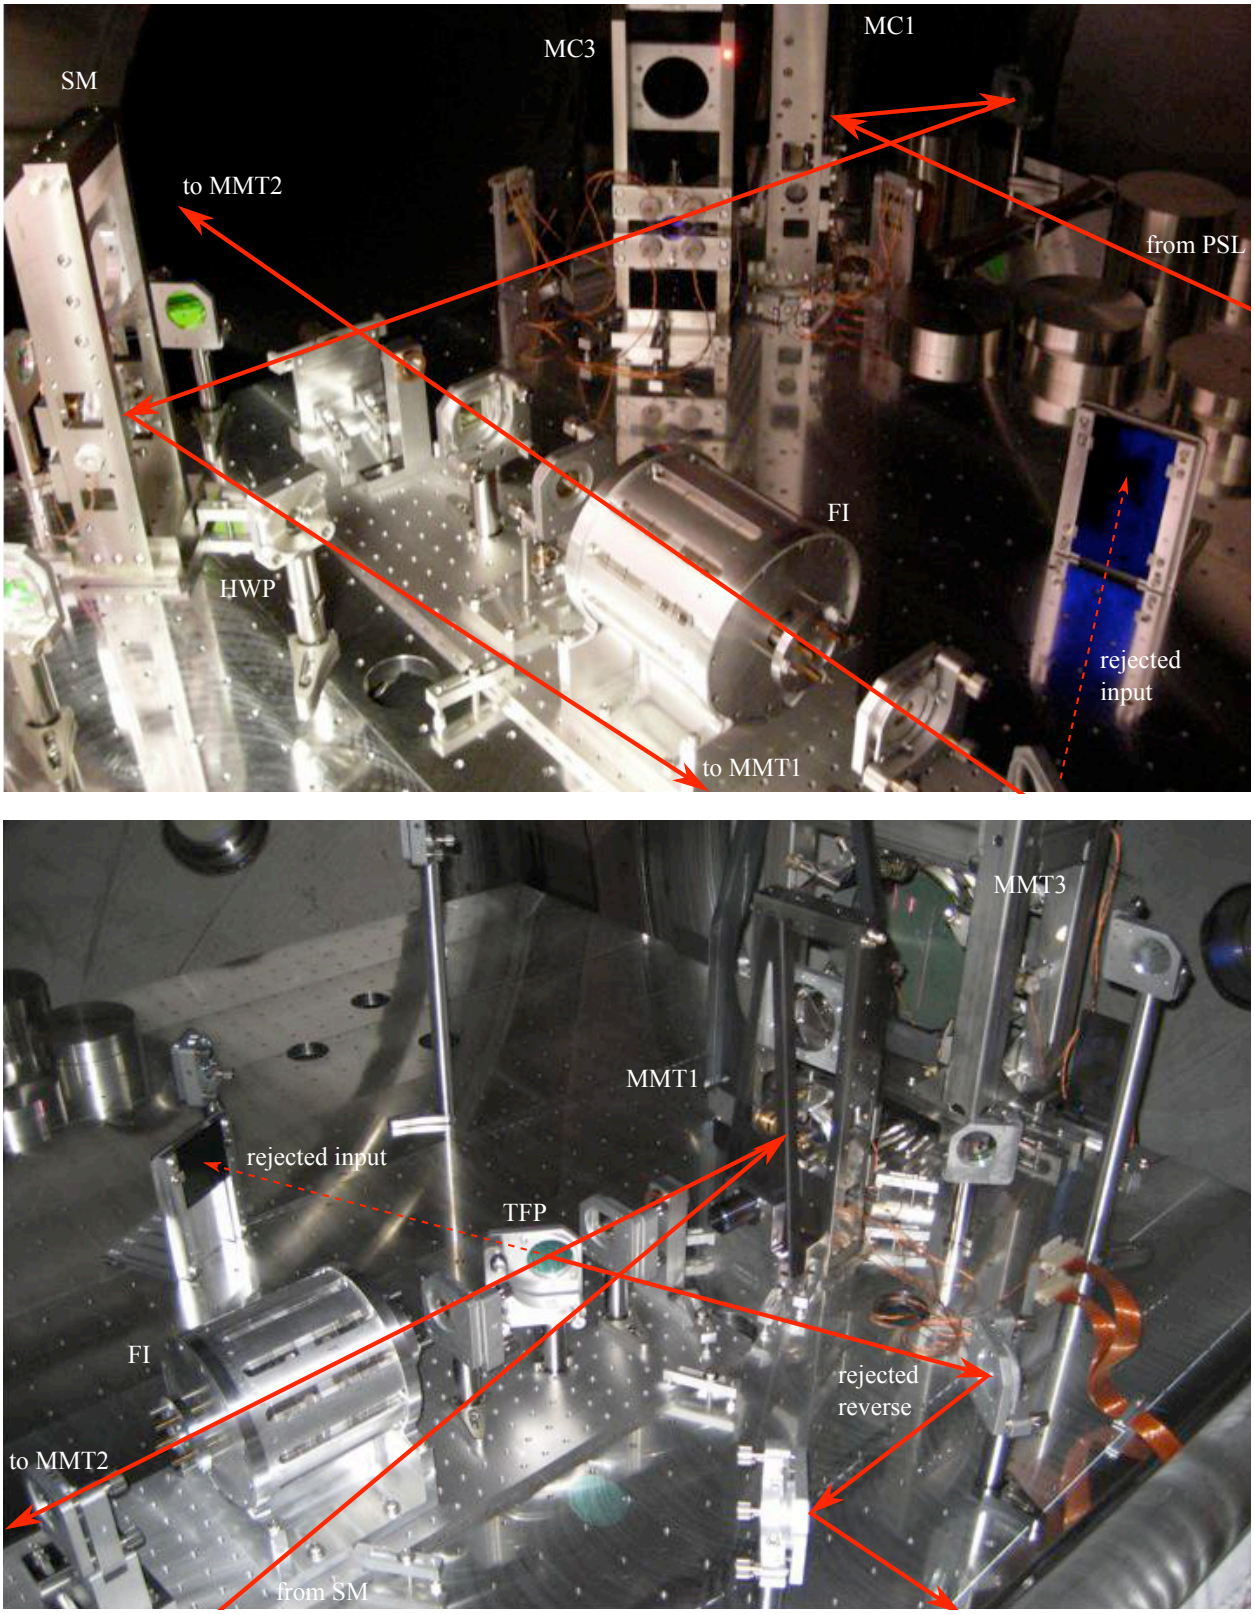

Figure 3-4. Photographs of the Enhanced LIGO HAM1 Input Optics *in situ* with a drawing of the beam path superimposed. Photographs courtesy of Katherine Dooley.

Table 3-1. Comparison of selected properties of the Initial and Enhanced LIGO EOM crystals, $LiNbO_3$ and RTP, respectively. RTP was preferred for Enhanced LIGO because of its lower absorption, superior thermal properties, and similar electro-optic properties [39].

	units	$LiNbO_3$	RTP
damage threshold	MW/cm^2	280	>600
absorption coeff. at 1064 nm	ppm/cm	<5000	<500
electro-optic coeff. ($n_z^3 r_{33}$)	pm/V	306	239
dn_y/dT	10^{-6}/K	5.4	2.79
dn_z/dT	10^{-6}/K	37.9	9.24

from Initial LIGO. At 200 W the RTP should produce a thermal lens of 200 m and higher order mode content of less than 1%, compared to the 3.3 m lens the $LiNbO_3$ produces at 10 W. The RTP has a minimal risk of damage, since it has both twice the damage threshold of $LiNbO_3$ and is subjected to a beam twice the size of that in Initial LIGO. RTP and $LiNbO_3$ have similar electro-optic coefficients. Also, RTP's dn/dT anisotropy is 50% smaller. Table 3-1 compares the properties of most interest of the two crystals.

We procured the RTP crystals from Raicol and packaged them into specially-designed, custom-built modulators. The crystal dimensions are $4 \times 4 \times 40$ mm and their faces are wedged by 2.85° and anti-reflection (AR) coated. The wedge serves to separate the polarizations and prevents an etalon effect, resulting in a suppression of amplitude modulation. Only one crystal is used in the EOM in order to reduce the number of surface reflections. Three separate pairs of electrodes, each with its own resonant LC circuit, are placed across the crystal in series, producing the three required sets of RF sidebands: 24.5 MHz, 33.3 MHZ and 61.2 MHz. A diagram is shown in Figure 3-5. Reference [43] contains further details about the modulator architecture.

3.3.2 Mode Cleaner Design

The mode cleaner is a suspended 12.2 m long triangular ring cavity with finesse $\mathscr{F}$=1280 (refer to Appendix A.2 for a measurement of the finesse) and free spectral range of 12.243 MHz. The three mirror architecture was selected over the standard two mirror linear filter cavity because it acts as a polarization filter and because it eliminates direct path back propagation to the laser [44]. A pick-off of the reflected beam is naturally facilitated for use in generating control

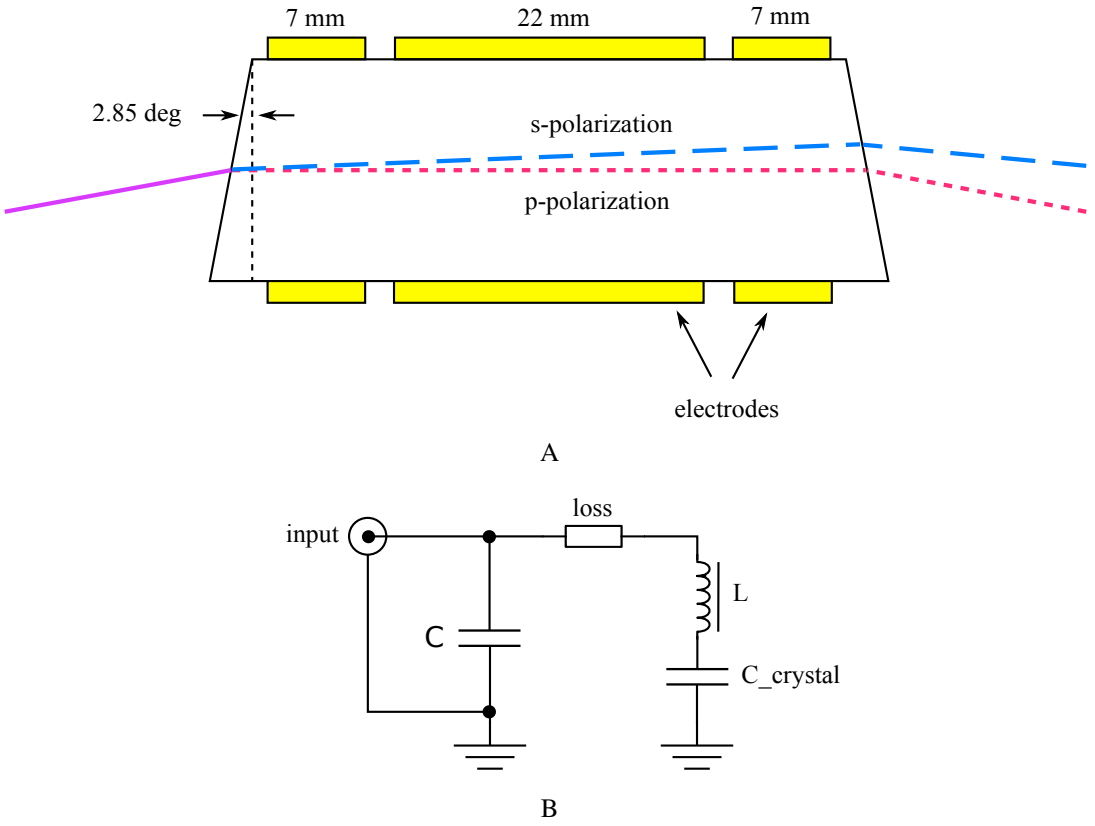

Figure 3-5. Electro-optic modulator design. A) The single RTP crystal is sandwiched between three sets of electrodes that apply three different modulation frequencies. The wedged ends of the crystal separate the polarizations of the light. The p-polarized light is used in the interferometer. B) A schematic for each of the three impedance matching circuits of the EOM. For the three sets of electrodes, each of which creates its own $C_{crystal}$, a capacitor is placed parallel to the LC circuit formed by the crystal and a hand-wound inductor. The circuits provide 50 Ω input impedance on resonance and are housed in a separate box from the crystal.

signals. A potential downside to the three mirror design is the introduction of astigmatism, but this effect is negligible due to the small opening angle of the MC.

The MC has a round-trip length of 24.5 m. The beam waist has a radius of 1.63 mm and is located between the two 45° flat mirrors, MC1 and MC3. See Figure 3-3). A concave third mirror, MC2, 18.15 m in radius of curvature, forms the far point of the mode cleaner's isosceles triangle shape. The power stored in the MC is 408 times the amount coupled in, equivalent to about 2.7 kW in Initial LIGO and at most 11 kW for Enhanced LIGO. The peak irradiances are 32 kW/cm^2 and 132 kW/cm^2 for Initial LIGO and Enhanced LIGO, respectively.

The mode cleaner mirrors are 75 mm in diameter and 25 mm thick. The substrate material is fused silica and the mirror coating is made of alternating layers of silica and tantala. In order to reduce the absorption of heat in these materials and therefore improve the transmission and modal quality of the beam in the mode cleaner, we removed particulate by drag wiping the surface of the MC mirrors with methanol and optical tissues. The mode cleaner was otherwise identical to that in Initial LIGO.

3.3.3 Faraday Isolator Design

The Enhanced LIGO Faraday isolator design required not only the use of low absorption optics, but additional design choices to mitigate any residual thermal lensing and birefringence. In addition, trade-offs between optical efficiency in the forward direction, optical isolation in the backwards direction, and feasibility of physical access of the return beam for signal use were considered. The result is that the Enhanced LIGO Faraday isolator needed a completely new architecture and new optics compared to both the Initial LIGO FI and commercially available isolators.

Figure 3-6 shows a schematic of the Enhanced LIGO Faraday Isolator. It begins and ends with low absorption calcite wedge polarizers (CWP). Between the CWPs is a thin film polarizer (TFP), a deuterated potassium dihydrogen phosphate (DKDP) element, a half-wave plate (HWP), and a Faraday rotator. The rotator is made of two low absorption terbium gallium garnet (TGG) crystals sandwiching a quartz rotator (QR) inside a 7-disc magnet with a maximum field strength of 1.16 T. The forward propagating beam upon passing through the TGG, QR, TGG, and HWP elements is rotated by $+22.5^\circ - 67.5^\circ + 22.5^\circ + 22.5^\circ = 0^\circ$. In the reverse direction, the rotation through HWP, TGG, QR, TGG is $-22.5^\circ + 22.5^\circ + 67.5^\circ + 22.5^\circ = 90^\circ$. The TGG crystals are non-reciprocal devices while the QR and HWP are reciprocal.

3.3.3.1 Thermal birefringence

Thermal birefringence is addressed in the Faraday rotator by the use of the two TGG crystals and one quartz rotator rather than the typical single TGG [45]. In this configuration, any thermal polarization distortions that the beam experiences while passing through the first TGG rotator

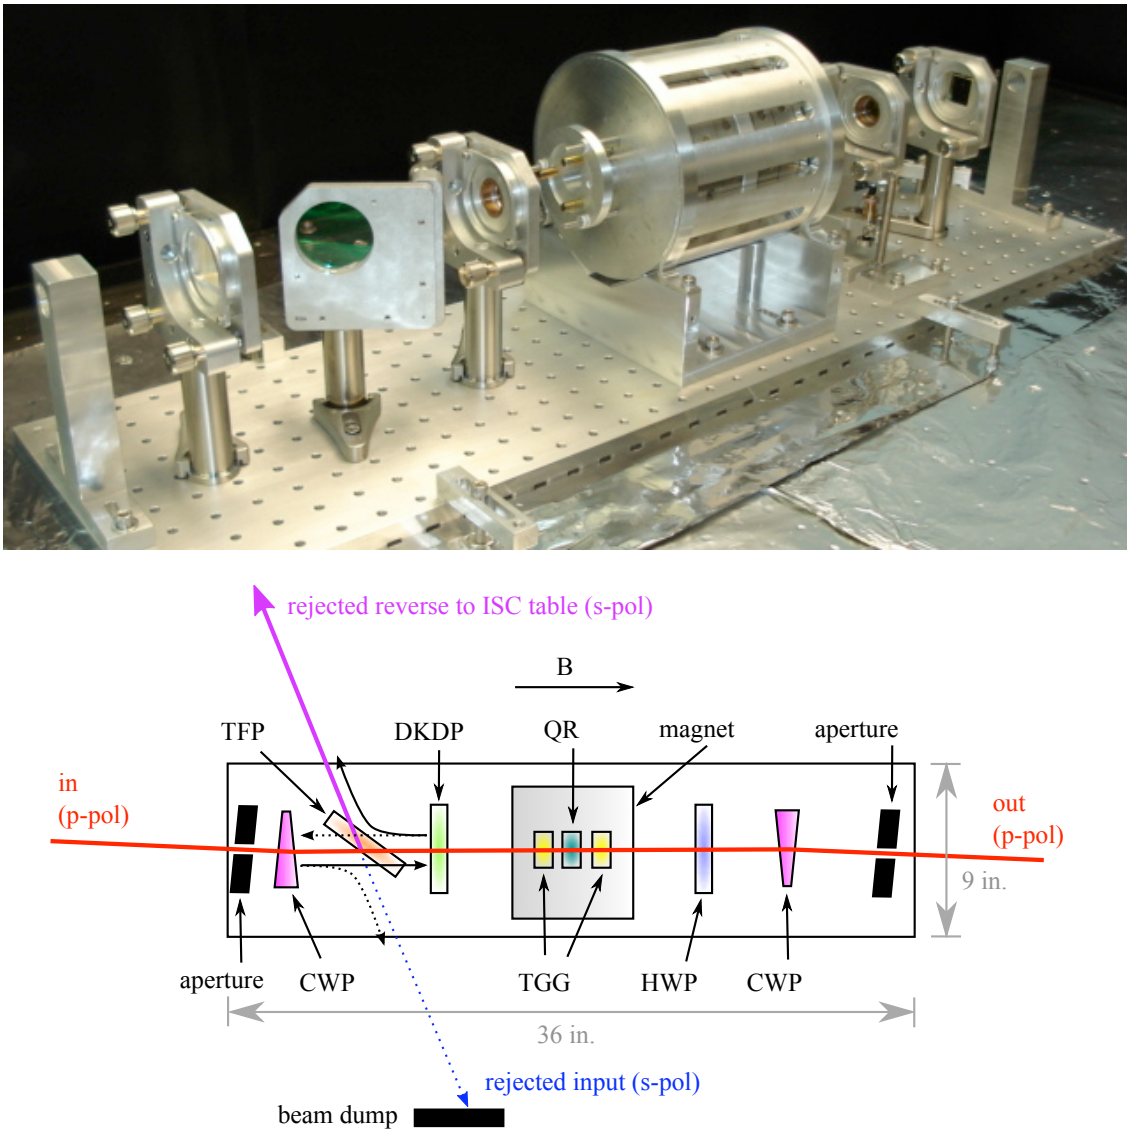

Figure 3-6. Faraday isolator photograph and schematic. The Faraday isolator preserves the polarization of the light in the forward-going direction and rotates it by 90 degrees in the reverse direction. Light from the MC enters from the left and exits at the right towards the interferometer. It is ideally p-polarized, but any s-polarization contamination is promptly diverted $\sim$ 10 mrad by the CWP and then reflected by the TFP and dumped. The p-polarized reflected beam from the interferometer enters from the right and is rotated to s-polarized light which is picked-off by the TFP and sent to the Interferometer Sensing and Control (ISC) table. Any imperfections in the Faraday rotation of the interferometer return beam results in p-polarized light traveling backwards along the original input path. Photograph courtesy of Katherine Dooley.

will be mostly undone upon passing through the second. The multiple elements in the magnet required a larger magnetic field than in Initial LIGO. The 7-disc magnet is 130 mm in diameter and 132 mm long and placed in housing 155 mm in diameter and 161 mm long. The TGG diameter is 20 mm.

3.3.3.2 Thermal lensing

Thermal lensing in the Faraday isolator is addressed by including DKDP, a negative dn/dT material, in the beam path. Absorption of light in the DKDP results in a de-focusing of the beam, which partially compensates for the thermal focusing induced by absorption in the TGGs [46, 47]. The optical path length (thickness) of the DKDP is chosen to slightly over-compensate the positive thermal lens induced in the TGG crystals, anticipating other positive thermal lenses in the system.

3.3.3.3 Polarizers

The polarizers used (two CWPs and one TFP) each offer advantages and disadvantages related to optical efficiency in the forward-propagating direction, optical isolation in the reflected direction, and thermal beam drift. The CWPs have very high extinction ratios ($> 10^5$) and high transmission ($> 99\%$) contributing to good optical efficiency and isolation performance. However, the angle separating the exiting orthogonal polarizations of light is very small, on the order of 10 mrad. This small angle requires the light to travel relatively large distances before we can pick off the beams needed for interferometer sensing and control. In addition, thermally induced index of refraction gradients due to the 4.95° wedge angle of the CWPs result in thermal drift. However, the CWPs for the Enhanced LIGO Faraday have a measured low absorption of 0.0013 cm^{-1} with an expected thermal lens of 60 m at 30 W and drift of less than 1.3 μrad/W [39].

The advantages of the thin film polarizer over the calcite wedge polarizer are that it exhibits negligible thermal drift when compared with CWPs and it operates at the Brewster angle of 55°, thus diverting the return beam in an easily accessible way. However, the TFP has a lower transmission than the CWP, about 96%, and an extinction ratio of only 10^3.

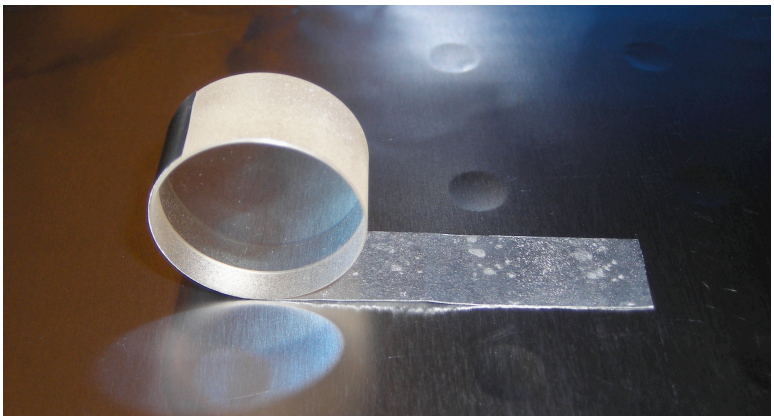

Figure 3-7. Photograph of TGG crystal with indium foil wrapping. Photograph courtesy of Katherine Dooley.

Thus, the combination of CWPs and a TFP combines the best of each to provide a high extinction ratio (from the CWPs) and ease of reflected beam extraction (from the TFP). The downsides that remain when using both polarizers are that there is still some thermal drift from the CWPs. Also the transmission is reduced due to the TFP and to the fact that there are 16 surfaces from which light can scatter.

3.3.3.4 Heat conduction

Faraday isolators operating in a vacuum environment suffer from increased heating with respect to those operating in air. Convective cooling at the faces of the optical components is no longer an effective heat removal channel, so proper heat sinking is essential to minimize thermal lensing and depolarization. It has been shown that Faraday isolators carefully aligned in air can experience a dramatic reduction in isolation ratio ($>$ 10-15 dB) when placed in vacuum [48]. The dominant cause is the coupling of the photoelastic effect to the temperature gradient induced by laser beam absorption. Also of importance is the temperature dependence of the Verdet constant–different spatial parts of the beam experience different linear polarization rotations in the presence of a temperature gradient [49].

To improve heat conduction away from the Faraday rotator optical components, we designed housing for the TGG and quartz crystals that provided improved heat sinking to the Faraday rotator. We also wrapped the TGGs with indium foil as pictured in Figure 3-7 to improve contact

with the housing, and we cushioned the DKDP and the HWP with indium wire in their aluminum holders. This has the additional effect of avoiding the development of thermal stresses in the crystals, an especially important consideration for the very fragile DKDP.

3.3.4 Mode-matching Telescope Design

The mode matching into the interferometer (at Livingston) was measured to be at best 90% in Initial LIGO. Because of the stringent requirements placed on the LIGO vacuum system to reduce phase noise through scattering by residual gas, standard opto-mechanical translators are not permitted in the vacuum; it is therefore not possible to physically move the mode matching telescope mirrors while operating the interferometer. Through a combination of needing to move the MMTs in order to fit the new Faraday isolator on the in-vacuum optics table and additional measurements and models to determine how to improve the coupling, a new set of MMT positions was chosen for Enhanced LIGO. Fundamental design considerations are discussed in Ref. [50].

3.4 Performance of the Enhanced LIGO Input Optics

The most convincing figure of merit for the Input Optics performance is that the Enhanced LIGO interferometers achieved low-noise operation with 20 W input power without thermal issues from the IO. Additionally, the Input Optics were operated successfully up to the available 30 W of power. (Instabilities with other interferometer subsystems limited the Enhanced LIGO science run operation to 20 W.) We present in this section detailed measurements of the Input Optics performance during Enhanced LIGO. Specific measurements and results presented in figures and the text come from Livingston; performance at Hanford was similar and is included in tables summarizing the results.

3.4.1 Optical Efficiency

The optical efficiency of the Enhanced LIGO Input Optics from EOM to recycling mirror was 75%, a marked improvement over the approximate 60% that was measured for Initial LIGO. A substantial part of the improvement came from the discovery and subsequent correction of a 6.5% loss at the second of the in-vacuum steering mirrors directing light into the MC (refer to

Table 3-2. Enhanced LIGO Input Optics power budget. Errors are ±1%, except for the TFP loss whose error is ±0.1%. The composite mode cleaner transmission is the percentage of power after the MC to before the MC and is the product of the MC visibility and transmission. Initial LIGO values, where known, are included in parentheses and have errors of several percent.

	Livingston	Hanford
Mode cleaner visibility	92%	97%
Mode cleaner transmission	88%	90%
Composite MC transmission	81% (72%)	87%
Faraday transmission	93% (86%)	94% (86%)
- Thin film polarizer loss	4.0%	2.7%
IO efficiency (PSL to RM)	75% (60%)	82%

Figure 3-3). A 45° reflecting mirror had been used for a beam with an 8° angle of incidence. Losses attributable to the mode cleaner and Faraday isolator are described in the following sections. A summary of the IO power budget is found in Table 3-2.

3.4.1.1 Mode cleaner losses

The mode cleaner was the greatest single source of power loss in both Initial and Enhanced LIGO. The mode cleaner visibility, defined here as

$$\text{visibility} = \frac{P_{\text{in}} - P_{\text{reflected}}}{P_{\text{in}}}, \tag{3–1}$$

the ratio of the amount of light coupled into the MC to the amount impinging the MC input mirror, was 92%. Visibility reduction is the result of higher order mode content of P_{in} and mode mismatch into the MC. The visibility was constant within 0.04% up to 30 W input power at both sites, providing a positive indication that thermal aberrations in the MC and upstream were negligible.

Of the light coupled into the MC, 88% was transmitted, corresponding to an average loss of 98 ppm per mirror. The scatter loss is expected to be 22 ppm/mirror based on the mirrors' measured root mean square surface microroughness of $\sigma_{rms} < 0.4$ nm [51]. Part of the discrepancy between expectation and measurement was determined to come from poor or damaged AR coatings. We measured a 1.3% reflection from the AR coatings on MC mirrors at

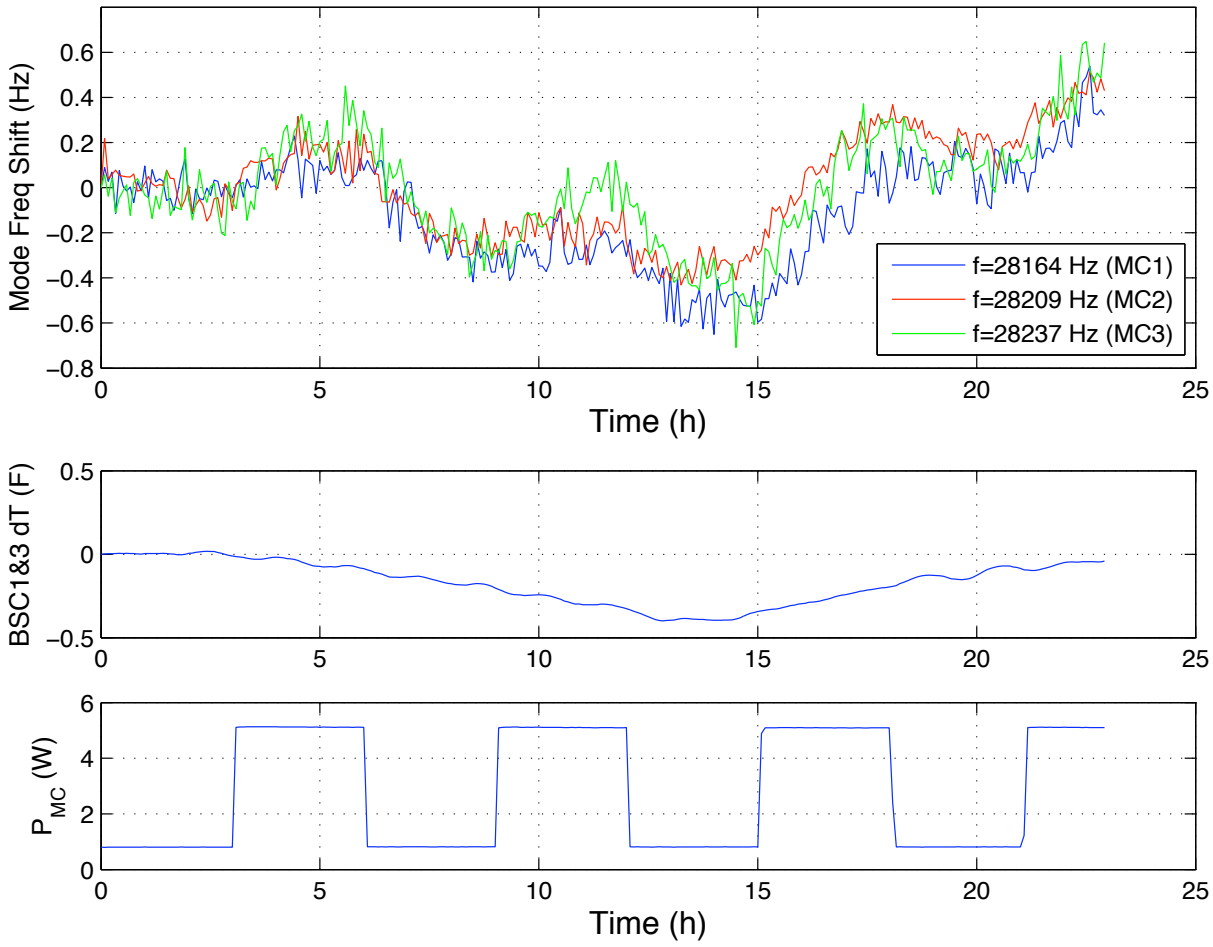

Figure 3-8. Data from the mode cleaner absorption measurement. Power into the MC was cycled between 0.9 W and 5.1 W at 3 hour intervals (bottom frame) and the change in frequency of the drumhead mode of each mirror was recorded (top frame). The ambient temperature (middle frame) was also recorded in order to correct for its effects.

both Livingston and Hanford, a transmitted power loss equivalent to 10 ppm of intracavity loss per mirror.

Another source of MC losses is via absorption of heat by particulates residing on the mirror's surface. We measured the absorption with a technique that makes use of the frequency shift of the thermally driven drumhead eigenfrequencies of the mirror substrate [52]. The frequency shift directly correlates with the MC absorption via the substrate's change in Young's modulus with temperature, dY/dT. A finite element model (COMSOL) was used to compute the expected frequency shift from a temperature change of the substrate resulting from the mirror

Table 3-3. Absorption values for the Livingston and Hanford mode cleaner mirrors before (in parentheses) and after drag wiping. The precision is ±10%.

mirror	Livingston	Hanford
MC1	2.1 ppm (18.7 ppm)	5.8 (6.1 ppm)
MC2	2.0 ppm (5.5 ppm)	7.6 (23.9 ppm)
MC3	3.4 ppm (12.8 ppm)	15.6 (12.5 ppm)

coating absorption. The measured eigenfrequencies for each mirror at room temperature are 28164 Hz, 28209 Hz, and 28237 Hz, respectively.

We cycled the power into the mode cleaner between 0.9 W and 5.1 W at 3 hour intervals, allowing enough time for a thermal characteristic time constant to be reached. At the same time, we recorded the frequencies of the high Q drumhead mode peaks as found in the mode cleaner frequency error signal, heterodyned down by 28 kHz. Figure 3-8 shows the measurement data. Correcting for ambient temperature fluctuations, we find a frequency shift of 0.043, 0.043, and 0.072 Hz/W. As a result of drag-wiping the mirrors, the absorption decreased for all but one mirror, as shown for both Hanford and Livingston in Table 3-3.

3.4.1.2 Faraday isolator losses

The Faraday isolator was the second greatest source of power loss with its transmission of 93%. This was an improvement over the 86% transmission of the Initial LIGO FI. The most lossy element in the Faraday isolator was the thin film polarizer, accounting for 4% of total losses. The integrated losses from AR coatings and absorption in the TGGs, CWPs, HWP, and DKDP account for the remaining 3% of missing power.

3.4.2 Faraday Isolation Ratio

The isolation ratio is defined as the ratio of power incident on the Faraday in the reverse direction (the light reflected from the interferometer) to the power transmitted in the reverse direction and is often quoted in decibels: isolation ratio = $10\log_{10}(P_{in-reverse}/P_{out-reverse})$. We measured the isolation ratio of the Faraday isolator as a function of input power both in air prior to installation and *in situ* during Enhanced LIGO operation.

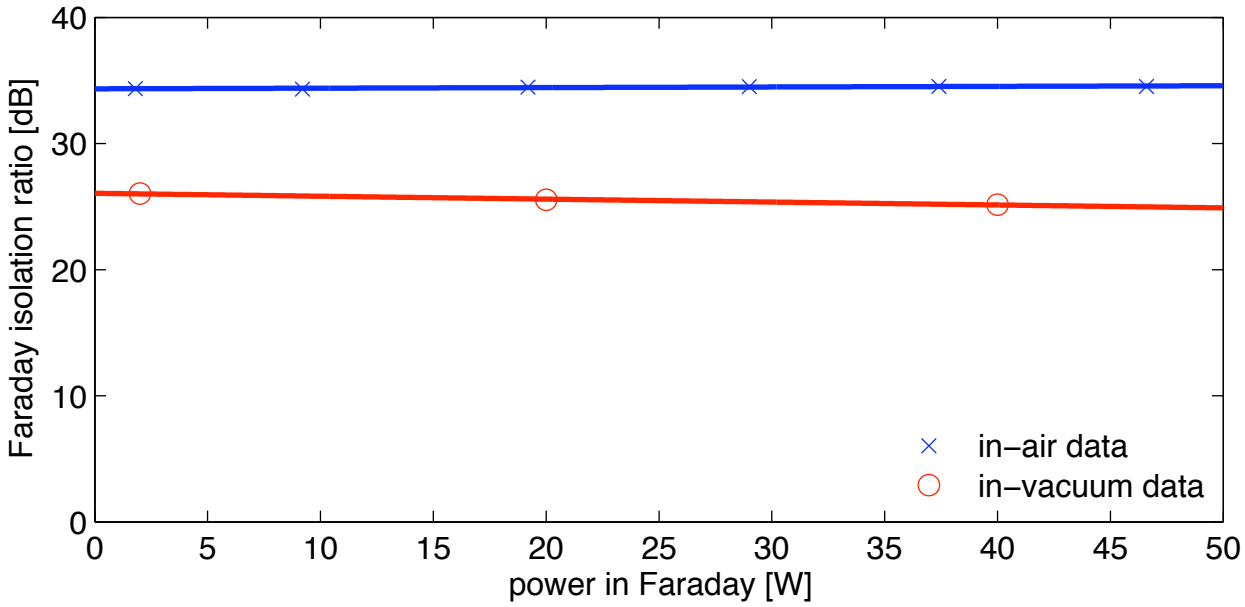

Figure 3-9. Faraday isolator isolation ratio as measured in air prior to installation and *in situ* in vacuum. The isolation worsens by a factor of 6 upon placement of the Faraday in vacuum. The linear fits to the data show a constant in-air isolation ratio and an in-vacuum isolation ratio degradation of 0.02 dB/W.

To measure the in-vacuum isolation ratio, we misaligned the interferometer arms so that the input beam would be promptly reflected off of the 97% reflective recycling mirror. This also has the consequence that the Faraday isolator is subjected to twice the input power. Our isolation monitor was a pick-off of the backwards transmitted beam taken immediately after transmission through the Faraday that we sent out of a vacuum chamber viewport. Refer to the "isolation check beam" in Figure 3-3. The in air measurement was done similarly, except in an optics lab with a reflecting mirror placed directly after the Faraday.

Figure 3-9 shows our isolation ratio data. Most notably, we observe an isolation decrease of a factor of six upon placing the Faraday isolator in vacuum, a result consistent with that reported by Ref. [48]. In air the isolation ratio is a constant 34.46 ± 0.04 dB from low power up to 47 W, and in vacuum the isolation ratio is 26.5 dB at low power. The underlying cause is the absence of cooling by air convection. If we attribute the loss to the TGGs, then based on the change in TGG polarization rotation angle necessary to produce the measured isolation drop of 8 dB and the temperature dependence of the TGG's Verdet constant, we can put an upper limit of 11 K

on the crystal temperature rise from air to vacuum. Furthermore, a degradation of 0.02 dB/W is measured in vacuum.

3.4.3 Thermal Steering

We measured the *in situ* thermal angular drift of both the beam transmitted through the mode cleaner and of the reflected beam from the Faraday isolator with up to 25 W input power. Just as for the isolation ratio measurement, we misaligned the interferometer arms so that the input beam would be promptly reflected off of the recycling mirror. The Faraday rotator was thus subjected to up to 50 W total and the MC to 25 W.

Pitch and yaw motion of the mode cleaner transmitted and interferometer reflected beams were recorded using the quadrant photodiode (QPD) on the Input Optics table and the RF alignment detectors on the Interferometer Sensing and Control table, as seen in Figure 3-3. There are no lenses between the MC waist and its measurement QPD, so only the path length between the two were needed to calibrate in radians the pitch and yaw signals on the QPD. The interferometer reflected beam, however, passes through several lenses. Thus, ray transfer matrices and the two alignment detectors were necessary to extract the Faraday drift calibration. Details of the calibration method are presented in Appendix A.4.

Figure 3-10 shows the calibrated beam steering data. The angle of the beam out of the mode cleaner does not change measurably as a function of input power in yaw (4.7 nrad/W) and changes by only 440 nrad/W in pitch. For the Faraday isolator, we record a beam drift originating at the center of the Faraday rotator of 1.8 µrad/W in yaw and 3.2 µrad/W in pitch. Therefore, when ramping the input power up to 30 W during a full interferometer lock, the upper limit on the drift experienced by the reflected beam is about 100 µrad. This is a thirty-fold reduction with respect to the Initial LIGO Faraday isolator and represents a fifth of the beam's divergence angle, $\theta_{div} = 490$ µrad.

3.4.4 Thermal Lensing

We measured the profiles of both the beam transmitted through the mode cleaner and the reflected beam picked off by the Faraday isolator at low ($\sim$ 1 W) and high ($\sim$ 25 W) input

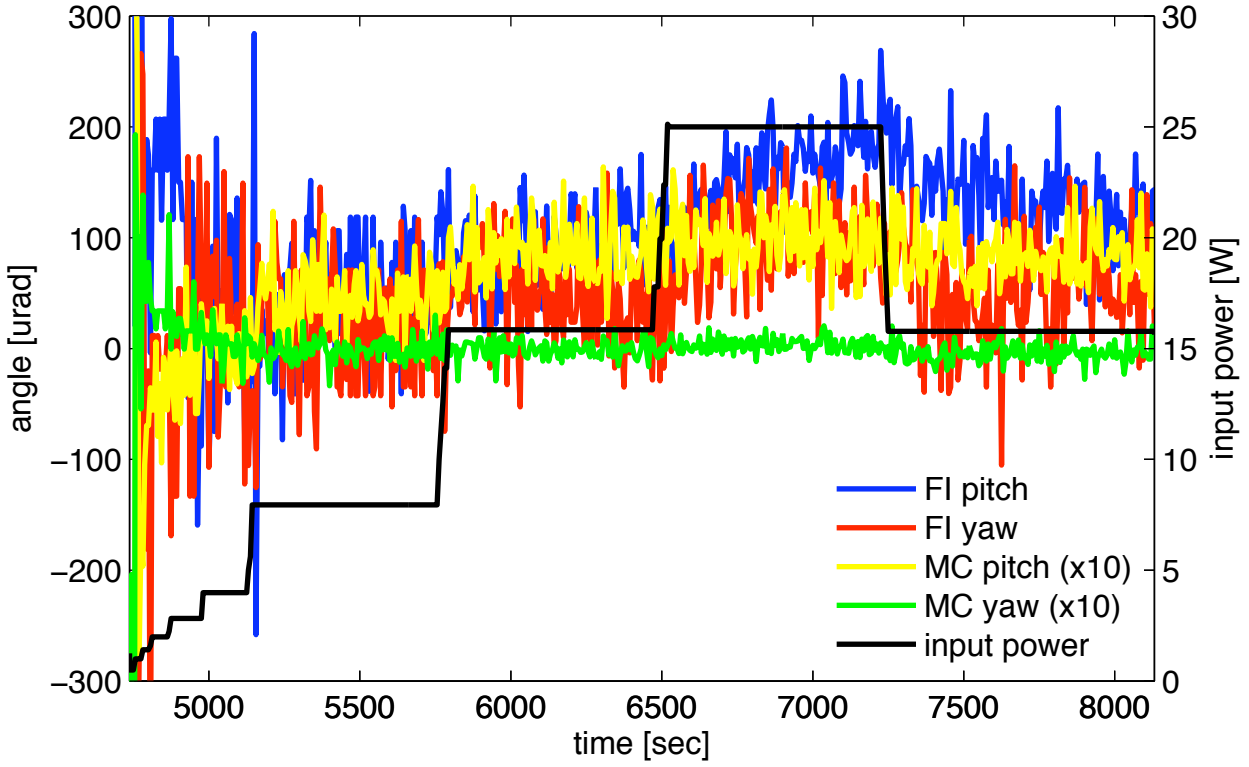

A

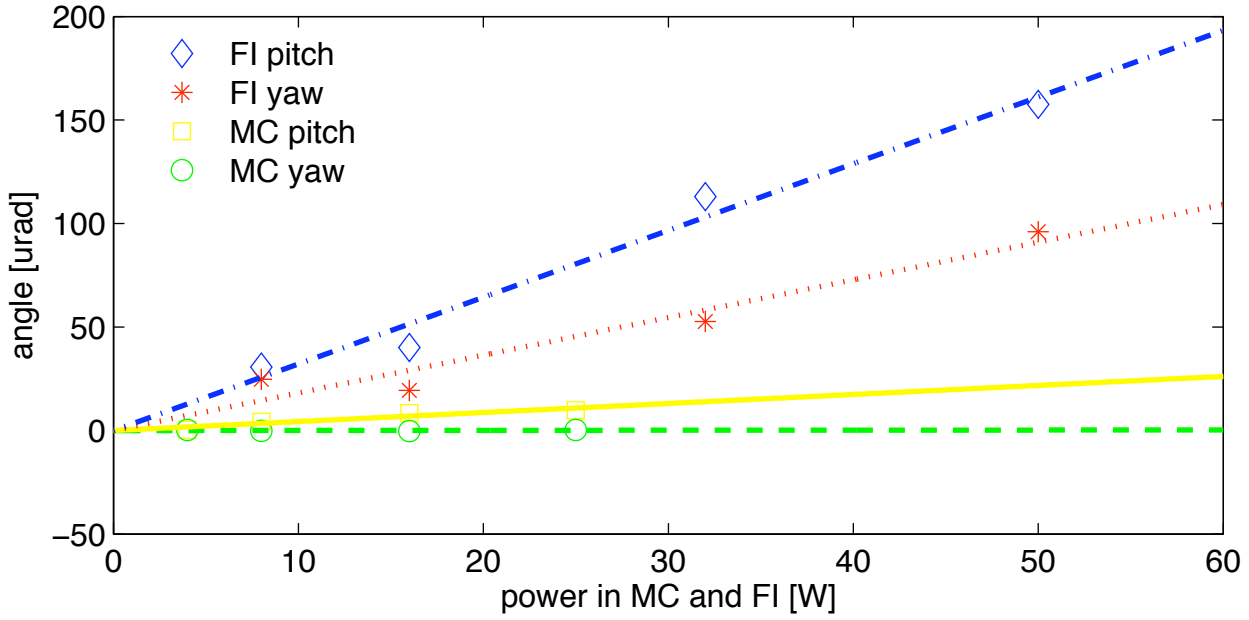

B

Figure 3-10. Mode cleaner and Faraday isolator thermal drift data. A) Angular motion of the beam at the MC waist and FI rotator as the input power is stepped. The beam is double-passed through the Faraday isolator, so it experiences twice the input power. B) Average beam angle per power level in the MC and FI. Linear fits to the data are also shown. The slopes for MC yaw, MC pitch, FI yaw, and FI pitch, respectively, are 0.0047, 0.44, 1.8, and 3.2 μrad/W.

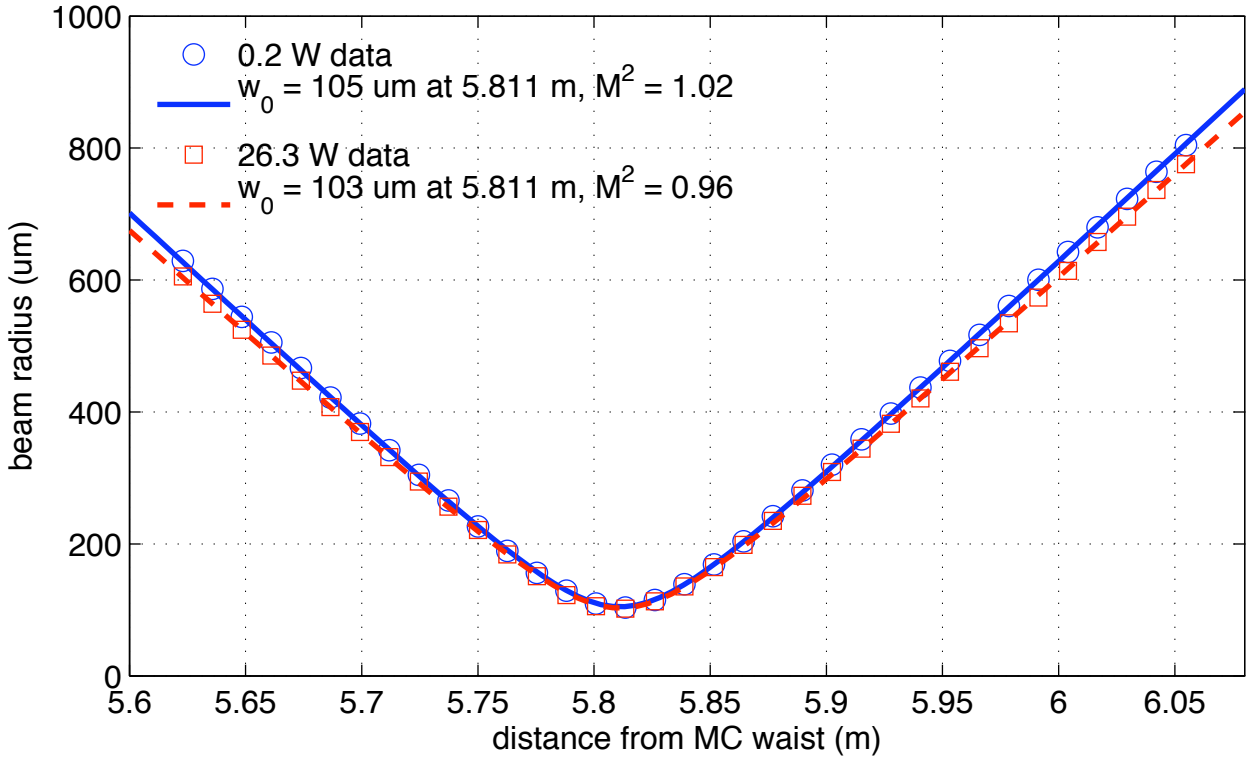

Figure 3-11. Profile at high and low powers of a pick-off of the beam transmitted through the mode cleaner. The precision of the beam profiler is ±5%. Within the error of the measurement, there are no obvious degradations.

powers to assess the degree of thermal lensing induced in the MC and FI. Again, we misaligned the interferometer arms so that the input beam would be promptly reflected off the recycling mirror. We picked off a fraction of the reflected beam on the Interferometer Sensing and Control table and of the mode cleaner transmitted beam on the Input Optics table (refer to Figure 3-3), placed lenses in each of their paths, and measured the beam diameters at several locations on either side of the waists created by the lenses. A change in the beam waist size or position as a function of laser power indicates the presence of a thermal lens.

As seen in Figure 3-11 and Figure 3-12, the waists of the two sets of data are collocated: no thermal lens is measured. For the Faraday isolator, the divergence of the low and high power beams differs, indicating that the beam quality degrades with power. The M^2 factor at 1 W is 1.04 indicating the beam is nearly perfectly a TEM_{00} mode. At 25 W, M^2 increases to 1.19, corresponding to increased higher-order-mode content. The percentage of power in higher-order

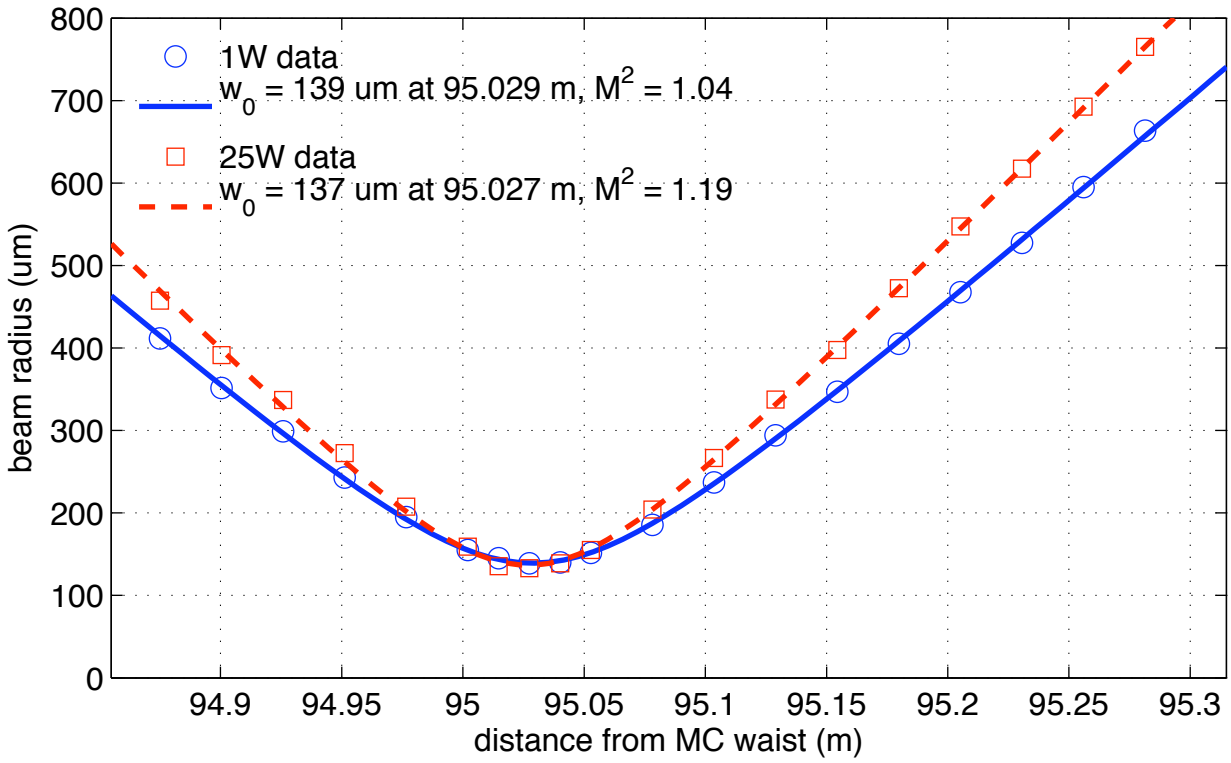

Figure 3-12. Faraday isolator thermal lensing data. With 25 W into the Faraday isolator (corresponding to 50 W in double pass), the beam has a steeper divergence than a pure TEM_{00} beam, indicating the presence of higher order modes. Errors are $\pm 5.0\%$ for each data point.

modes depends strongly on the mode order and relative phases of the modes, and thus cannot be determined from this measurement [53].

The results for the mode cleaner data are consistent with no thermal lensing. The high and low power beam profiles are within each other's error bars and well below our requirements.

3.4.5 Mode-matching

We measured the effectiveness of the mode-matching telescope by taking the ratio of power at the reflected port when all of the interferometer cavities are on resonance to the power in the reflected beam when the cavities are unlocked. Since the impedance matching is near perfect, all light at the reflected port during interferometer lock is attributable to a mode mismatch. Initially, anywhere between 10% and 17% of the light was rejected by the cavity due to poor, power-dependent mode matching. After translating the mode-matching telescope mirrors during

a vacuum chamber incursion and upgrading the other IO components, the ratio we measured was 8% independent of input power. The MMT succeeds at coupling 92% of the light into the interferometer at all times, marking both an improvement in MMT mirror placement and success in eliminating measurable thermal issues. Appendix A.6 presents details of the mode-matching measurement.

3.5 Implications for Advanced LIGO

As with other Advanced LIGO interferometer components, Enhanced LIGO served as a technology demonstrator for the Advanced LIGO Input Optics, albeit at lower laser powers than will be used there. The performance of the Enhanced LIGO Input Optics components, at 20 W of input power allows us to infer their performance in Advanced LIGO. The requirements for the Advanced LIGO Input Optics demand are for similar performance to Enhanced LIGO, but with almost 8 times the laser power.

The Enhanced LIGO electro-optic modulator showed no thermal lensing, degraded transmission, nor damage in over 1 million hours of sustained operation at 30 W of laser power. Measurements of the thermal lensing in RTP at powers up to 160 W show a relative power loss of $< 0.4\%$, indicating that thermal lensing should be negligible in Advanced LIGO. Peak irradiances in the EOM will be approximately four times that of Enhanced LIGO (a 45% larger beam diameter will somewhat offset the increased power). Testing of RTP at 10 times the expected Advanced LIGO irradiance over 100 hours show no signs of damage or degraded transmission.

The mode cleaner showed no measurable change in operational state as a function of input power. This bodes well for the Advanced LIGO mode cleaner. Compared with the Enhanced LIGO mode cleaner, the Advanced LIGO mode cleaner is designed with a lower finesse (520) than Initial LIGO (1280). For 150 W input power, the Advanced LIGO mode cleaner will operate with 3 times greater stored power than Initial LIGO. The corresponding peak irradiance is 400 kW/m^2, well below the continuous wave coating damage threshold. Absorption in the Advanced LIGO mode cleaner mirror optical coatings has been measured at 0.5 ppm, roughly four times less than the best mirror coating absorption in Enhanced LIGO, so the expected

thermal loading due to coating absorption should be reduced in Advanced LIGO. The larger Advanced LIGO mode cleaner mirror substrates and higher input powers result in a significantly higher contribution to bulk absorption, roughly 20 times Enhanced LIGO, however the expected thermal lensing leads to small change ($< 0.5\%$) in the output mode [42].

The Enhanced LIGO data obtained from the FI allows us to make several predictions about how it will perform in Advanced LIGO. The measured isolation ratio decrease of 0.02 dB/W will result in a loss of 3 dB for a 150 W power level expected for Advanced LIGO relative to its cold state. However, the Advanced LIGO FI will employ an *in situ* adjustable half wave plate which will allow for a partial restoration of the isolation ratio. In addition, a new FI scheme to better compensate for thermal depolarization and thus yield higher isolation ratios may be implemented [54]. The maximum thermally induced angular steering expected is 480 μrad (using a drift rate of 3.2 μrad/W), approximately equal to the beam divergence angle. This has some implications for the Advanced LIGO length and alignment sensing and control system, as the reflected FI beam is used as a sensing beam. Operation of Advanced LIGO at high powers will likely require the use of a beam stabilization servo to lock the position of the reflected beam on the sensing photodiodes. Although no measurable thermal lensing was observed (no change in the beam waist size or position), the measured presence of higher order modes in the FI at high powers is suggestive of imperfect thermal lens compensation by the DKDP. This fault potentially can be reduced by a careful selection of the thickness of the DKDP to better match the absorbed power in the TGG crystals.

3.6 Summary

In summary, we have presented a comprehensive investigation of the Enhanced LIGO Input Optics, including the function, design, and performance of the IO. Several improvements to the design and implementation of the Enhanced LIGO IO over the Initial LIGO IO have lead to improved optical efficiency and coupling to the main interferometer through a substantial reduction in thermo-optical effects in the major IO optical components, including the electro-optic modulators, mode cleaner, and Faraday isolator. The IO performance in Enhanced LIGO enables

us to infer its performance in Advanced LIGO, and indicates that high power interferometry will be possible without severe thermal effects.

CHAPTER 4
ANGULAR MOTION OF THE INTERFEROMETER MIRRORS

For light to resonate in the interferometer, the mirrors need to point at one another and remain stationary with respect to this pointing. It is necessary to actively align the mirrors for several reasons:

- to find the optimal DC alignment of the interferometer
- to suppress any motion that results from external disturbances
- to counteract a high power static instability

There are 16 angular degrees of freedom (DOFs) that need to be sensed and controlled, making the ASC one of the most complex interferometer subsystems. This chapter presents the ASC design, and discusses the causes of mirror angular displacement and the effects of residual mirror motion on the interferometer.

4.1 Tolerance for Angular Motion

In a most general sense, it is the controlled angular motion of the mirrors that matters, which is in contrast with DARM for which both the controlled and uncontrolled length displacements matter. The residual angular motions must be small enough to both keep the inteferometer locked and to not impact strain sensitivity. The requirements for how much motion is tolerable stem from two effects of misalignment that directly couple to strain sensitivity: failure to achieve maximum power gain, and angle to length coupling. The misalignment tolerances are dictated by what is necessary to prevent the strain sensitivity of the perfectly aligned interferometer from degrading by more than 0.5% in the detection band of 40 to 7000 Hz [55].

Because the strain sensitivity is proportional to the power at the beam splitter (refer to Eq. 2–13), a decrease in circulating power directly results in a decrease of shot-noise-limited DARM. Furthermore, differing power fluctuations in the two arm cavities results in a changing contrast defect, a difference in the amount of light returning from one arm compared to the other. A changing contrast defect creates power fluctuations at the AS port, making it indistinguishable from gravitational waves. Also, the DC power of the contrast defect contributes to increasing the shot noise noise-floor. Power fluctuations couple quadratically to DARM, and in the power

recycling cavity in particular, power fluctuations make for inconsistent signal to noise ratios for the signals that depend on sideband power. To maintain a power buildup within 1% of maximum, the core optics must have an angular displacement of less than 10^{-8} rad rms with respect to the cavity axis [56]. The derivation of power buildup as a function of mirror angle displacement is found in Appendix C.3.

Another deleterious effect of poor alignment is angle to length coupling. When the axis of rotation of a mirror coincides with the center of the beam, any tilt of the mirror about this axis does not affect the path length of the reflected beam. However, if there is a mismatch between rotation axis and beam location, then the light will pick up a longitudinal phase shift when the mirror is tilted. During a full interferometer lock, this is recorded by DARM. Explicitly, when the beam is located a distance x away from the center of the mirror, an angular displacement of the mirror θ about its center results in a path length change of the beam

$$\Delta L = x\theta. \tag{4–1}$$

Therefore, the alignment specifications must include not only tolerable levels of angular motion, but requirements for the physical centering of the beam spots on the mirrors. As detailed in Ref. [56], the beams must be centered on the core optics within 1 mm. At DC, for x = 1 mm and $\theta = 10^{-7}$ rad, $\Delta L = 10^{-10}$ m which is four orders of magnitude below the DARM rms of 10^{-6} m.

4.2 Sources of Angular Mirror Motion

There is a continuous stream of changing torque inputs to the mirrors that cause them to twist and turn in pitch and in yaw. Some torque inputs exist regardless of the state of the interferometer, while others are a direct consequence of the control systems. The primary torque inputs are introduced here, and further discussion of some of them is found later in the chapter. The list includes:

- ground motion
- coil actuators (length to angle)
- noise impression from the angular control system
- radiation pressure

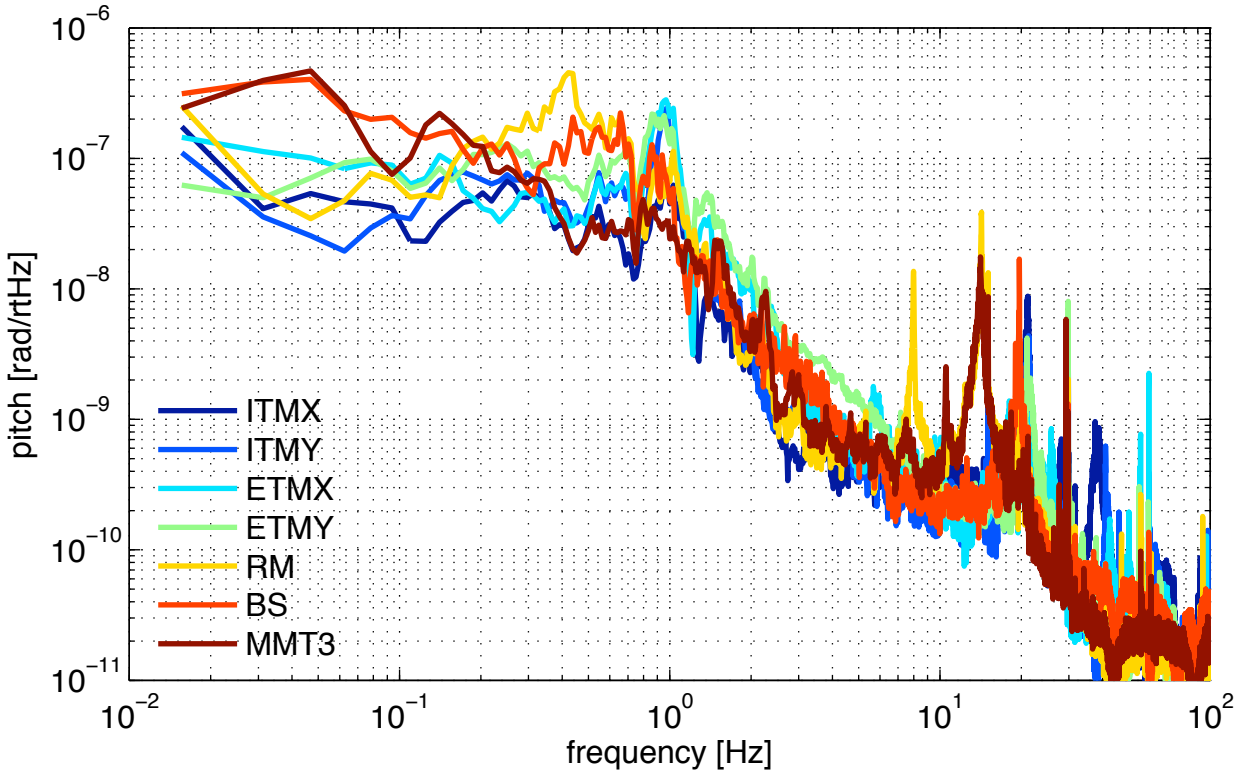

Figure 4-1. Typical angular motion of the core suspended mirrors in the absence of interferometric control. Velocity damping provided by the OSEMs and the optical levers is present. Once the interferometer is locked, the OSEM damping is turned off. Ground motion at the time of this measurement is shown in Figure C-4.

4.2.1 Ground Motion

The most egregious of these torque inputs is ground motion that makes its way through the multiple stages of seismic isolation to the mirror suspensions and thence to the mirrors. Ground motion is the only source of angular motion that is present regardless of the state of interferometer operation. An example of the shape and amount of angular motion experienced by the core optics due to seismic noise during a relatively quiet seismic time is shown in Figure 4-1.[1] . The rms mirror motion is of the order 10^{-7} rad. This is the motion that needs to be controlled interferometrically.

[1] These spectra include the effect of local velocity damping because the optical levers are always on. See Section 2.4.3.

I obtained these spectra using the optical levers as witnesses of the mirror motion, which in turn is expected to come from the ground's motions. The optical lever signal, may, however, be contaminated by optical lever sensing (electronics) noise or acoustic couplings. It is therefore informative to embark on a study of how much of the optical lever spectra are in fact due to ground motion.

I used a method of Wiener filtering [57] to show that the motion of Figure 4-1, indeed originates primarily from the ground. The Wiener filter is the best estimate of the contribution of one signal to another. The contributions of the three seismometers to each large optic optical lever are shown in Figure 4-2. We see that the optical lever signals are almost entirely explained by ground motion, with the notable exceptions being a peak between 3 and 4 Hz for ITMY, and 0.3 and 0.7 Hz for the RM, and broadband extra motion for the BS.

Furthermore, the relative magnitudes of the contributions from the three seismometers make sense. Greater than 1 Hz, coherence is greatest with the nearest seismometer, and less than 1 Hz, seismometers contribute more equally (except for the ITMs, which are everywhere very coherent with the corner station seismometer).

4.2.2 Coil Actuators

Imperfections in the forces applied by the actuators on the rear of the test masses can convert piston (pure translation) drives into torque. The lengths of the cavities are carefully controlled (that's what we strive to be most sensitive to!) and any imbalances between the four electromagnets on a single mirror will create a coupling from length drive to angle (L2A). This effect is measurable, and then is carefully tuned out through selecting appropriate digital gains for each of the coils. Relative gain variation is 10% of average. The ability to tune the gains perfectly is limited and the residual length to angle coupling is about 1%. For the typical rms length drive of 1 μm on a core optic and OSEMs separated by a distance of $\sqrt{2}R$ where R is the radius of the optic, the 1% L2A coupling results in a 10^{-8} radian displacement:

$$\theta = \frac{0.01 * 10^{-6}\ \text{m}}{\sqrt{2} * 0.125\ \text{m}} \approx 10^{-8}\ \text{rad.} \tag{4–2}$$

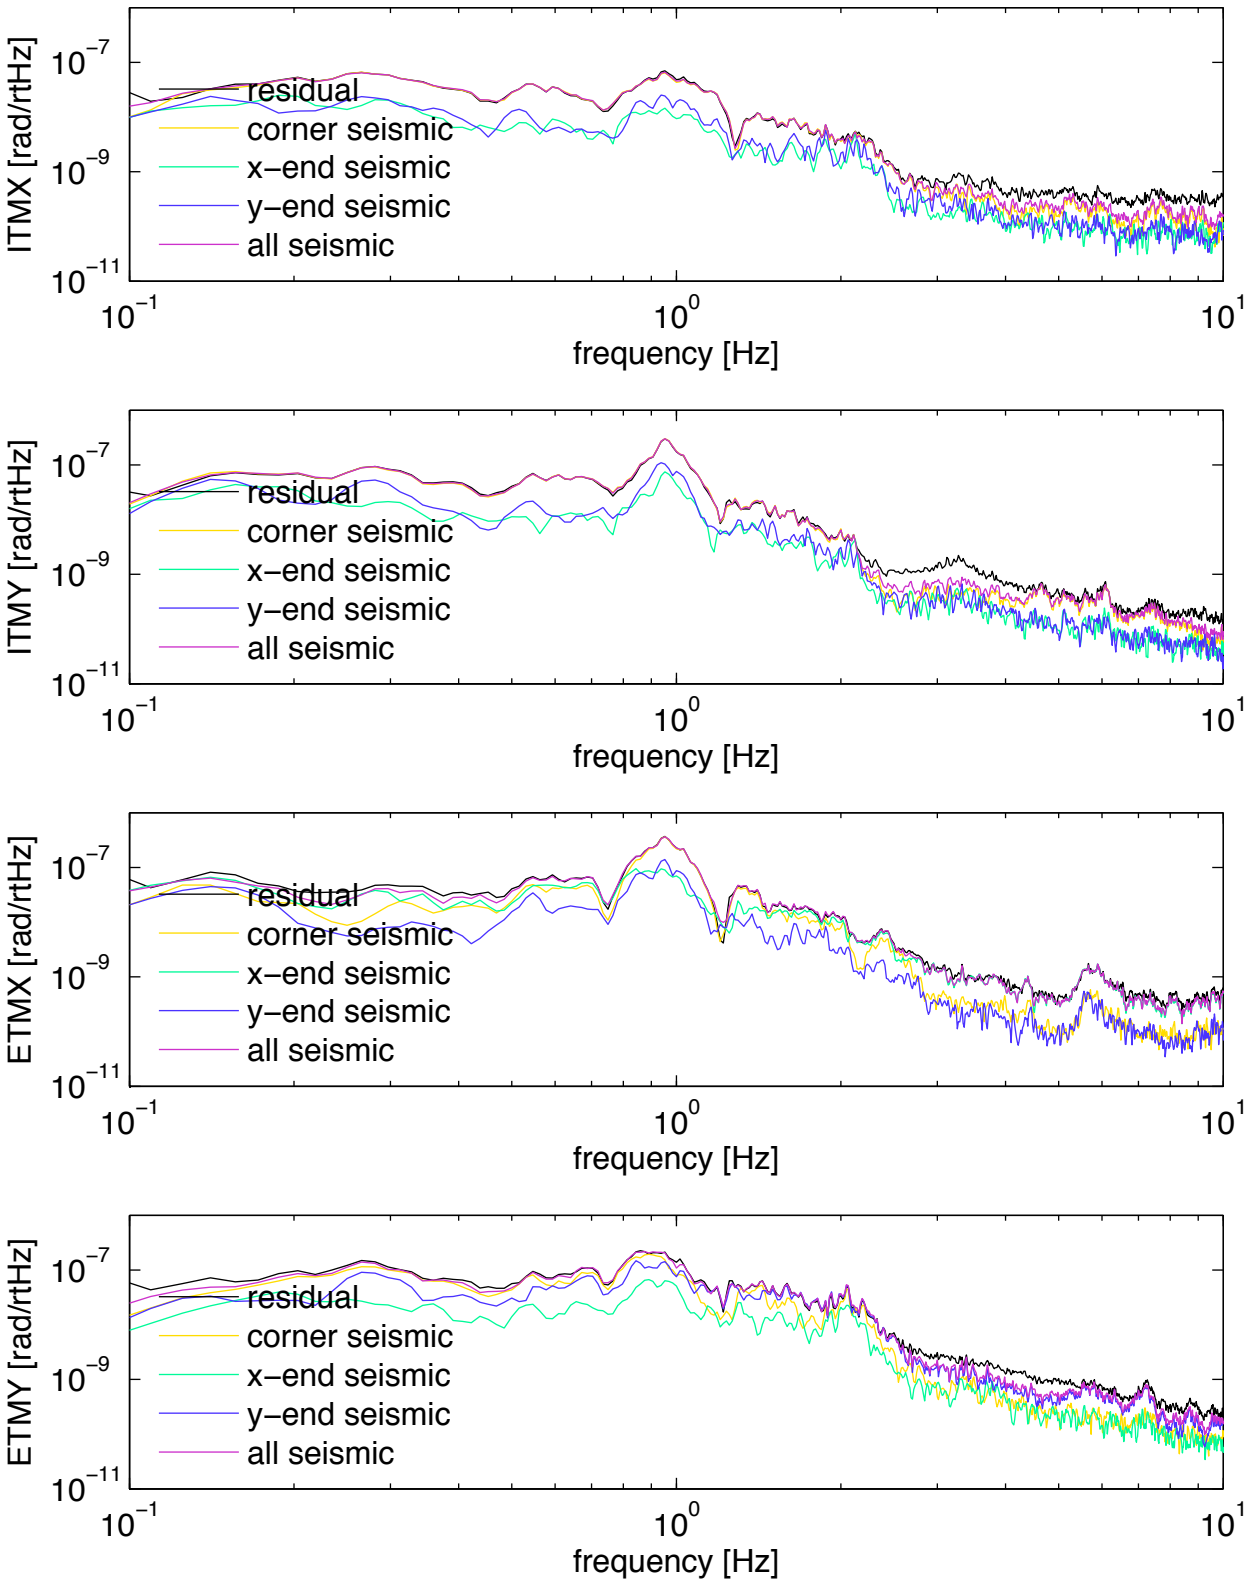

Figure 4-2. Contribution of seismic noise to optical lever error signal (called residual). The interferometer was unlocked and optical lever and OSEM AC damping present.

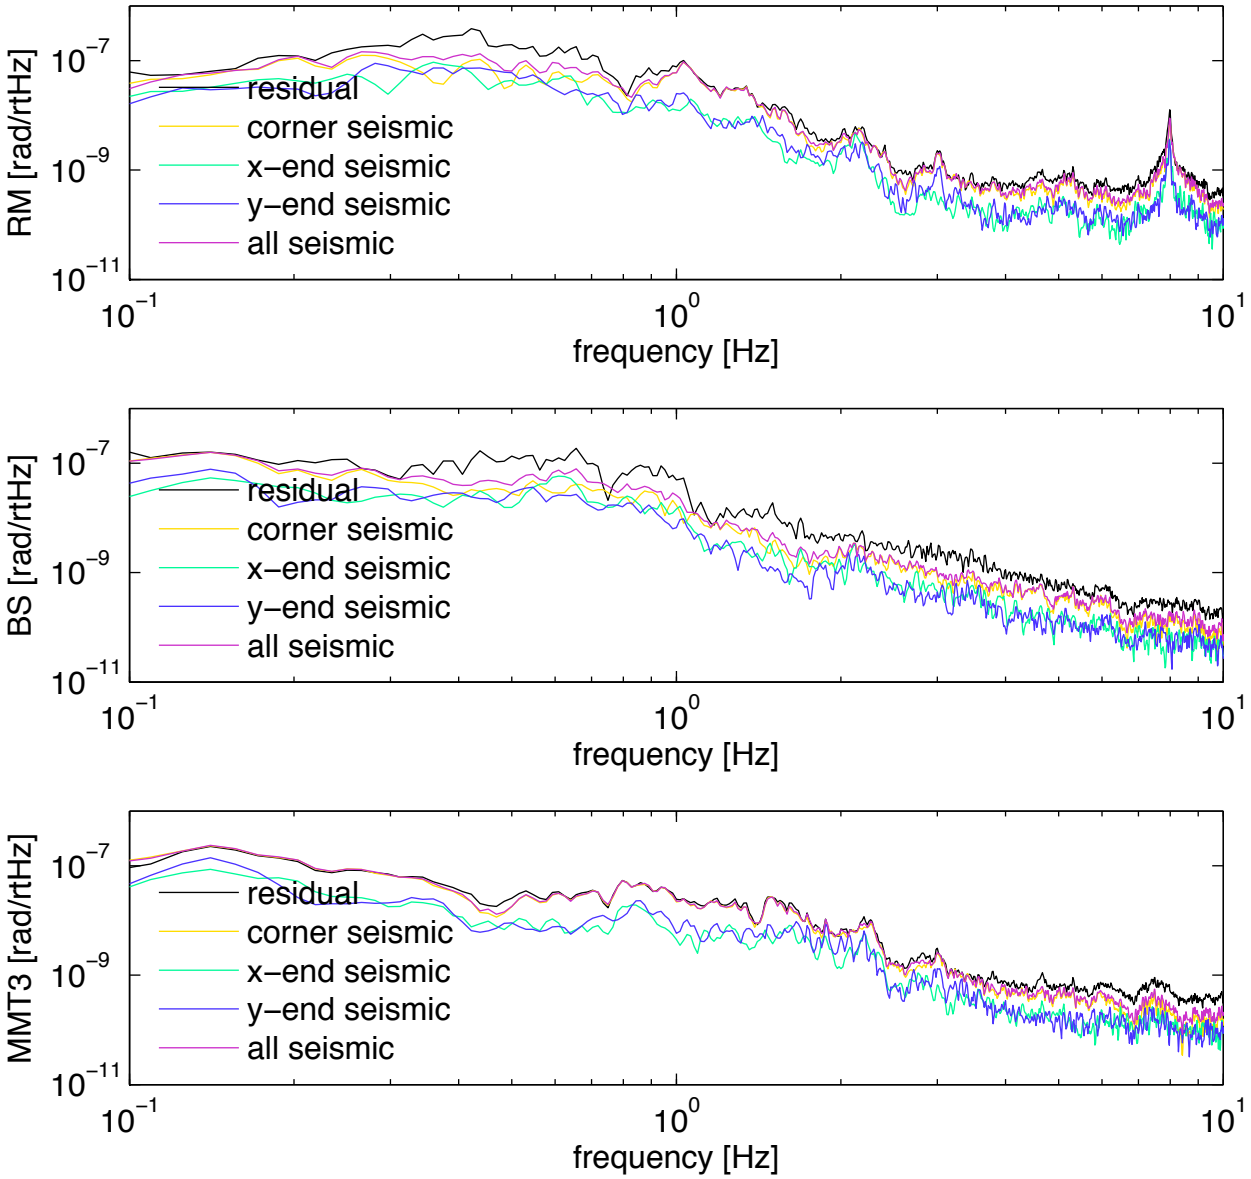

Figure 4-2 continued.

4.2.3 Noise from Angular Control

The angular control system, which strives to counteract the above torque inputs to reduce angular motion, introduces angular mirror motion itself. The primary way it contributes noise is through imperfect sensing of the angular displacements. The alignment is also under-controlled, which ends up allowing the control system to impress input beam motion on the mirrors. These issues are explained in more detail in Section 4.6.

4.2.4 Radiation Pressure

Radiation pressure creates a torque when the beam impinges the mirror off-center. The force on the mirror due to radiation pressure is derived from the change in momentum of a photon upon reflection off the mirror and results in:

$$F_{rp} = \frac{2P}{c} \tag{4–3}$$

where P is the power of the light reflected by the mirror and c is the speed of light. Because the beam of photons strikes the mirror perpendicular to its surface, the torque exerted on a mirror due to radiation pressure is

$$\tau_{rp} = \frac{2Px}{c} \tag{4–4}$$

where x is the distance of the beam from the mirror's center of mass. For a 40 kW beam 1 mm off-center, the torque is on the order of 10^{-7} Nm, yielding an angular displacement of the order 10^{-7} rad as determined by the pendulum torque to angle transfer function. Refer to Section 4.3.

Amongst the various torque inputs, radiation pressure plays a unique role in mirror motion because the torque it exerts depends on the angles of the mirrors. A result of the geometric coupling between beam displacements and mirror angles, radiation pressure therefore acts as an angular spring. It is best treated not as an external torque, but as a modification to the pendulum torque to angle transfer function. Chapter 5 dedicates a discussion to the physics of radiation pressure torques. In all, radiation pressure shapes the angular dynamics of the mirrors in LIGO and plays an important role in the design of an angular control system.

4.3 The Mirror as a Torsion Pendulum

In order to design a control system that reduces the angular motion of the interferometer mirrors to the levels necessary for stable interferometer operation and minimal impact on strain sensitivity, the angular response of the mirrors to external torque must be fully understood. A model of the mirror's torque to angle transfer function is thus required. To start, the mirrors in LIGO may be regarded as torsion pendula. The mirror may twist an angle θ about a horizontal

axis passing through its center of mass to create motion in *pitch* and about a vertical axis passing through its center of mass to create a motion in *yaw*.

The angular equation of motion of the mirror is governed by the sum of all torques on the mirror. First, let's consider the most simplistic scenario where there is only a pendulum restoring torque $\tau_p = -\kappa_p \theta$, where κ_p is the pendulum's torsional constant. The equation of motion is

$$I\ddot{\theta} + \kappa_p \theta = 0, \tag{4–5}$$

which has a solution of $\theta(t) = \sin(\omega_0 t)$, where $\omega_0 = \sqrt{\kappa_p / I}$ is the resonant angular frequency and I is the mirror's moment of inertia. The pendulum torsional constant serves to make the mirror oscillate indefinitely about its equilibrium position upon the slightest of displacements.

We are particularly interested in the pendulum's angular response to an external torque, such as seismic noise. In order to calculate the torque to angle transfer function, we must include an external torque term, τ_{ext}, in the equation of motion:

$$I\ddot{\theta} + \gamma\dot{\theta} + \kappa_p \theta = \tau_{ext}. \tag{4–6}$$

A velocity damping term, γ, is also included. Taking the Laplace transform to convert from the time domain to the frequency domain, we have:

$$Is^2\Theta + \gamma s\Theta + \kappa_p \Theta = \tau_{ext} \tag{4–7}$$

where s is a complex parameter. The transfer function is then defined as

$$H(s) := \frac{\Theta(s)}{\tau_{ext}(s)} = \frac{1}{Is^2 + \gamma s + \kappa_p}. \tag{4–8}$$

We are only interested in examining the transfer function for a pure sine wave excitation, $e^{i\omega t}$, so we substitute $s = i\omega$ to get

$$H(\omega) = \frac{1/I}{\omega_0^2 - \omega^2 + i\gamma\omega/I}. \tag{4–9}$$

The resonant frequency of this damped system can be computed by finding the ω at which the amplitude of the transfer function, $[I^2[\omega^2-\omega_0^2]^2+\gamma^2\omega^2]^{-1/2}$, is maximized:

$$\omega_{res}=\sqrt{\omega_0^2-\frac{\gamma^2}{2I^2}}. \tag{4–10}$$

We note that damping reduces the resonant frequency, although the effect is usually insignificant.

A quantity that is more familiar than γ for describing the losses of a system with a real resonance is the quality factor, $Q := \omega_{res}/\text{FWHM}$, where FWHM is the full-width-half-max of the transfer function's amplitude-squared resonance. When the losses are small, $\omega_{res}\approx\omega_0$ and FWHM $\approx \gamma/I$ [58, 23-4]. The quality factor is then well approximated by $Q=\sqrt{\kappa_p I}/\gamma$. The transfer function written in terms of Q is

$$H(\omega)=\frac{1/I}{\omega_0^2-\omega^2+i\omega\omega_0/Q}. \tag{4–11}$$

Figure 4-3 shows the pendulum torque to angle transfer function (for pitch) using the parameters of a LIGO core optic. For external torques applied to the mirror well above its resonant frequency, the mirror acts like a free mass, one that is not held in place by suspension wires nor subject to damping. For torques applied to the mirror below its resonant frequency, the mirror's angle is determined by the inverse of the torsional constant.

4.4 Overview of Interferometer Alignment

There are 8 mirrors whose pitch and yaw angles must be sensed and controlled. The sensing is accomplished by 8 sensors, which fall into three groups:

- camera image (BS): senses the position of the beam on the BS
- quadrant photodiodes (QPDX, QPDY): sense the position of the light transmitted through the ETMs
- wavefront sensors (WFS1, WFS2[2] , WFS3, WFS4): sense the angular misalignment of the cavities with respect to their input beams

[2] Two signals are derived from WFS2.

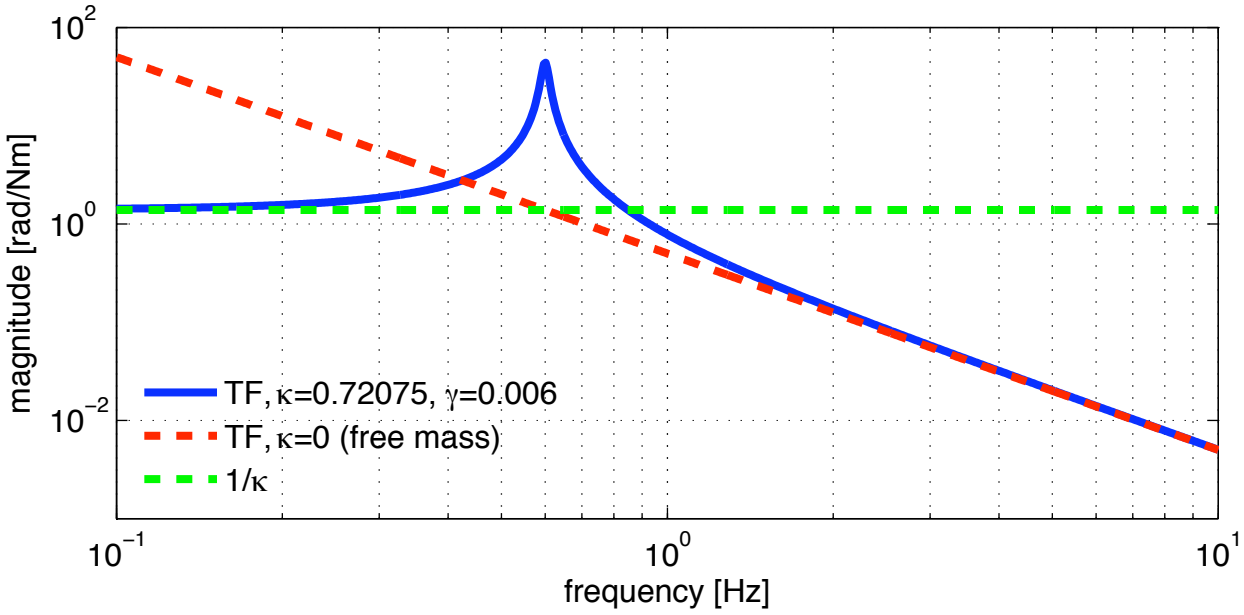

Figure 4-3. Torque to pitch transfer function of a LIGO core optic (blue). The optic acts like a free mass at high frequencies (red) and the DC magnitude of the transfer function is determined by the inverse torsional constant (green). A damping constant $\gamma = 0.006$ ($Q = 32$) was selected for pictorial representation only. The resonant frequency of LIGO core optics in yaw is 0.5 Hz.

Figure 4-4 shows the locations of these sensors and the 8 mirrors they must control. The alignment scheme can be simplified by considering it in two basic units: the input beam and the power-recycled Fabry-Perot Michelson (FPM). The alignment of the latter is self-contained (with the exception of the BS) and uses the WFS to keep the mirrors aligned to one another from DC up to several Hz. The alignment of the input beam to the power recycled FPM unit is accomplished via the BS camera image and QPDX at frequencies well below the pendulum resonances. Finally, as part of the input beam alignment, QPDY keeps the beam that is reflected from the BS pointed at the y-arm. The alignment scheme is depicted in Figure 4-5 and explained in more detail in the remainder of this section.

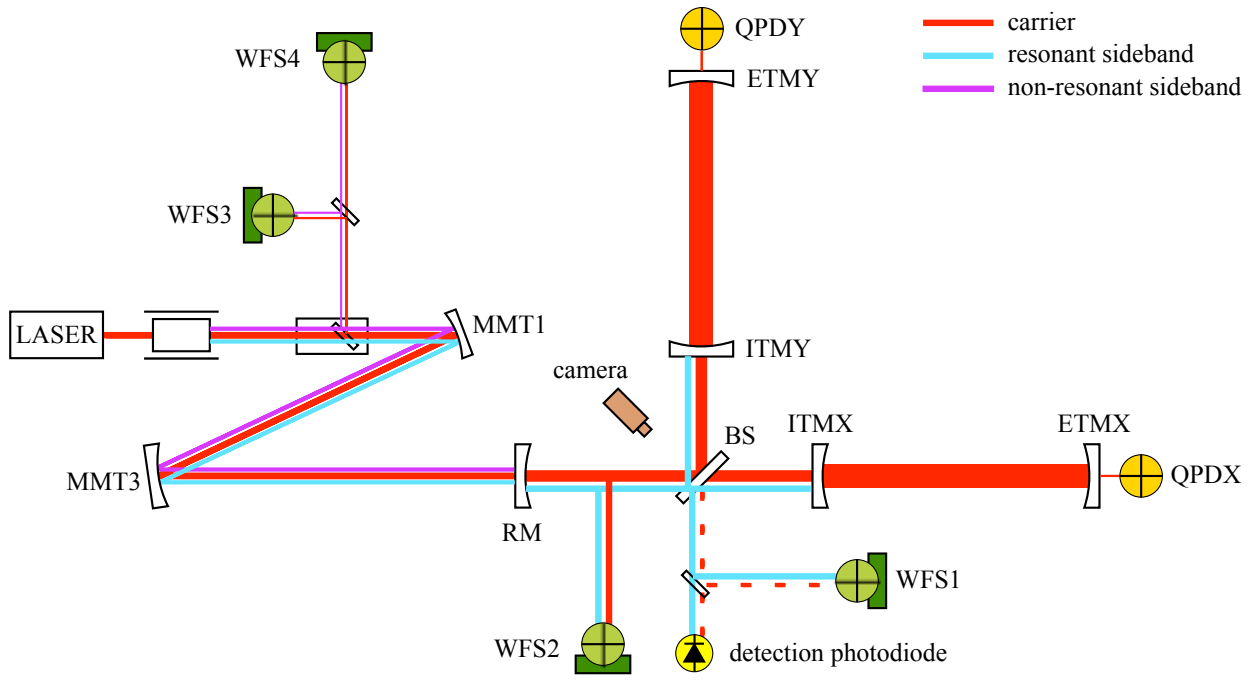

Figure 4-4. Layout of ASC sensors (WFS, QPDs, and camera) and the 8 mirrors they must control (ITMs, ETMs, MMTs, BS, and RM).

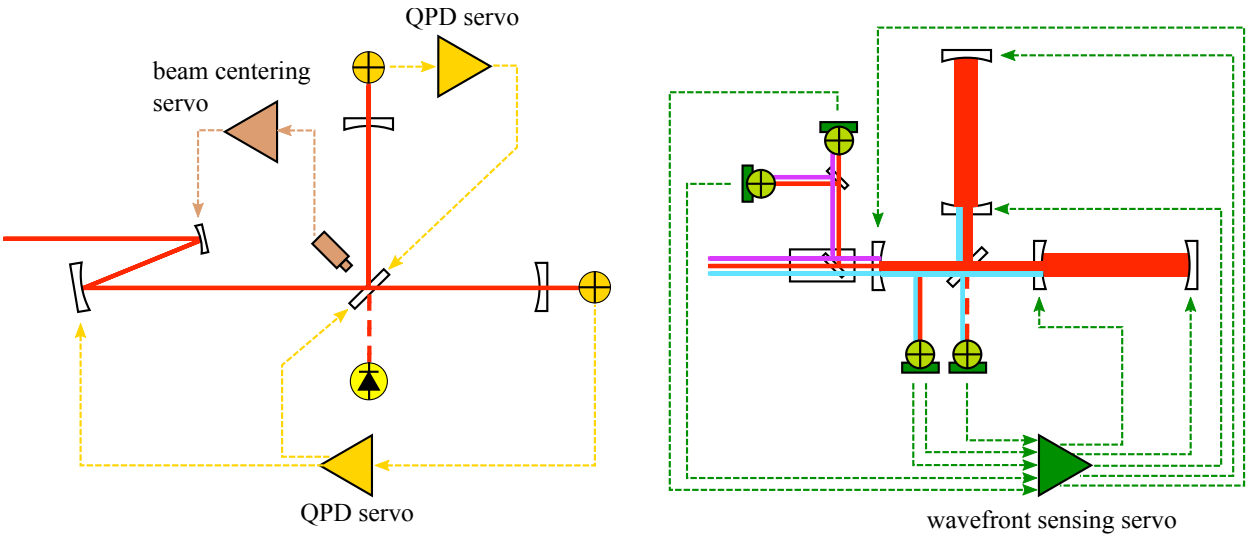

Figure 4-5. Schematic of the alignment sensing and control system, viewed as two different units. A) The QPD servo and beam centering servo (BCS) together direct the input beam on minute time scales. B) The wavefront sensing servo maintains the alignment of the power-recycled FPM mirrors with respect to each other up to several Hz.

The self-contained alignment of the power-recycled Fabry-Perot Michelson is realized through the set of wavefront sensors (WFS)[3] which provide the most sophisticated form of measuring angular motion of the mirrors and whose principle of operation is described in Section 4.5.1. The pitch and yaw motions of the five mirrors in this unit, ETMX, ETMY, ITMX, ITMY, and the PRM, are sensed by the pitch and yaw of five WFS signals, WFS1Q, WFS2I, WFS2Q, WFS3I, and WFS4I, where I and Q denote in-phase and quadrature demodulation, respectively. These WFS look at light at the AS port, at the reflected port, and in the power recycling cavity. Their error signals are used to control the relative motions of these five mirrors up to a couple Hz. The original design of wavefront sensing for a power-recycled Michelson with Fabry-Perot cavities in the arms was created and tested by Mavalvala [59, 60].

The MMT-directed input beam and the power-recycled FPM unit need to be aligned to one another such that the input beam is perfectly reflected upon itself. Over minute time scales, the input beam follows the interferometer and over faster time scales the interferometer follows the input beam. This is achieved via a blend of all sensors. The low frequency matching of the input beam to the interferometer is realized through the pitch and yaw signals of QPDX, a QPD which monitors the position of the light transmitted through the x-arms, and the pitch and yaw signals of a camera that monitors the location of the beam spot on the beam splitter. These two alignment sensors adjust the pointing of MMT1 and MMT3 at about 30 mHz. The camera's beam centering servo (BCS) works by taking the image of the speckle of light reflected off of the beam splitter as shown in Figure 4-6 and feeding it into a labview program which integrates the intensity of the image to identify the coordinates of the center of the beam spot. These coordinates are compared with a hardcoded desired center location. A mirror upstream, MMT1, is moved to redirect the input beam, minimizing the difference between the desired and actual beam spot location on the BS. The QPD servos work in a similar fashion. The higher frequency matching of the input beam

[3] Pronounced "woofs".

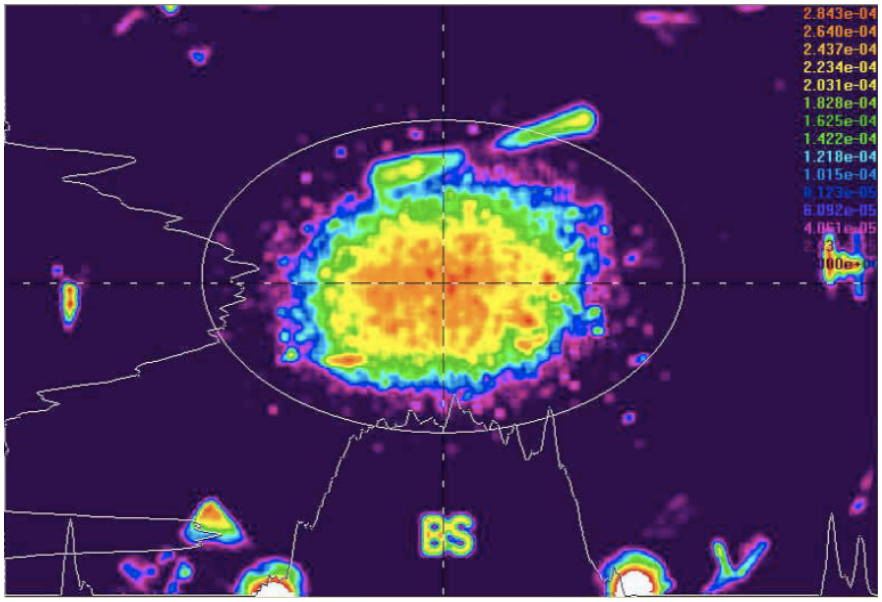

Figure 4-6. Image of beam on beam splitter as used in the beam centering servo. The beam appears stretched because the camera's viewing angle is at 45° with respect to the mirror surface. The color scale is arbitrary, but orange is strong, violet is weak.

and the interferometer to one another is achieved by the reflected port wavefront sensors (WFS3 and WFS4), up to a couple Hz.

The one additional step needed for full interferometer alignment is to maintain the relative alignment of the y-arm to the x-arm as the x-arm and input beam together move around. This is accomplished through the pitch and yaw signals of QPDY, the QPD that monitors the light transmitted through the y-arm, which sense how the beam splitter should be pointed. QPDX also sends a signal to the BS to compensate for the signal it sends to MMT3.

All mirror angles are of course inter-dependent and they must track each other. However, a rough hierarchy of who follows who can be established because ultimately the interferometer is bolted to the ground and necessarily maintains some DC orientation. This orientation comes from QPDX and QPDY, which are physically attached to piers standing on the ground and force the beams transmitted through the ETMs to stay put at a certain location on their sensors. In all, the input beam must make it to those two exact places and the other mirrors are left to line themselves up accordingly.

This alignment process involving the WFS, QPDs, and BS camera relies on the entire interferometer already being locked. It manages the continuous fine-tuning of mirror angles so that maximal power buildup in the interferometer is maintained, and so that the interferometer does not wander from its linear operating point (described in Section 2.2.1). How to achieve the initial alignment of all of the mirrors is an interesting process in itself and is documented in Appendix C.4.

4.5 The Angular Sensing and Control Servo

When the interferometer is locked, the optical levers and the wavefront sensors provide simultaneous feedback to the mirror angles. The optical levers provide local damping, and the WFS maintain both the alignment of the power-recycled FPM with the input beam and the internal alignment of the power-recycled FPM. The study and characterization of the elements of this angular control system when the laser power is increased is the topic of the remainder of this chapter and the next. First, it is helpful to present an overview of the servo without the effects of radiation pressure.

Figure 4-7 is a block diagram of the major components to the angular control servo. The torsion-like pendula of the interferometer are subject to external torque which is converted to mirror angle. The interferometer, in turn, turns the mirror angles into $\mathrm{TEM}_{10}/\mathrm{TEM}_{01}$ modes at its various ports, and these modes are converted into error signals by the wavefront sensors. The voltages produced by the wavefront sensors are digitized and manipulated by the front-end computers to create control signals that are then converted back to analog signals for actuation on the mirrors. The result of this process is a suppression of the error signal seen by the WFS and therefore a suppression of the physical relative mirror motions.

4.5.1 The Wavefront Sensing Scheme

The wavefront sensors provide the most sophisticated form of measuring angular motion of the mirrors. They are quadrant photodiodes equipped with RF electronics that rely on the Pound–Drever–Hall locking scheme to produce error signals [61]. WFS are the angular equivalent of the single element PDs used for length sensing. Their frame of reference is the fundamental

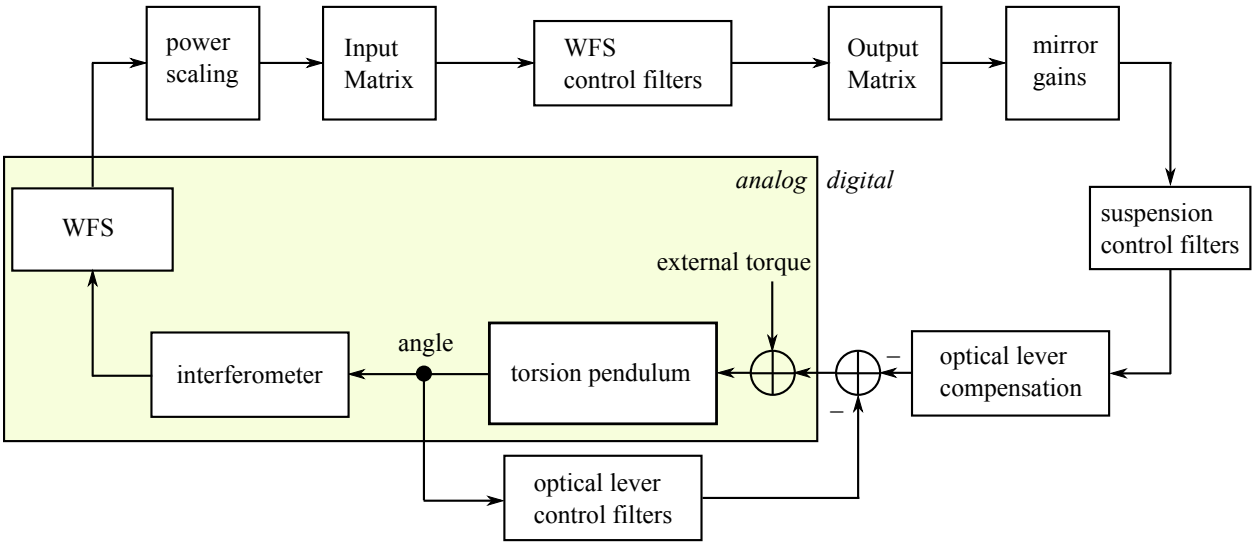

Figure 4-7. ASC control servo, showing both the WFS and optical lever loops. The external torque, torsion pendulum, interferometer, and WFS are analog; all other blocks are digital. Mirror angles are converted by the interferometer into signals the wavefront sensors can detect. These WFS error signals are digitized, filtered, and converted into analog control signals for individual mirrors. The details of each of the digital blocks are explained in Chapter 6.

Gaussian mode of the interferometer cavities; angular misalignments of the cavities generate TEM_{01} and TEM_{10} spatial modes of the carrier light as explained in Ref. [62] and derived in Appendix C.3. The overlap of the higher order mode carrier field with a TEM_{00} reference sideband field produces an excess of power on one half of the detector compared to the other when demodulated at the sideband frequency. This is the error signal. The sidebands

Reference [63] contains a thorough description of the wavefront sensing scheme. In summary, the idea is that the WFS are located where optical gain, the amount of laser power produced at some interferometer port for a given physical change of some aspect of the interferometer, is high. The WFS optical gain in, particular, is determined by how much TEM_{01}/TEM_{10} mode shows up at the locations of the WFS when a mirror (or specific combination of mirrors) moves in angle. One WFS is placed at the anti-symmetric port where differential-mode signals are transmitted, one is placed at a pick-off of the beam from the recycling cavity which contains

common and differential information, and two are placed at the reflected port, where common-mode signals are transmitted. The precise locations are determined from the Gouy phases of the light and details are found in Ref. [28].

4.5.2 The Digital Path

The digital portion of the servo, from the power scaling through the optical lever compensation, experiences two changes of basis which are effected by the input matrix and the output matrix. The input matrix takes the appropriate linear combination of WFS error signals to create error signals for which control filters may be designed easily. The output matrix creates a linear combination of the digital control signals in order to distribute them appropriately to specific mirrors. Both the power scaling and the mirror gains are diagonal matrices that provide minor modifications to the amplitude of the error and control signals, respectively. The power scaling is updated in real time to compensate for any changes in laser power on the WFS. The mirror gains are hard-coded scaling factors to reflect deviations from the theoretical cavity geometry. See Sections 6.4.1 and 6.1.3, respectively, for details.

4.5.3 Optical Lever Compensation

Each of the large optics has its own optical lever servo that provides velocity damping. Because both the WFS and the optical levers control the mirror motion when the interferometer is locked, we must consider the interaction of the two servos. Whether one views the interaction as the WFS controlling the optical-lever-controlled mirrors or as the optical levers controlling the WFS-controlled-mirrors is arbitrary. For the purpose of explanation here, I use the former viewpoint.

Figure 4-8 zooms in on the portion of the Figure 4-7 control loop picture that shows the optical lever controlling the pendulum. The torque to angle transfer function of the pendulum is our plant, P, the optical lever control is O, and the optical lever compensation is shown as $1 + OP$. A result of basic control theory is that the optical lever controlled pendulum loop can be

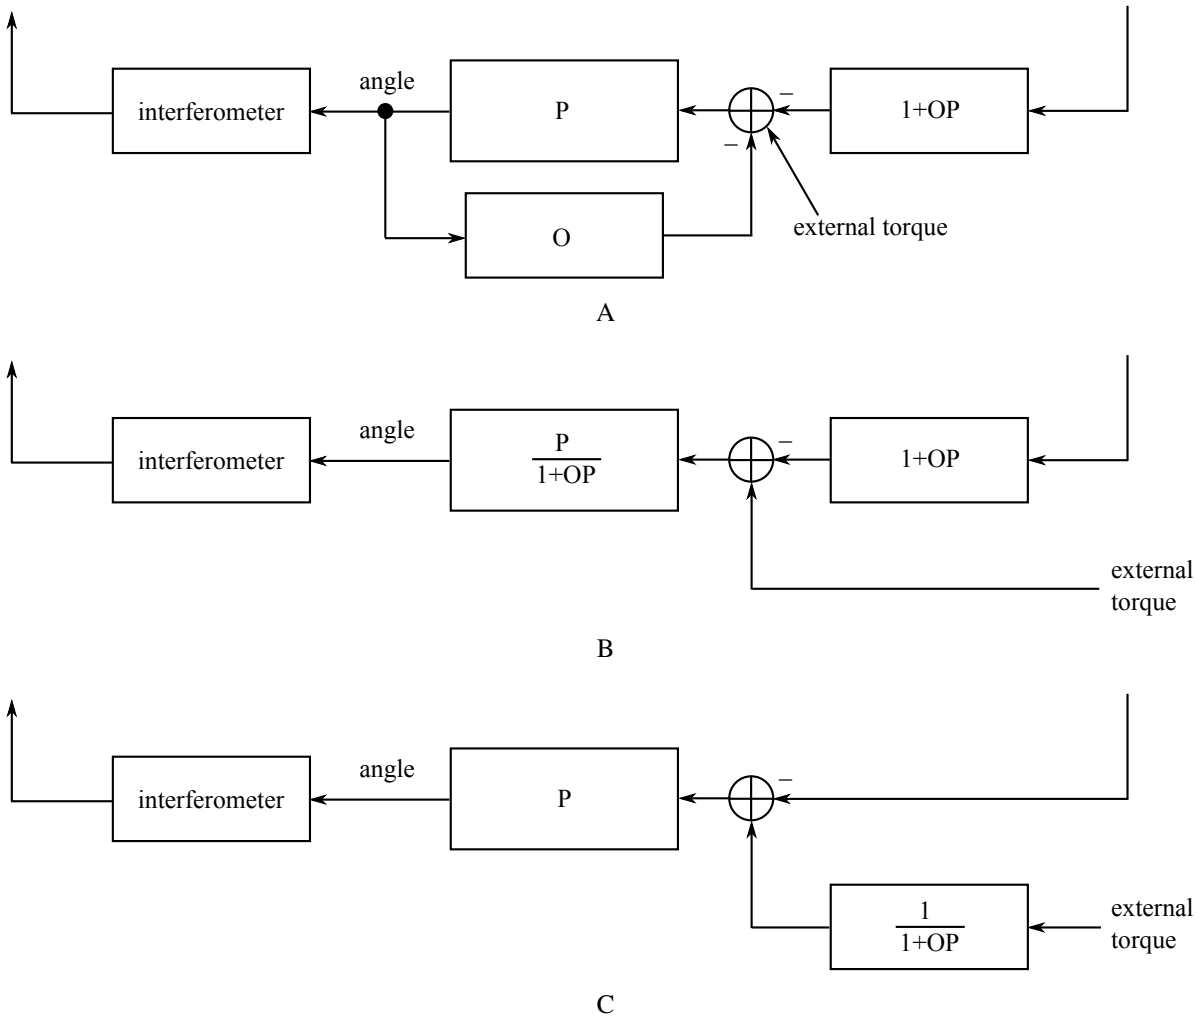

Figure 4-8. Optical lever compensation scheme. Subfigures A, B, and C are equivalent. A) Zoom of the bottom part of Figure 4-7. *P* is the pendulum and *O* is the optical lever digital control filters. B) Replacement of the optical lever controlled plant with the closed loop gain representation. C) Optical lever compensation removes the optical levers from the picture of the pendulum. External torque is still suppressed by the optical lever closed loop gain.

represented by its closed loop transfer function:

$$\mathrm{TF}_{\mathrm{closed}} = \frac{P}{1+OP} \tag{4–12}$$

as shown in Figure 4-8 B.

One will notice that it is not a coincidence that the optical lever compensation has the same form as the denominator of the closed loop transfer function. Together, they cancel each

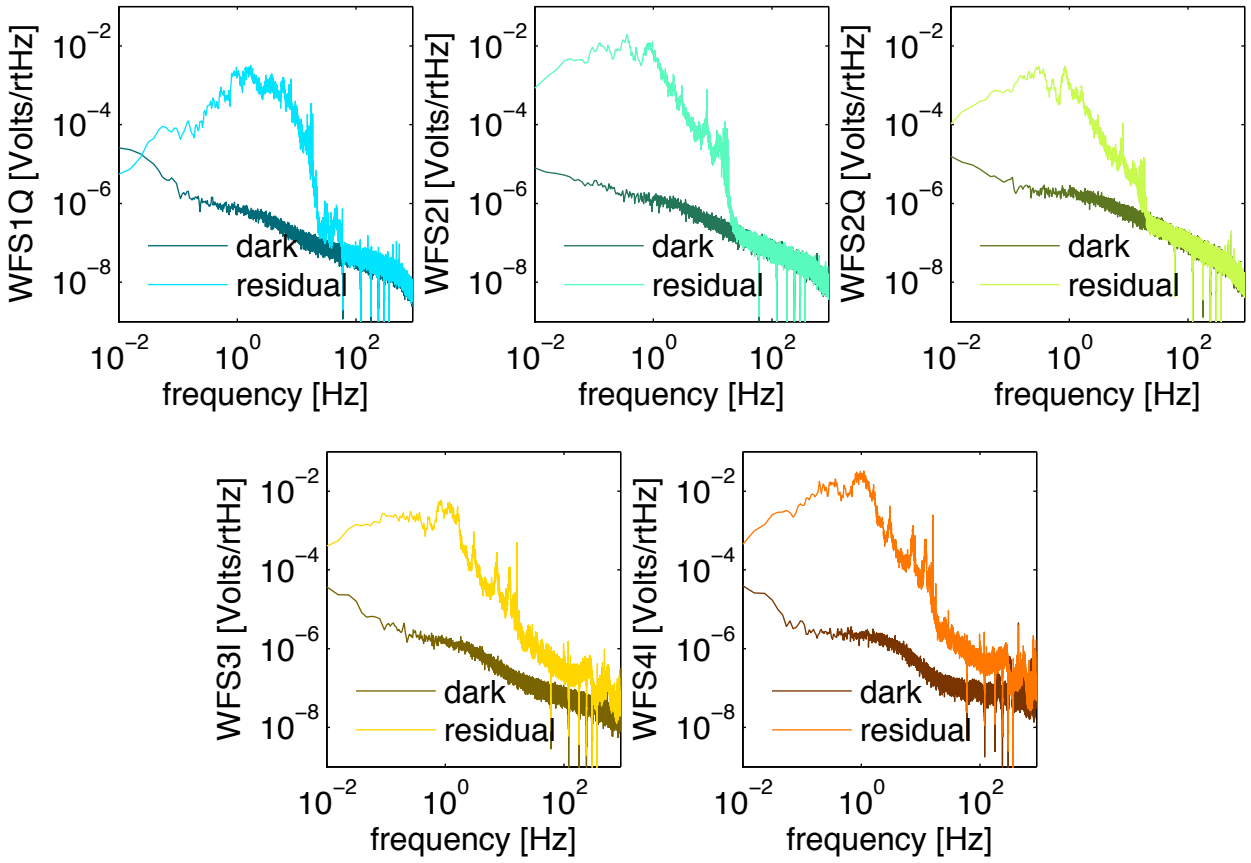

Figure 4-9. WFS dark noise compared to typical error signal. The excess signal above the dark noise in WFS3I and WFS4I above 20 Hz is likely acoustic noise, although this hypothesis has not been verified. WFS1 and WFS2 are on a floating table in a sound proof chamber, while WFS3 and WFS4 are on a non-seismically isolated table without a sound proof enclosure. Seismic noise at time of residual spectra is in Figure C-5.

other out, and all that remains as the plant is the simple pendulum. Including the optical lever compensation in the ASC loop is not necessary, but it is a useful technique for simplifying the design process of the WFS control filters. One need only consider how to control the simple pendulum torque to angle transfer function, rather than the optical-lever controlled pendulum.

4.6 Angular Control Limitations

The limits for how well we can control the angular motion of the interferometer are governed by how well we are able to sense the angular motion. Several of the wavefront sensors' signals are dark-noise-limited above 20 to 25 Hz, as seen in Figure 4-9. Moreover, depending on

the power level, WFS1Q may instead be limited by shot noise (refer to Eq. 2–12). Any control signal derived from frequencies in the sensing noise limited regime will impress the sensor noise on the mirrors. The impression of sensor noise cannot be avoided entirely in the presence of feedback, but can be mitigated by including amongst the control filters a steep cut-off beginning at frequencies where sensor noise dominates.

Besides the sensing noise, there are also sometimes real signals that result in more harm than good when used as feedback. The HAM seismic isolation tables used by the Input Optics (the core optics are suspended from BSC tables) have stack modes at 0.8 to 3 Hz that ring up the MMTs. At low frequencies, around 1 Hz, some of the WFS signals are dominated by these angular fluctuations of the input beam. The resulting attempt of the mirrors to follow the input beam jitter leads to a magnification of the motion because of the drastically different length scales. Large power fluctuations in both arms and the power recycling cavity ensue, leading to departure from the linear error signal regime and often lock loss.

Other limitations to the reduction of mirror motion result from the nature of control loops. The cut-off filter, for example, reduces the phase margin of the open loop gain, necessarily pushing down the unity gain frequency (UGF) and therefore the magnitude of suppression at all frequencies below the UGF. A less aggressive cut-off filter, while improving the servo's stability and allowing for higher overall gains, leads to more impression of sensing noise on the optics. Also, if the phase margin of the loop is low, mirror motion will be amplified through gain peaking.

A final point that should be noted is that there is an angular drive to length coupling due to OSEM imbalances that results from the same principle of the length to angle coupling presented in Section 4.2.2. In all, we will show in Chapter 6 that it is ultimately the WFS noise floor that determines the best possible achievable suppression because through angular control the impression of WFS sensing noise plays a major role in the reduction of DARM sensitivity up to 55 Hz.

CHAPTER 5
THE EFFECT OF HIGH LASER POWER ON INTERFEROMETER ALIGNMENT

The torque induced by radiation pressure, as introduced in Section 4.2.4, couples the angular motion of the suspended mirrors, complicating the plant for which controls must be designed. The derivation of the angular response of the mirrors to external torque in the presence of radiation pressure is presented in several publications, but I provide my own derivation here for completeness. Radiation pressure torque is, after all, the foundation for this work's investigation of high power effects in the angular sensing and control. We derive a set of eigenfunctions that diagonalize the linear cavity's response to radiation pressure and describe the mirrors' equations of motion in this new eigenbasis. We show that the torque to angle transfer functions of the new eigenmodes are modified such that one mode is statically unstable at Enhanced LIGO powers.

5.1 The Radiation Pressure Angular Spring

The geometric axis of a cavity formed by two spherical mirrors is dictated by the line joining the two centers of curvature. Only if the mirrors are pointed directly at one another will the cavity geometric axis pass through the centers of the mirrors. Should a laser beam resonate in the cavity, it will do so along this geometric axis. Thus, if the mirrors are tilted away from one another, the beam spot on each mirror will not be centered. The relationship between the positions of the beams on the mirrors relative to center, x_i, and the angles of the mirrors, θ_i, is given by:

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \frac{L}{1-g_1 g_2} \begin{bmatrix} g_2 & 1 \\ 1 & g_1 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}. \tag{5–1}$$

The g-factor is defined as $g_i = 1 - R_i/L$ where R_i is the mirror radius of curvature, and L is the length of the cavity. The g-factor is 0.73 (0.71) for the LLO (LHO) ITM and 0.54 (0.45) for the LLO (LHO) ETM.

We saw in the previous chapter that the radiation pressure torque on a mirror depends on the position of the beam on the mirror, $\tau_{rp} = 2Px/c$ (Eq. 4–4). Based on Eq. 5–1 the radiation pressure torque on a mirror that is part of a Fabry-Perot cavity is therefore dependent on the angle

of both the mirror of interest and the second mirror forming the cavity:

$$\begin{bmatrix} \tau_{rp,1} \\ \tau_{rp,2} \end{bmatrix} = \frac{2PL}{c(1-g_1g_2)} \begin{bmatrix} g_2 & 1 \\ 1 & g_1 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}. \tag{5–2}$$

This is more succinctly expressed as

$$\vec{\tau}_{rp} = -\mathbf{K_{rp}}\vec{\theta}, \tag{5–3}$$

where $\mathbf{K_{rp}}$ is the *torsional stiffness matrix*. Equation 5–3 is the expression that describes the radiation pressure angular spring.

5.1.1 Diagonalizing the Modified Equations of Motion

The radiation pressure spring modifies the pendulum angular equation of motion and therefore the torque to angle transfer function through the addition of an angle-dependent torque term. Re-writing Eq. 4–6 in matrix form and with the radiation pressure spring term, the two equations that describe the motion of two mirrors forming a Fabry-Perot cavity is:

$$\mathbf{I}\ddot{\vec{\theta}} + \gamma\dot{\vec{\theta}} + \kappa_\mathbf{p}\vec{\theta} - \frac{2PL}{c(1-g_1g_2)} \begin{bmatrix} g_2 & 1 \\ 1 & g_1 \end{bmatrix} \vec{\theta} = \vec{\tau}_{ext}. \tag{5–4}$$

$\mathbf{I}$, γ, and κ_p are 2×2 diagonal matrices and $\vec{\theta}$ and $\vec{\tau}_{ext}$ are 2×1 vectors as in the previous section. Due to the non-diagonal matrix in Eq. 5–4, the motions of each of the mirrors forming the cavity are tied to one another. The natural way to work with such a system is to rotate the coupled equations into a new basis. The resulting de-coupled equations of motion will described specific combinations of mirror tilts instead of the tilt of an individual mirror. Vectors in the rotated basis are written with primes.

In order to decouple the two equations of Eq. 5–4, we need to diagonalize $\mathbf{K_{rp}}$. The subscripts a and b are used to denote the elements of the diagonalized basis, to contrast the 1 and

2 which denote the mirror basis. Ignoring the constants of matrix $\mathbf{K_{rp}}$, its eigenvalues are

$$\lambda_a = \frac{g_1 + g_2 + \sqrt{(g_1 - g_2)^2 + 4}}{2} \tag{5–5}$$

$$\lambda_b = \frac{g_1 + g_2 - \sqrt{(g_1 - g_2)^2 + 4}}{2} \tag{5–6}$$

and its eigenvectors are

$$\vec{v}_a = \begin{bmatrix} 1 \\ \frac{g_1 - g_2 + \sqrt{(g_1 - g_2)^2 + 4}}{2} \end{bmatrix} \tag{5–7}$$

$$\vec{v}_b = \begin{bmatrix} \frac{-g_1 + g_2 - \sqrt{(g_1 - g_2)^2 + 4}}{2} \\ 1 \end{bmatrix}. \tag{5–8}$$

Therefore, the matrix

$$\mathbf{S} = \begin{bmatrix} \vec{v}_a & \vec{v}_b \end{bmatrix} = \begin{bmatrix} 1 & \frac{-g_1 + g_2 - \sqrt{(g_1 - g_2)^2 + 4}}{2} \\ \frac{g_1 - g_2 + \sqrt{(g_1 - g_2)^2 + 4}}{2} & 1 \end{bmatrix} \tag{5–9}$$

diagonalizes $\mathbf{K_{rp}}$ such that

$$\mathbf{S^{-1}K_{rp}S} = \mathbf{D} = \begin{bmatrix} \lambda_a & 0 \\ 0 & \lambda_b \end{bmatrix} = \begin{bmatrix} \frac{g_1 + g_2 + \sqrt{(g_1 - g_2)^2 + 4}}{2} & 0 \\ 0 & \frac{g_1 + g_2 - \sqrt{(g_1 - g_2)^2 + 4}}{2} \end{bmatrix}. \tag{5–10}$$

The matrix of eigenvectors, $\mathbf{S}$, is the basis transformation matrix. It serves to define the torque and angle vectors in the new basis. For example,

$$\vec{\theta}' = \begin{bmatrix} \theta_a \\ \theta_b \end{bmatrix} = \mathbf{S^{-1}} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \mathbf{S^{-1}}\vec{\theta}. \tag{5–11}$$

Rearranging Eq. 5–10 to the form $\mathbf{K_{rp}} = \mathbf{SDS^{-1}}$ and substituting it into Eq. 5–4, we have:

$$\mathbf{I}\ddot{\vec{\theta}} + \gamma\dot{\vec{\theta}} + \kappa_\mathbf{p}\vec{\theta} - \frac{2PL}{c(1 - g_1 g_2)}\mathbf{SDS^{-1}}\vec{\theta} = \vec{\tau}_{ext} \tag{5–12}$$

Table 5-1. Geometric parameters of the LIGO arm cavity eigenmodes. x_i are the beam locations on the mirrors relative to center, a is the cavity axis displacement at the waist, and α is the cavity axis angle with respect to a line joining the centers of the mirrors. Differences between LLO and LHO arise from the mirrors at each site having different radii of curvature. Quantities are expressed as a function of the amount of tilt in a particular mode.

cavity parameter	unit	LLO $\vec{v}_a$ mode	LLO $\vec{v}_b$ mode	LHO $\vec{v}_a$ mode	LHO $\vec{v}_b$ mode
$\lvert x_1\rvert$	mm/urad	9.88	2.44	8.20	2.51
$\lvert x_2\rvert$	mm/urad	10.84	2.22	9.35	2.20
$\lvert a\rvert$	mm/urad	10.17	1.01	8.48	1.34
$\lvert\alpha\rvert$	urad/urad	0.24	1.17	0.29	1.18

Multiplying on the left by $\mathbf{S^{-1}}$, taking advantage of the diagonal $\mathbf{I}$, γ, and κ_p matrices, and using $\mathbf{S^{-1}}$ to change the basis of each of the vectors, the de-coupled equations of motion are:

$$\mathbf{I}\ddot{\vec{\theta}}' + \gamma\dot{\vec{\theta}}' + \kappa_{\mathbf{p}}\vec{\theta}' - \frac{2PL}{c(1-g_1g_2)}\begin{bmatrix}\lambda_a & 0\\ 0 & \lambda_b\end{bmatrix}\vec{\theta}' = \vec{\tau}'_{ext}. \tag{5–13}$$

The radiation pressure torsion constant, κ_{rp}, is

$$\kappa_{rp} = -\frac{2PL}{c(1-g_1g_2)}\lambda \tag{5–14}$$

where $\lambda = \lambda_a$ or λ_b, depending on the mode in question.

The angular motion of the Fabry-Perot cavity is no longer described by the motions of its individual mirrors. Due to radiation pressure, the cavity is treated as a unit and the two orthogonal modes of angular motion are combinations of the two mirrors' angles. The eigenvectors $\vec{v}_a$ and $\vec{v}_b$ describe these two sets of orthogonal mirror tilts, and the eigenvalues λ_a and λ_b (along with their common constants) quantify the magnitude of the radiation pressure torsional spring constant for each of the modes. While the equations of motion had been identical for each of the individual mirrors, the decoupled equations in the presence of radiation pressure breaks that symmetry.

Table 5-1 outlines the characteristics of these two eigenmodes for the specific geometry of the LIGO arm cavities. The amount of beam displacement on each mirror is given as a function of the amount of tilt in one eigenmode or the other. Furthermore, the amount of cavity axis

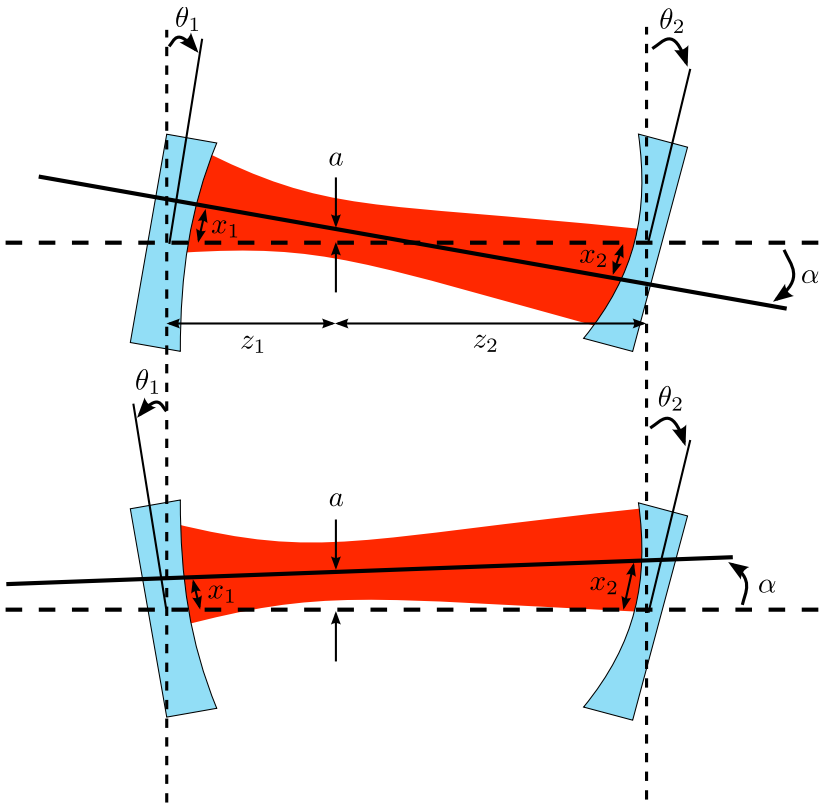

Figure 5-1. Illustration of the orthogonal modes of cavity tilt. The upper diagram shows tilts given by eigenvector $\vec{v}_b$ and the lower diagram shows $\vec{v}_a$.

displacement a and cavity axis tilt α is also calculated for each eigenmode using the geometric relationship between a set of mirror tilts and their cavity axis as derived in Appendix C.2. Figure 5-1 illustrates a cavity in each of the two eigenmodes when using the parameters from Table 5-1.

5.1.2 Soft and Hard Modes

The torque to angle transfer function of each of these eigenmodes has the same form as that of a single pendulum (Eq. 4–8), but the torsion constant is modified. More importantly, the spring constant is modified differently for each mode, yielding distinct behaviors of the two eigenmodes. In this section, we analyze these behaviors and accordingly introduce the names *soft* and *hard* to use in place of a and b for describing the two modes.

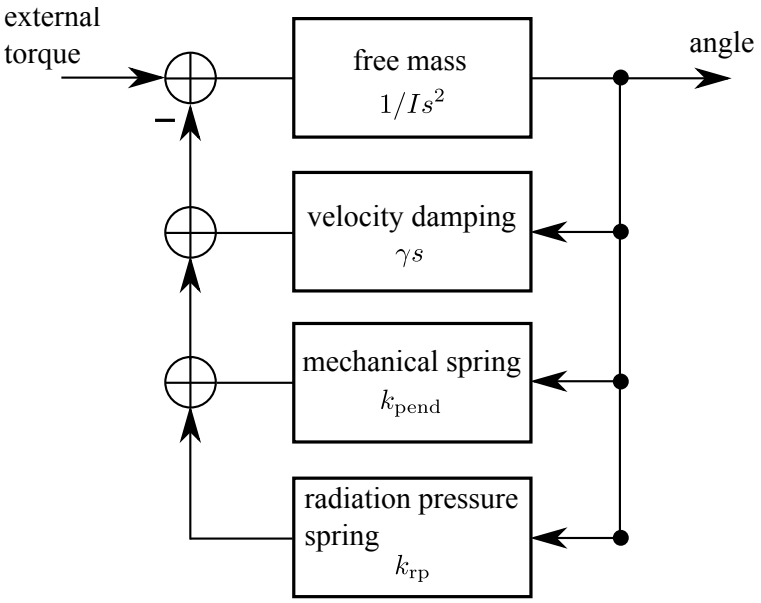

Figure 5-2. Demonstration of how radiation pressure modifies the torque to angle transfer function of a Fabry-Perot cavity's eigenmodes.

Just as in Section 4.3, we can take the Laplace transform of each of the equations in Eq. 5–13 to get the general form of the modal torque to angle transfer function:

$$H'(s) = \frac{\Theta'(s)}{\tau'_{ext}(s)} = \frac{1}{Is^2 + \gamma s + \kappa_p + \kappa_{rp}}. \tag{5–15}$$

Figure 5-2 shows the control theory view of the addition of the radiation pressure spring constant to the transfer function.

The magnitude and sign of the total torsional spring constant, $\kappa_{tot} = \kappa_p + \kappa_{rp}$, conveys critical information about the stability of the cavity and the nature of its response to external torque. Recalling the equation of an angular spring, $\tau = -\kappa_{tot}\theta$, a restoring torque is provided only if $\kappa_{tot} > 0$, which is equivalent to the condition for stability. If $\kappa_{tot} < 0$, the spring is an anti-spring, resulting in an unstable, run-away situation. Furthermore, while κ_{tot} is positive, its magnitude directly relates to the stiffness of the spring.

The stability criteria for the coupled cavity eigenmodes depend on the relationship between κ_p and κ_{rp}:

$$\text{stable: } k_{tot} > 0 \implies \frac{2PL}{c(1-g_1g_2)}\lambda < \kappa_p \tag{5–16}$$

$$\text{unstable: } k_{tot} < 0 \implies \frac{2PL}{c(1-g_1g_2)}\lambda > \kappa_p. \tag{5–17}$$

The pendulum spring constant, κ_p, is always positive, so we can conclude with certainty that the cavity eigenmode is stable as long as the quantity on the left-hand side of Eq. 5–17 is negative. However, if this quantity is positive, then its magnitude compared to κ_p determines stability. Because P, L, and c are all positive numbers and the g-factor is restricted to $0 < g_1g_2 < 1$ [1] , the sign of the left-hand side is determined solely by that of λ. From the g-parameter restriction, it can be shown that λ_a is always positive and that λ_b is always negative. Therefore, the mode whose mirror angles are described by $\vec{v}_a$ is either stable or unstable, and the mode described by $\vec{v}_b$ will always be stable.

The precise situation for the potentially unstable mode depends on the one non-constant variable, the circulating power P. There is a critical power at which κ_{rp} = - κ_p, and at any greater power, instability ensues. In general, as power increases, the total spring constant for the potentially unstable mode decreases, creating a softer spring, and the total spring constant for the unconditionally stable mode increases, creating a stiffer spring. Thus arise the terms *soft* and *hard* to describe the two eigenmodes that have been referred to by $\vec{v}_a$ and $\vec{v}_b$, respectively.

Figure 5-3 shows the dependence of κ_{tot} on circulating power for the soft and hard modes of a LIGO arm cavity. Without power in the cavity, the modes are identical and their spring constants are simply that of the individual pendula. The symmetry-breaking effect of radiation pressure comes into play as soon as light resonates in the cavity: the hard mode's spring constant increases and the soft mode's spring constant decreases. The critical power at which the soft

[1] This is the necessary condition for a two mirror resonator to form a stable periodic focusing system. [64, p. 747]

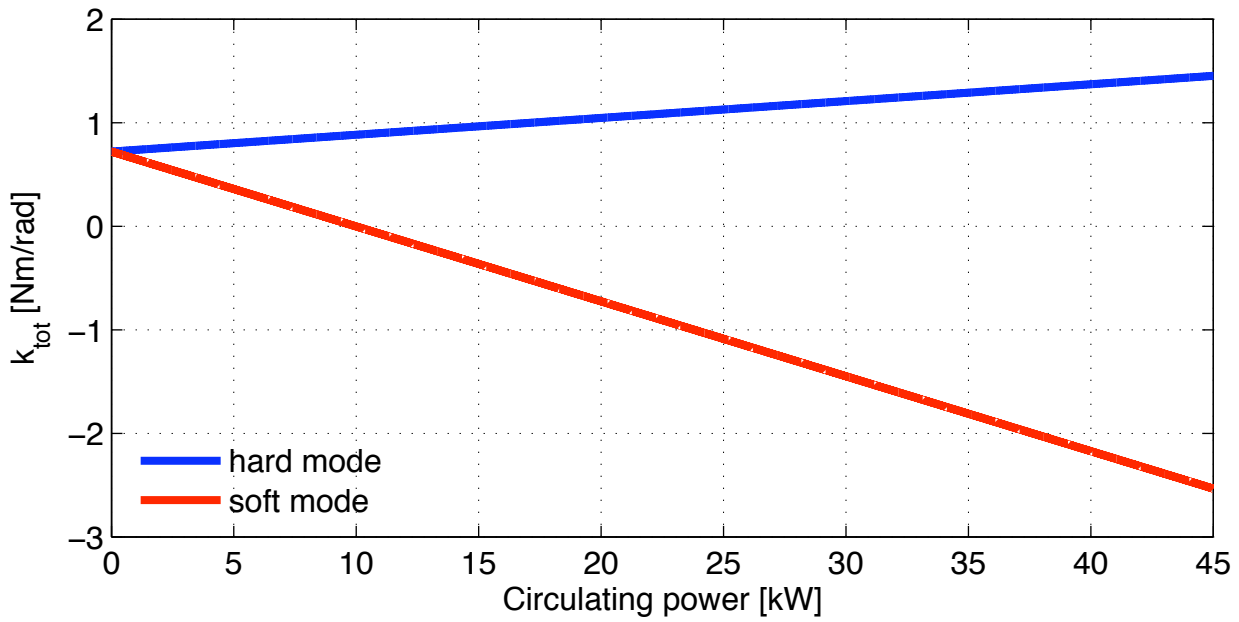

Figure 5-3. Torsional spring constants (pitch) of an optically coupled cavity for LLO parameters. The soft mode is unstable when the spring constant is negative.

Table 5-2. Torsional spring constants (pitch) for the soft and hard modes of a typical Initial LIGO power and the highest of Enhanced LIGO powers. The soft mode in Enhanced LIGO is unstable. The κ_p values assume a resonant frequency of 0.6 Hz.

	P_{circ}	κ_p	κ_{tot}, soft mode	κ_{tot}, hard mode
Initial LIGO	9 kW	0.721 Nm/rad	0.0734 Nm/rad	0.867 Nm/rad
Enhanced LIGO	40 kW	0.721 Nm/rad	-2.18 Nm/rad	1.38 Nm/rad

mode becomes unstable is 10 kW, which corresponds to approximately 6 W input power (for Enhanced LIGO efficiencies) to the interferometer. Above the critical power, radiation pressure creates an optical anti-spring.

Table 5-2 highlights the values of the spring constants for the typical power that was used in Initial LIGO (9 kW) and for the highest of powers achieved in Enhanced LIGO (40 kW). The corresponding transfer functions for these spring constants is found in Figure 5-4. The resonant frequency, $\omega_0 = \sqrt{\kappa_{tot}/I}$, increases with power for the hard modes and decreases for the soft modes. Once κ_{tot} becomes negative, as is the case for the Enhanced LIGO soft mode, there is no real resonant frequency. A summary of the opto-mechanical parameters for LLO and LHO is found in Table 5-3.

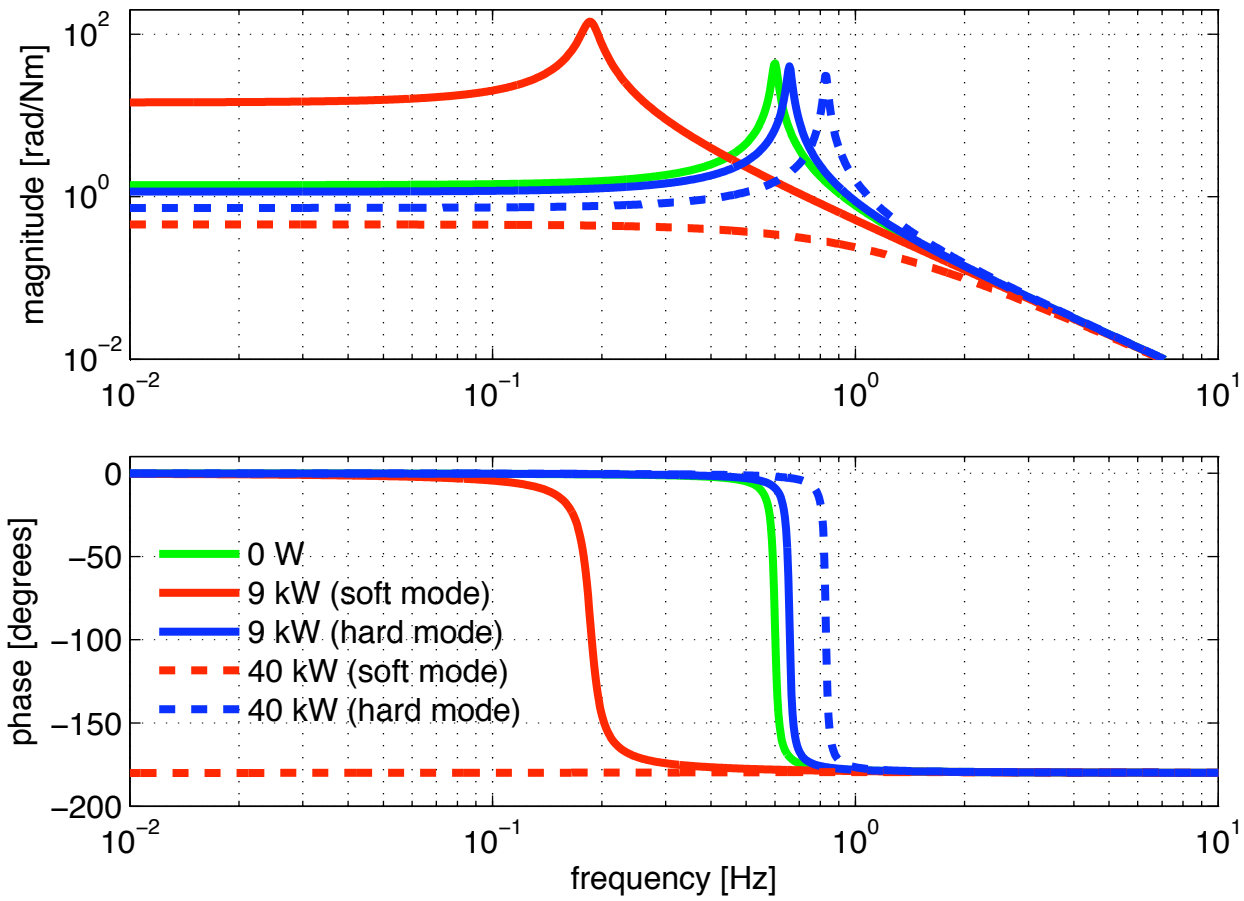

Figure 5-4. Single cavity opto-mechanical transfer function for pitch. The resonant frequency increases with power for the hard mode, but decreases for the soft mode, eventually becoming imaginary. $P_{circ} = 9$ kW (5.25 W input) was a typical operating power for Initial LIGO and $P_{circ} = 40$ kW (23.5 W input) is the highest of powers reached for Enhanced LIGO.

Table 5-3. Opto-mechanical parameters for the LIGO Livingston and LIGO Hanford cavities. Differences result because the mirrors at each site have different radii of curvature.

parameter	label	LLO	LHO
ITM g-factor	g_1	0.73	0.71
ETM g-factor	g_2	0.54	0.45
arm length	L	3995 m	3995 m
test mass moment of inertia	I	0.0507 kg m^2	0.0507 kg m^2
pendulum torsion constant	κ_p	0.72 Nm/rad (pitch)	0.72 Nm/rad (pitch)
		0.50 Nm/rad (yaw)	0.50 Nm/rad (yaw)
soft mode eigenvector	$\vec{v}_a$	(1, 1.10)	(1, 1.14)
hard mode eigenvector	$\vec{v}_b$	(-1.10, 1)	(-1.14, 1)
power when soft mode is unstable	P_{crit}	10.0 kW (pitch)	11.6 kW (pitch)
		7.0 kW (yaw)	8.0 kW (yaw)

Table 5-4. Conditions on total torsional constant κ_{tot} for determining system stability.

κ_{tot} condition	pole s_+	impulse response
$\kappa_{tot} < 0$	real positive	statically unstable
$\kappa_{tot} = 0$	zero	
$0 < \kappa_{tot} < \gamma^2/4I$	real negative	stable decay
$\kappa_{tot} > \gamma^2/4I$	real negative, and imaginary	stable, oscillatory

5.1.3 Pole Analysis

One final comment about the analysis of the modified transfer function (Eq. 5–15) is that the two poles,

$$s = s_\pm = \frac{-\gamma \pm \sqrt{\gamma^2 - 4I\kappa_{tot}}}{2I}, \tag{5–18}$$

provide an alternative way to view the stability of the system. As long as the poles are negative, the impulse response will decay or be sinusoidal. However, if a pole is positive, the system's motion will experience exponential growth. The constraints for $s_\pm$ to be in a particular half of the s-plane are easily derived from Eq. 5–18. Note that s_- will always be in the left half of the plane and that s_+ is the pole that has the potential of falling in the right half of the plane. Table 5-4 show how the s-plane locations for s_+ depend on κ_{tot}. The sign of κ_{tot} determines stability, as expected, and we see that the nature of the stable response depends on the damping coefficient. Figure 5-5 plots the pole locations for a range of κ_{tot} experienced while powering up Enhanced LIGO.

5.2 Implications

The Enhanced LIGO goal of increasing the input power to 30 W from the Initial LIGO 7 W makes radiation pressure torques cross into the realm of significance. In particular, the soft opto-mechanical mode which just approached instability for Initial LIGO powers actually becomes unstable for Enhanced LIGO powers. The transfer functions for which controls must be designed are no longer those of a pendulum with resonance at 0.6 Hz (pitch) or 0.5 Hz (yaw), but those of the soft and hard eigenmodes, whose resonances change with power. For Enhanced LIGO, the angular control servo plant of Figure 4-7 is treated as the cavity's radiation pressure-modified torque to angle transfer function, not as a simple stand-alone pendulum transfer function. An

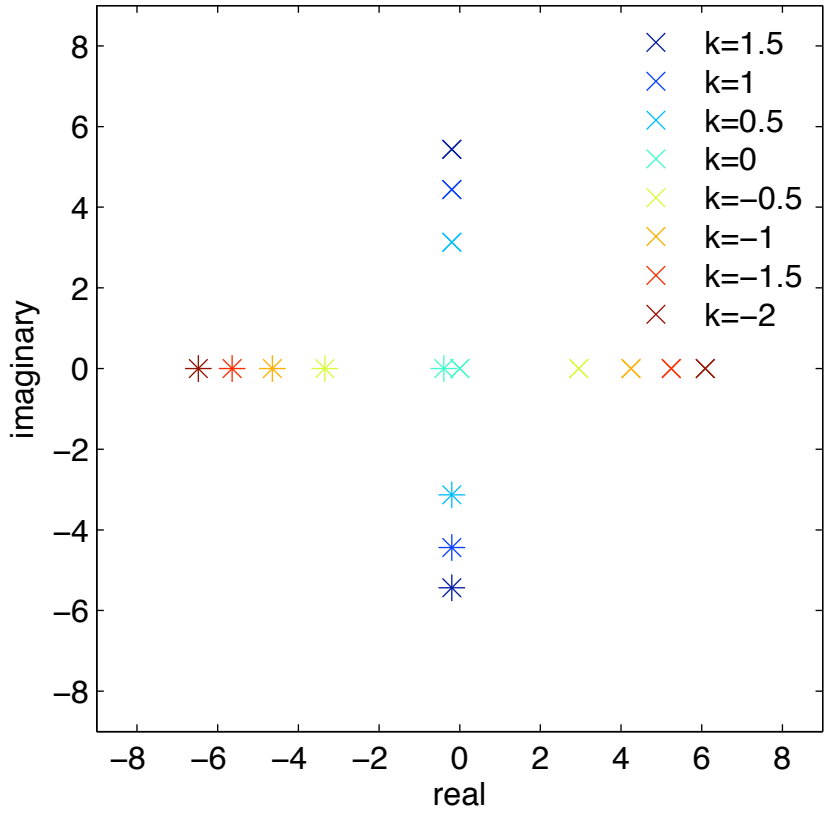

Figure 5-5. Poles of the torque to pitch transfer function as a function of torsional constant, κ_{tot}. Crosses show s_+ and asterisks show s_-. Poles in the right half of the s-plane indicate the system is unstable.

elegant implication of the purely geometric description of the cavity axis is that the radiation pressure eigenmodes are orthogonal independently of power. Therefore, the modes remain independent and the control system need not be updated with power changes.

One complication of having a radiation pressure modified and therefore power dependent plant, P_{rp}, is that the optical lever compensation (described in Section 4.5.3) is no longer valid. The compensation of $1 + OP$ is hard-coded in the digital control system, where the model for P is that of a simple pendulum. The only way to achieve perfect compensation would be to load a new model of P into the compensation filter bank for each new power and for each optic. Doing this in Enhanced LIGO was not practical, so only at low powers does the optical lever

compensation actually cancel out the effect of the optical-lever-controlled radiation pressure-modified plant, $P_{rp}/1+OP_{rp}$, and leave simply P_{rp} for the WFS to control. It this turns out that as power increases, the imperfect compensation actually make the loops more stable [65].

The sensors in Initial LIGO were not tuned to specifically look for the combined mirror motions that create the soft and hard modes. The only way to provide adequate control for both modes would be to increase the gain of all of the angular control loops. Because some of the sensors are not as good as others, this would result in excessive impression of sensor noise on DARM. To minimize impact on strain sensitivity while reducing the angular motion of the interferometer's mirrors to the levels necessary for stable operation, we need to pick out the combination of sensors that together sense specifically the hard mode or the soft mode, and then design controls that specifically address the characteristics of just one mode. This is the foundation of the ASC work for Enhanced LIGO: switching the WFS control to the radiation pressure eigenmode basis, and increasing the gains of only those loops that require it.

CHAPTER 6
ANGULAR SENSING AND CONTROL CHARACTERIZATION AND PERFORMANCE IN THE RADIATION PRESSURE EIGENBASIS

The primary aspect of the Enhanced LIGO ASC upgrade was to switch the control servo from the sensor basis to the natural radiation pressure eigenmode basis, and to keep the contamination to DARM at a minimum. I present in this chapter the design of the new basis and measurements I made to characterize it and its effect on DARM.

A critical aspect of the characterization of any system is to calibrate the data in physical units to facilitate comparison to models and to make meaningful statements. Because the LIGO data is collected digitally, the units are naturally in digital counts. Part of my work was therefore to calibrate each of the relevant ASC channels to physical units. I include the details of the calibrations in Appendix B.

Also, as shown in Section 4.2.1 almost all of the mirror motion is in fact due to the ground. Therefore, the measurements I made of the ASC are very sensitive to the particular state of seismic noise. I include in Appendix C.7 seismic spectra from the time of each measurement I present.

6.1 The ASC Change of Basis

The basis for angular control in Initial LIGO was that of the physical sensors. The wavefront sensors are located such that they sense common and differential ETM and common and differential ITM angular motions. The input matrix was diagonal and the control filters were designed to feed back to those sets of motions via the output matrix. This does not, however, lend itself to easily handling radiation pressure torque. Because the sensor basis is not the radiation pressure eigenbasis of Section 5.1.1, each control servo handled combinations of the soft and hard modes. Due to radiation pressure, each mirror has two resonances, a more complicated plant than that offered by the eigenbasis which has a single resonance for each mode. The change in control basis from Initial LIGO to Enhanced LIGO was a rather straightforward operation of changing only the ASC input matrix and the ASC output matrix as described in the following subsections. Refer to Figure 4-7 for the locations of these matrices in the servo.

Although we changed the meaning of the WFS control signals for Enhanced LIGO, we did not update the convention for their naming as seen in the channel names: WFS1, WFS2A, WFS2B, WFS3, and WFS4. This can be misleading because of the similarity to the sensor names (WFS1Q, WFS2I, WFS2Q, WFS3I, and WFS4I), potentially leading one to assume a one-to-one correspondence. Although a one-to-one correspondence was accurate for Initial LIGO, it is not true for Enhanced LIGO. Therefore, for clarity in this dissertation, I refer to the WFS control signals by the radiation pressure eigenbasis degrees of freedom (DOFs) they represent:

- WFS1 $\rightarrow$ differential soft (dSoft)
- WFS2A $\rightarrow$ common soft (cSoft)
- WFS2B $\rightarrow$ differential hard (dHard)
- WFS3 $\rightarrow$ common hard (cHard)
- WFS4 $\rightarrow$ recycling mirror (RM)

Differential and common refer to the comparison of the sign of the mode in each arm.

6.1.1 WFS Input Matrix

The optical gain for each of the optics' motions is not concentrated at any one particular port, although it can appear more strongly in one location compared to another. In order to make use of as much of the information as possible, we must use all sensors that witness a particular motion. However, when there are multiple signals at a particular place, the WFS error signals tell us the sum of all optical gains at its location. The amount of optical gain at each detector for each motion at one particular frequency forms the *sensing matrix*. The inverse of the sensing matrix is known as the *input matrix*, which tells how to take the appropriate weighted sum of signals in order to reconstruct a particular motion. The procedure for measuring the sensing matrix is as follows:

- excite one of the mirrors (or specific combination of mirrors) at frequency f
- demodulate each of the WFS signals at f
- normalize to the phase of the excitation readback
- repeat for each mirror (or sets of mirrors)

A key aspect of the measurement is that a notch filter at frequency f is engaged so the control servo does not suppress our excitation.

Table 6-1. Angular optical gain at 9.7 Hz in units of Volts per degree of freedom microradian (pitch). Numbers in gray are the measurement results that have coherences less than 0.9. Boxes highlight the elements actually used in the control system. All other elements are set to zero.

WFS1Q	WFS2Q	WFS2I	WFS3I	WFS4I	
2.0	0.03	0.06	-0.008	0.01	dSoft
0.31	-0.03	-0.04	0.002	-0.01	dHard
0.02	-0.01	0.18	-0.02	-0.10	cSoft
0.17	-0.01	-0.21	0.007	-0.12	cHard
0.09	-0.01	-0.21	0.04	-0.21	RM

An example calibrated sensing matrix in the radiation pressure eigenbasis as taken during a 10 W lock is shown in Table 6-1.[1] Rows represent excitation and columns are the wavefront sensors. Before inverting to create the input matrix, the smallest of the elements (which are more or less equivalent to the elements for which optical gain is also expected to be weak), are artificially set to zero. This avoids the contamination of strong signals by those with weak signal-to-noise ratios. The elements that remain after this process are highlighted by boxes. Note that the sensing matrix is in fact composed of two sub-matrices: one for the differential degrees of freedom, and one for the common degrees of freedom. Also, WFS1Q has particularly strong signal compared to the other wavefront sensors. We see in Section 6.5 that this allows us to provide much more control to the differential soft mode compared to the other modes.

6.1.2 WFS Output Matrix

The WFS output matrix determines how to convert the radiation pressure eigenbasis control signals into individual mirror control signals. It is the basis transformation matrix, S, as defined in Eq. 5–9. The matrix is arbitrarily normalized so the largest element is 1, and it is repeated with appropriate sign changes to form differential and common soft and hard modes of the two arms. The output matrix is shown in Table 6-2, where the r is 0.91 for Livingston and 0.87 for Hanford.

[1] Refer to Appendix B.3 and B.4 for a description of how the calibration is done.

Table 6-2. WFS output matrix (pitch). For the Livingston cavity geometry $r = 0.91$ and for Hanford $r = 0.87$.

dSoft	dHard	cSoft	cHard	RM	
1	r	1	r	0	ETMX
-1	-r	1	r	0	ETMY
r	-1	r	-1	0	ITMX
-r	1	r	-1	0	ITMY
0	0	0	0	1	RM

6.1.3 Diagonalizing the WFS Drive Matrix

The mirror gain matrix (introduced in Fig. 4-7) is a diagonal matrix that modifies the amplitude of the control signal to each optic by some factor close to unity. It corrects for experimentally measured deviations from the theoretical output matrix. The purpose of the mirror gain matrix, **M**, is to ensure that when a particular WFS is driven, only its DOF is excited. We call the matrix of observed DOF motions per WFS excitation the drive matrix, **D**.

The mirror gain matrix is what we tune to make the drive matrix diagonal. The drive matrix can be represented as:

$$\mathbf{D} = \mathbf{UMC}. \tag{6–1}$$

A WFS control signal is multiplied first by the output (control) matrix, **C**, then by the mirror gains, and then put into hardware with an unknown transfer function, **U**, to create a physical torque on the mirrors. To make **D** diagonal, we start with $\mathbf{M} = \mathbf{1}$, experimentally construct the drive matrix, and calculate the new mirror gain matrix, $\mathbf{M}_{\text{new}}$, which must have the property:

$$\mathbf{M}_{\text{new}} = \mathbf{U}^{-1}\mathbf{C}^{-1} = \mathbf{MCD}^{-1}\mathbf{C}^{-1}. \tag{6–2}$$

Note that the set of unknown transfer functions **U** is eliminated.

I measured **D** by recording the demodulated optical lever error signals during the sensing matrix measurement (refer to Sec. 6.1.1). Because the optical levers provide a record of the motion of the test masses, combining their calibrated responses via the output matrix (Table 6-2)

Table 6-3. Mirror gains for diagonalization of drive matrix.

ETMX	ETMY	ITMX	ITMY	RM
1.33	1.38	0.96	0.87	1.0

Table 6-4. Actual eigenbasis motion during sensing matrix excitations as witnessed by the optical levers. Columns are excitations and rows are the pitch motions at 9.7 Hz in units of µrad. The mirror gains are selected to make this matrix as diagonal as possible.

dSoft	dHard	cSoft	cHard	RM	
5.1e-06	-5.2e-08	6.1e-07	-3.8e-08	-1.8e-07	dSoft
-3.4e-07	**5.0e-06**	7.3e-07	-1.0e-06	2.4e-07	dHard
-4.1e-07	-3.3e-08	**5.9e-06**	6.8e-07	2.5e-07	cSoft
-6.4e-07	-5.9e-07	1.1e-06	**5.7e-06**	4.7e-07	cHard
-1.6e-07	-1.8e-06	-5.5e-07	2.6e-06	**5.6e-06**	RM

allows one to convert the individual mirror motions during each DOF excitation to motions in the radiation pressure DOF basis. [2]

The mirror gain matrix is shown in Table 6-3. The 30% difference with the model is a result of uncertainty in the cavity geometry. The resulting drive matrix, corresponding to the same measurement time as the sensing matrix of Table 6-1, is presented in Table 6-4. Columns are excitations and rows are the pitch motions at 9.7 Hz in units of µrad. The absolute amplitude of DOF motions for each excitation is not of significance; only the amplitudes in each column should be compared. Ideally, this drive matrix should be diagonal and it is up to at most a factor of two.

6.2 Sensing Matrix Stability

As the interferometer conditions change, so does the sensing matrix. The inverse of the sensing matrix, the input matrix, is hardcoded in the digital control servo and not actively updated. Therefore, it can be expected that the ASC performance may not be stable.

By design, the input matrix is not exactly the inverse of the sensing matrix, meaning the system is not completely diagonal. For example, the input matrix times the sensing matrix is, by

[2] The optical lever calibration procedure and measurement results are found in Appendix B.2.

design:

$$\begin{bmatrix} 0.79 & 0 & 0.30 & 0 & 0 \\ 0 & 1.0 & 0 & 0 & 0 \\ 0 & 0 & 1.0 & 0 & 0 \\ 0 & 0 & 0 & 1.0 & 0 \\ 0 & 0 & 0 & 0 & 1.0 \end{bmatrix} \tag{6–3}$$

Over time, the sensing matrix changes enough that the system is even less diagonal. Only 10 minutes after having measured and created the input matrix, the product of the input matrix with a newly measured new sensing matrix is:

$$\begin{bmatrix} 0.79 & 0 & 0.32 & 0 & 0 \\ 0 & 1.0 & 0 & 0 & 0 \\ -0.02 & 0 & 1.0 & 0 & 0 \\ 0 & -0.02 & 0 & 1.0 & 0.02 \\ 0 & 0.02 & 0 & 0.05 & 1.03 \end{bmatrix} \tag{6–4}$$

After one week, it is:

$$\begin{bmatrix} 0.67 & 0 & 0.27 & 0 & 0 \\ 0 & 0.91 & 0 & 0.03 & -0.03 \\ 0.09 & 0 & 0.84 & 0 & 0 \\ 0 & 0.15 & 0 & 0.83 & 0.27 \\ 0 & 0.06 & 0 & 0.07 & 1.04 \end{bmatrix} \tag{6–5}$$

And after three weeks, it is:

$$\begin{bmatrix} 1.4 & 0 & 0.55 & 0 & 0 \\ 0 & 1.8 & 0 & -0.09 & 0.15 \\ -0.08 & 0 & 2.0 & 0 & 0 \\ 0 & -0.22 & 0 & 1.4 & 0.64 \\ 0 & 0.30 & 0 & -0.72 & 2.9 \end{bmatrix} \tag{6–6}$$

Despite these significant changes, the interferometer remained stable and the sensitivity remained constant. This shows that the ASC is a very robust sensing and control system.

6.3 Input Beam Motion

The beam centering and QPD servos operate up to only about 50 mHz, meaning the beam-centering degree of freedom is uncontrolled at higher frequencies. Because beam spot motion on the test masses couples to DARM, anything that causes the beam's position on the test mass to change on time scales faster than half a minute becomes itself a direct noise source for DARM. The HAM seismic isolation tables from which the input optics are suspended have resonant "stack" modes from about 0.8 Hz to 3 Hz. The excess table motion at these frequencies is transmitted to the MMTs. Jitter on the pointing of the input beam is thus a primary contender for beam spot motion on the test masses.

The wavefront sensor servos are the mechanism by which input beam motion is impressed on the test masses; they are responsible, among other things, for making the interferometer follow the input beam up to several Hz. The WFS detect differences between the angle of the cavity (as determined by the angles of the mirrors) and the angle of the beam impinging the cavity. If either the input beam or the cavity angle changes, the WFS will move the mirrors to correct for the angle mismatch. Thus, even if the mirrors are perfectly quiet, a non-stationary input beam will result in mirror motion and mirror motion in turn creates beam spot motion as seen from Eq. 5–1.

I measured the impression of the input beam motion on the mirrors by increasing the gain of the common-degree-of-freedom WFS servos (cHard, cSoft, RM) for about 10 minutes. Comparing the amount of angular motion of the mirrors as witnessed by the optical levers from this time of high common WFS gain to a time with nominal WFS gain and similar seismic motion, we can see the effect directly. Figure 6-1 shows comparison spectra, demonstrating how there is higher test mass motion around 1 Hz when the common WFS gains are higher. The rms mirror motion also increases by about 20%.

It is possible for the extra mirror motion to result from gain peaking of the WFS servos. However, a plot of the WFS error signals during the time of nominal gain and high gain shows

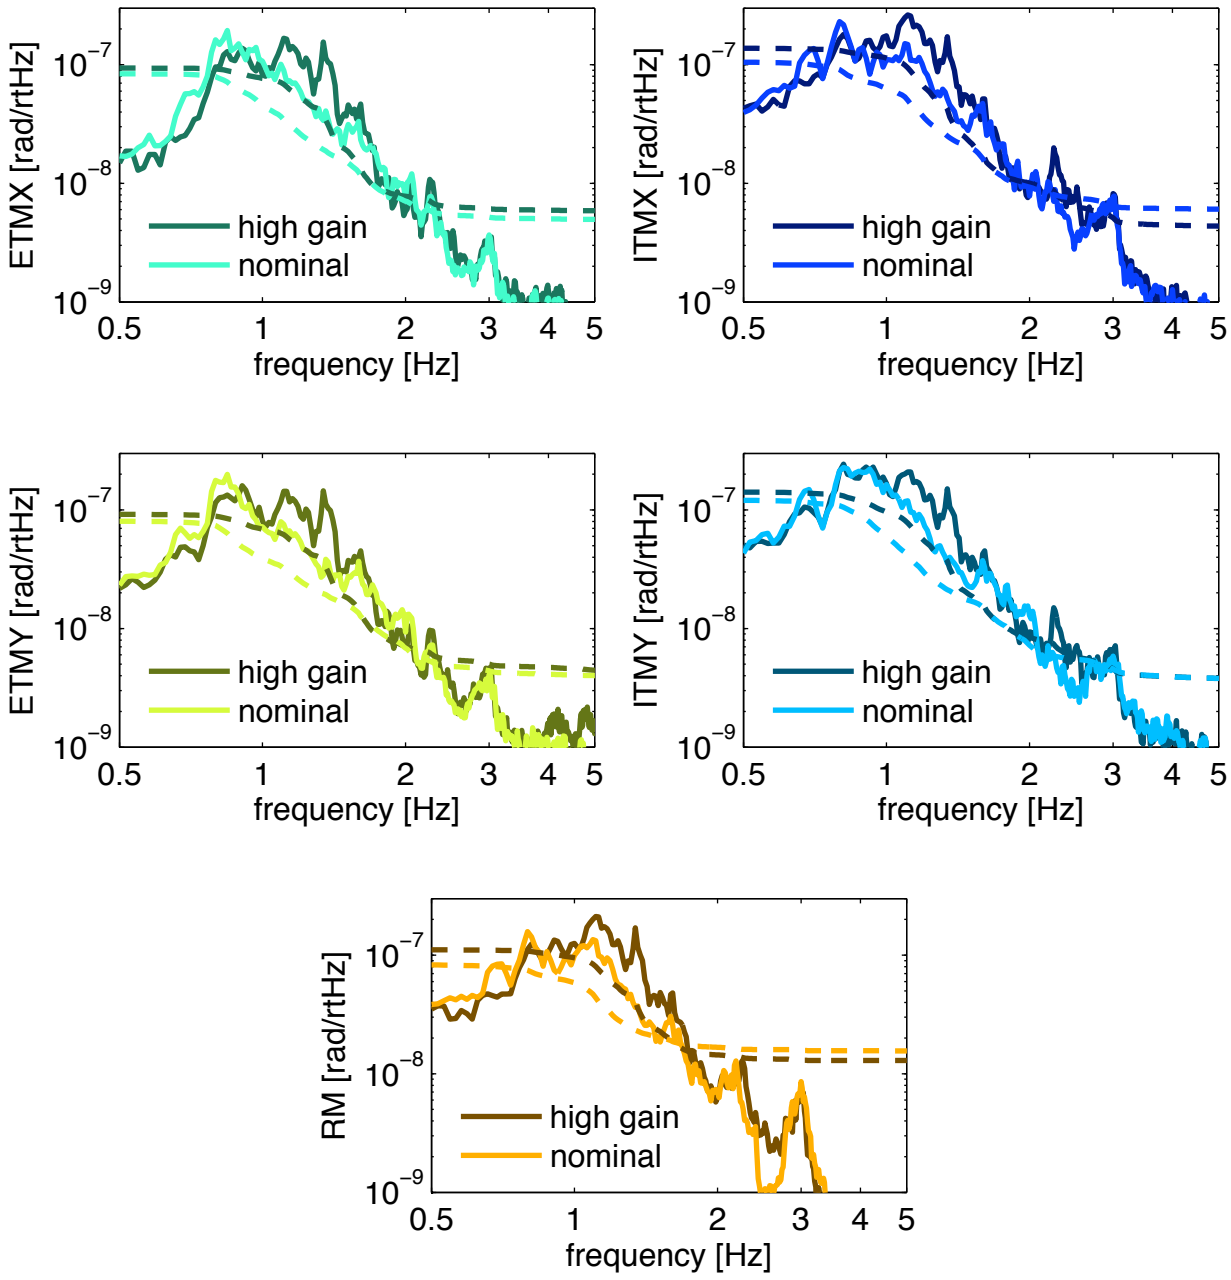

Figure 6-1. Input beam motion impression on the core mirrors (pitch). Residual mirror motion as witnessed by the optical levers when the common WFS gains (cSoft, cHard, RM) are increased to 2.5× nominal is compared to residual mirror motion when the WFS gains are nominal. Dashed lines are the root-mean-square of the amplitude spectral density integrated from the right. Both spectra come from a time of similar seismic activity (typical weekday afternoon noise), shown in Figure C-6.

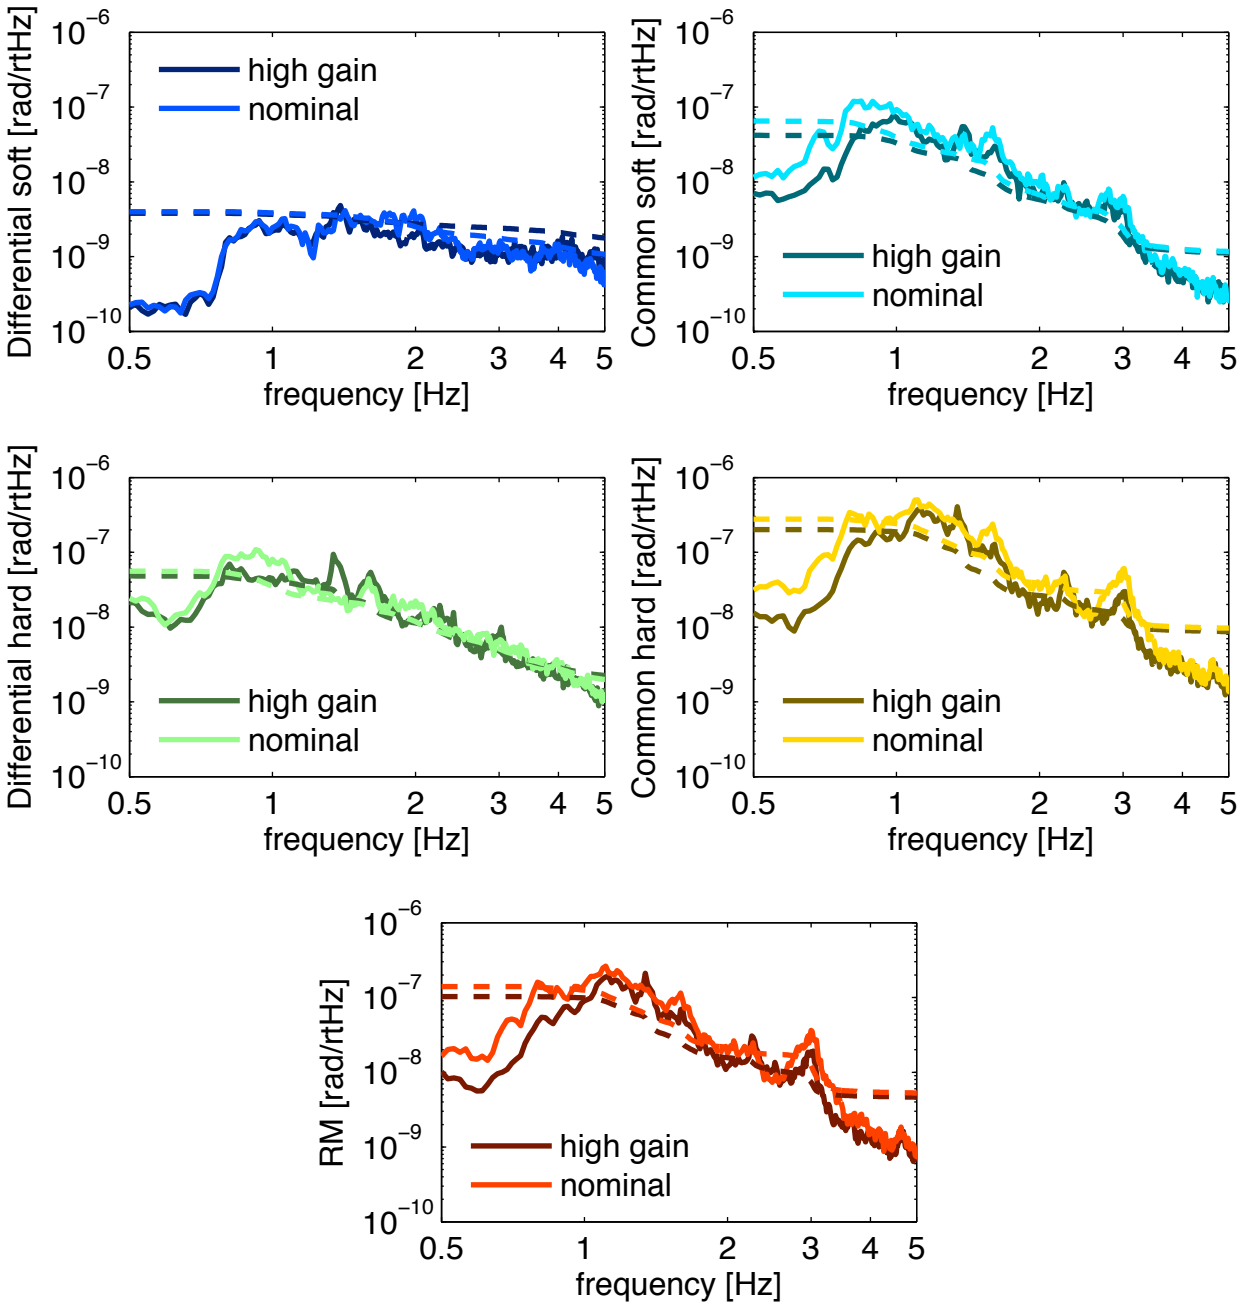

Figure 6-2. Comparison of WFS error signals (the residual motion) during a time of normal operation and a time when the common WFS gains were 2.5× higher than nominal. This excludes gain peaking as a cause of the extra mirror motion witnessed during the time of high gain (Figure 6-1). Dashed lines are the root-mean-square of the amplitude spectral density integrated from the right. Figure C-6 shows the ground motion spectra at the time of this measurement.

that there is not, in fact, any evidence of gain peaking. The common WFS spectra in Figure 6-2 do not show any extra noise when their gains are higher. It is worth noting that the higher gain is evident in the common WFS spectra by the extra suppression seen below 1 Hz. Also, the differential WFS spectra are unchanged, as expected. It can therefore be concluded with reasonable certainty that the increase in test mass motion between 1 and 2 Hz during this test is indeed due to the WFS impressing input beam motion on the mirrors.

6.4 The Marginally-stable Power Recycling Cavity

The power recycling cavity (PRC) is the linear cavity formed by the RM and ITMs. Because the radius of curvature of both the RM and the ITMs points in the same direction and the waist is well outside the Rayleigh range of the mirrors, the cavity is geometrically unstable. For example, in its cold state at LLO the g-factor of the cavity is 1.00005 and at LHO it's 1.00003. The beam in the PRC is not spatially contained and the cavity is degenerate with respect to higher order modes. The heating of the ITMs from the kilowatts of power in the arm cavities together with the ITM thermal compensation system (TCS) serve the role of making the PRC geometrically stable for interferometer operation. The heating and cooling of the ITMs is a very complicated process and therefore not very precise, so the value of the hot PRC's g-factor is usually not constant.

The changing g-factor has potentially severe consequences for the ASC. Because of its geometry, the power build-up in the PRC is very sensitive to both the mirror angles and the g-factor. Power fluctuation is detrimental because the signal to noise ratios of the sensors that probe the PRC light degrade due to the presence of increased junk light that contributes shot noise but not signal. WFS1Q, WFS2I, and WFS2Q are the most sensitive to the PRC because their signals are derived from the 25 MHz sidebands. Their sensitivity to mirror motion is therefore subject to change. Because achieving a flat power build-up in the PRC is a difficult task (too much motion in the PRC is quite often a cause of lock loss when making measurements), we must update the real-time control system to reflect their changing sensitivities. Otherwise, the mirror angles will not be controlled accurately.

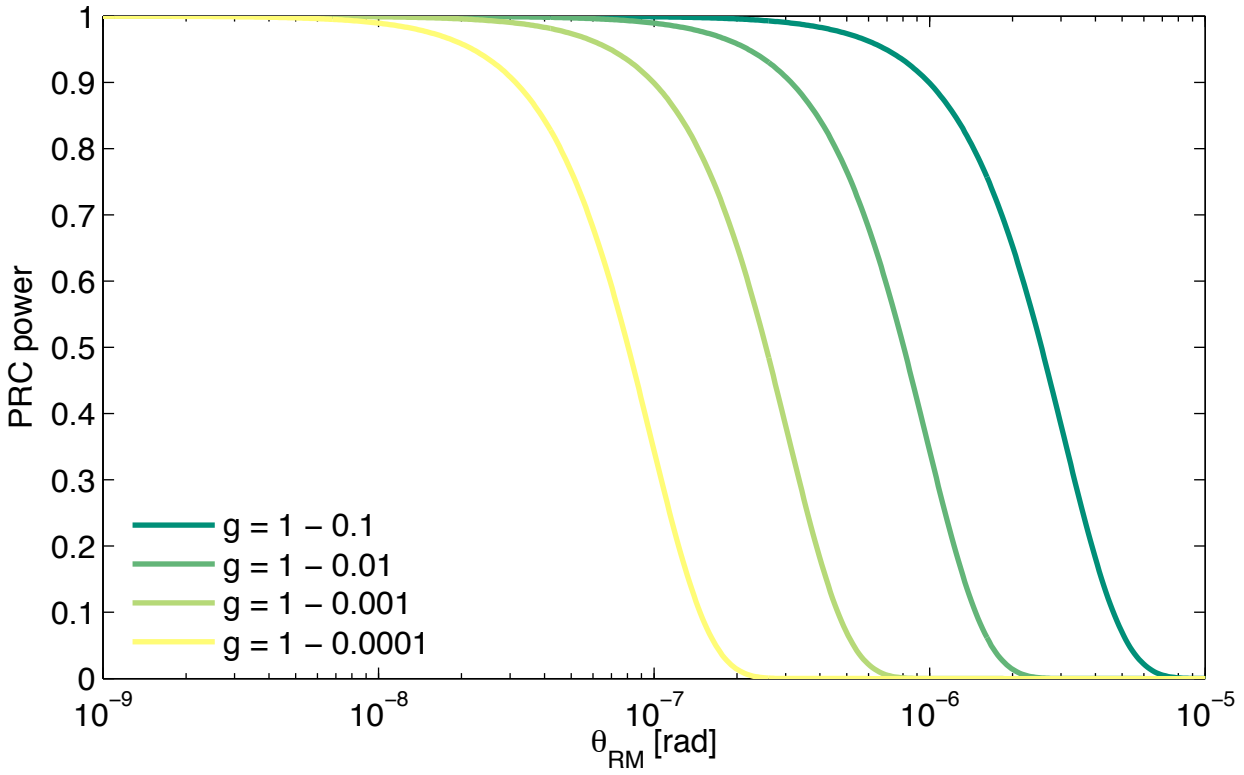

Figure 6-3. Dependence of power build-up in the power recycling cavity on the PRC's g-factor and the RM tilt. TCS is necessary for stabilizing the PRC's geometry and therefore its sensitivity to mirror motion. For simplicity, the ITM is assumed stationary in these plots.

An estimate of the expected power fluctuations based on the g-factor and RM motion is a straightforward exercise when using Eq. C–4 and Eq. C–31 as derived in Appendix C. If we estimate the g-factors of the RM and ITM as $g_{RM} = 1+\delta$ and $g_{ITM} = 1-\delta$ (where $\delta = 6\times10^{-4}$ for LLO the cold state) and approximate the distance of each mirror to the cavity waist as z because the two mirrors are very close to each other compared to the waist location, then Eq. C–4 reduces to:

$$\begin{bmatrix} a_{PRC} \\ \alpha_{PRC} \end{bmatrix} = \begin{bmatrix} z(2+\delta)/\delta & z(2-\delta)/\delta \\ -1/\delta & -1/\delta \end{bmatrix} \begin{bmatrix} \theta_{RM} \\ \theta_{ITM} \end{bmatrix}. \tag{6–7}$$

Figure 6-3 plots the power in the PRC as a function of θ_{RM} for several values of δ, demonstrating the sensitivity of the PRC to the ITM heating. For example, the typical RM angular

displacement of 10^{-7} rad results in a 66% power loss when the PRC g-factor is very near instability with a value of $1 - 0.0001$. Only as the g-factor moves further from 1 does the angular motion of the RM have less and less of an effect on the power build-up.

6.4.1 Power Scaling

The signal at the wavefront sensors is proportional to the amplitude of the sidebands, or the square root of the sideband power. Thus, as the PRC g-factor and therefore the power in the recycling cavity changes, so do the WFS1 and WFS2 optical gains. In order to compensate for this g-factor dependence, we multiply the WFS$\{$1Q, 2I, 2Q$\}$ error signals in real-time by

$$\frac{1}{P_{in}}\left[\frac{\text{NSPOB}}{350}\right]^{-1/2} \tag{6–8}$$

and WFS3I and WFS4I by $1/P_{in}$. NSPOB is the normalized sideband power in the PRC as measured by the $2f$-demodulated POB signal, and the 350 is the reference NSPOB, treated as nominal. Thus, during interferometer operation, all WFS signals are normalized to input power and are not dependent on the PRC power. This correction to the WFS signals is called power scaling.

To verify that the WFS optical gains do indeed scale with the sideband power as expected, I tracked the WFS optical gain as g changes. I excited three of the test masses (ETMX, ITMX, RM) at three different frequencies (9.7 Hz, 10.7 Hz, and 11.7 Hz, respectively) during a full interferometer lock and changed the TCS settings so that over the course of 15 minutes the g-factor steadily changed. Demodulating each of the WFS signals at each of the three excitation frequencies as a function of time shows how the strength of the signal at the WFS due to the motion of these three mirrors changes. To compensate for the difference in pendulum responses to the excitations, I multiplied the demodulated signals for a particular excitation f by $(f/9.7)^2$. I also normalized the response by the phase of the mirror's motion as witnessed by the optical levers.

The results are shown in Figure 6-4, and include a plot of how NSPOB changed over time. As expected, WFS1Q, WFS2I, and WFS2Q show dependence on the PRC power, and therefore

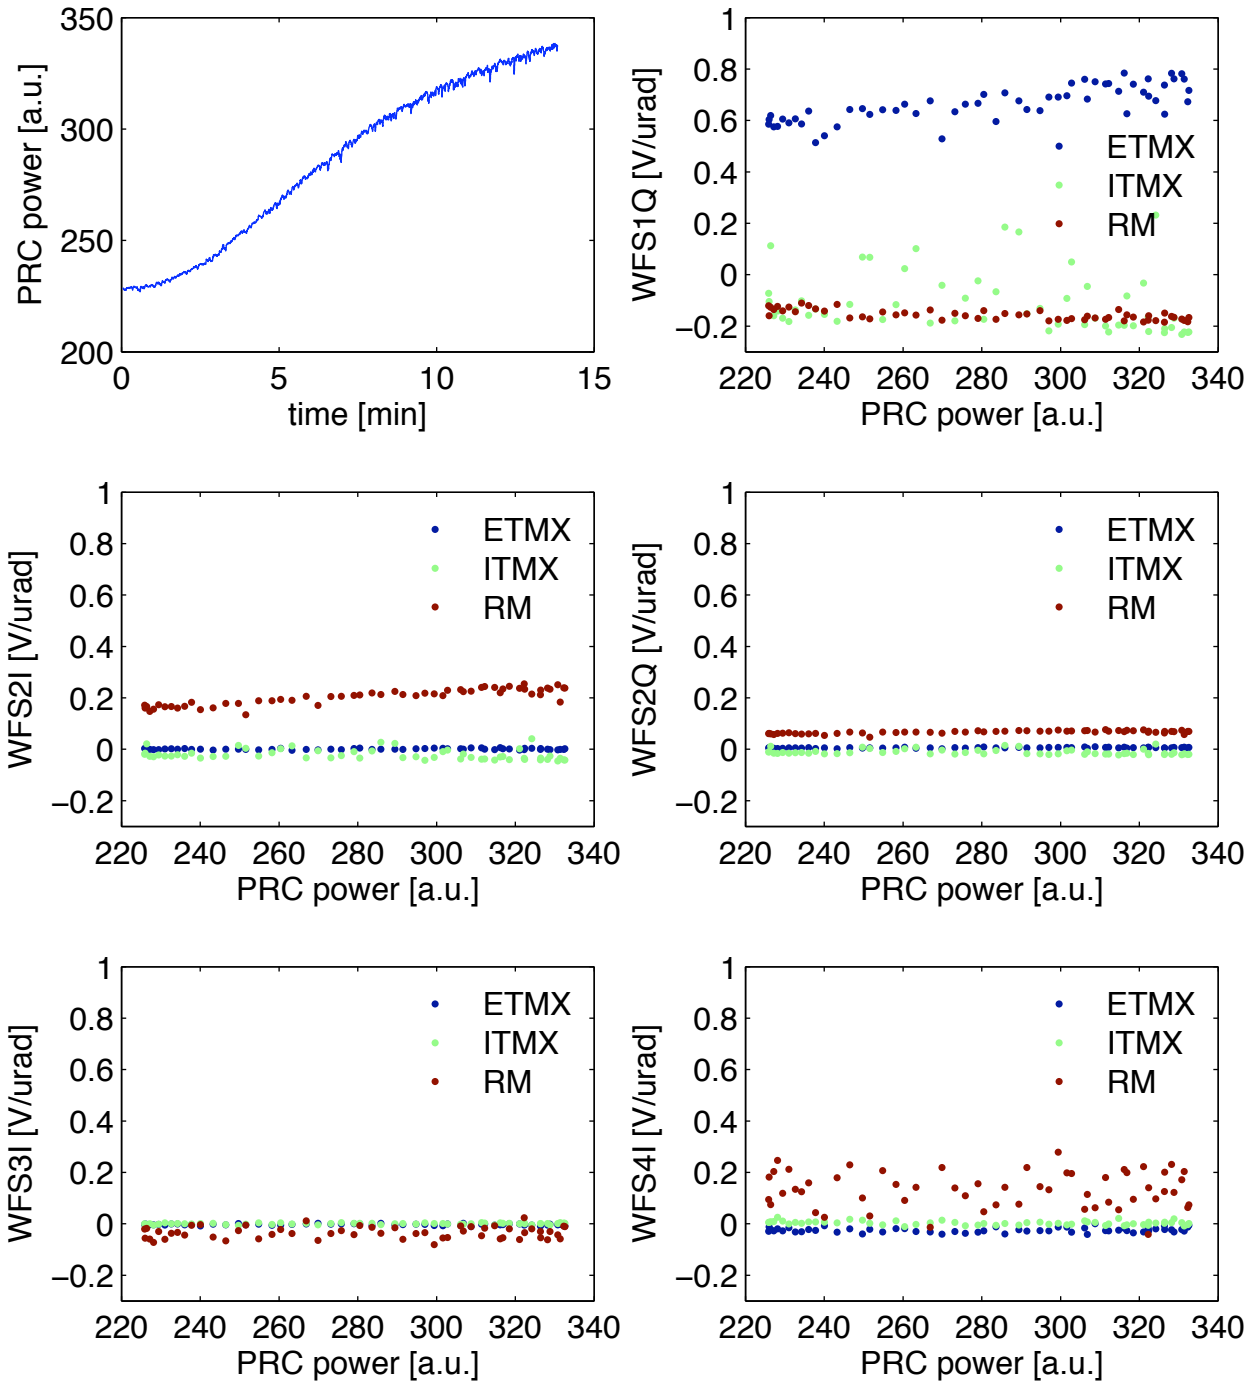

Figure 6-4. WFS optical response to test mass motion as a function of power recycling cavity geometry. WFS1Q, WFS2I, WFS2Q are more sensitive to test mass motion as the power in the recycling cavity increases. Therefore, to achieve a dependable feedback system, we scale the error signals in real-time, forcing their responses to be flat with power. This range of PRC power is low for normal operations.

the g-factor. The WFS3 and WFS4 sensing elements are flat. Fitting lines to each of the tracked elements, we find a good fit with the expected power of $1/2$ dependence.

6.4.2 Sideband Imbalance

Another important effect of the PRC on the ASC signals is the balancing of the upper and lower 24.4 MHz sideband amplitudes. SPOB is the product of the amplitude of the upper and lower sidebands, but the total sideband power is the sum of the power in the lower sideband and the power in the upper sideband. Therefore, if the upper and lower sidebands are not the same, SPOB does not accurately represent the power in the cavity. This creates inaccuracies in the WFS power scaling.

We set up a temporary optical spectrum analyzer at a pickoff of the anti-symmetric port beam in order to measure amplitudes of the sidebands. Without any TCS, we saw that with less than 6 W input power, the lower sideband is smaller than the upper sideband, at 6 W the amplitudes are equal, and above 6 W, the upper sideband is smaller than the lower. Thus, if TCS is not tuned perfectly at all times, we can expect unequal sidebands.

6.5 WFS Servo Open Loop Transfer Functions

The open loop transfer function, or open loop gain (OLG), is an informative measure of the characteristics of a control servo. It is the product of each of the elements of the loop, and is often summarized as being the product of the plant, H, with the control filters, G:

$$\mathrm{OLG} = HG. \tag{6–9}$$

In our case, the plant is the radiation-pressure-modified pendulum transfer function (Eq. 5–15) and the control filters (found in Appendix C.6) are the digital filters between the input and output matrices. The amplitude of the OLG tells the suppression the control loop provides and the phase tells about the stability of the system.

We measure the open loop transfer function while the loop is closed. This is done by injecting a swept-sine excitation just before the digital filters and taking the ratio of the signals just before and just after the excitation. Figure 6-5 shows the open loop transfer functions of

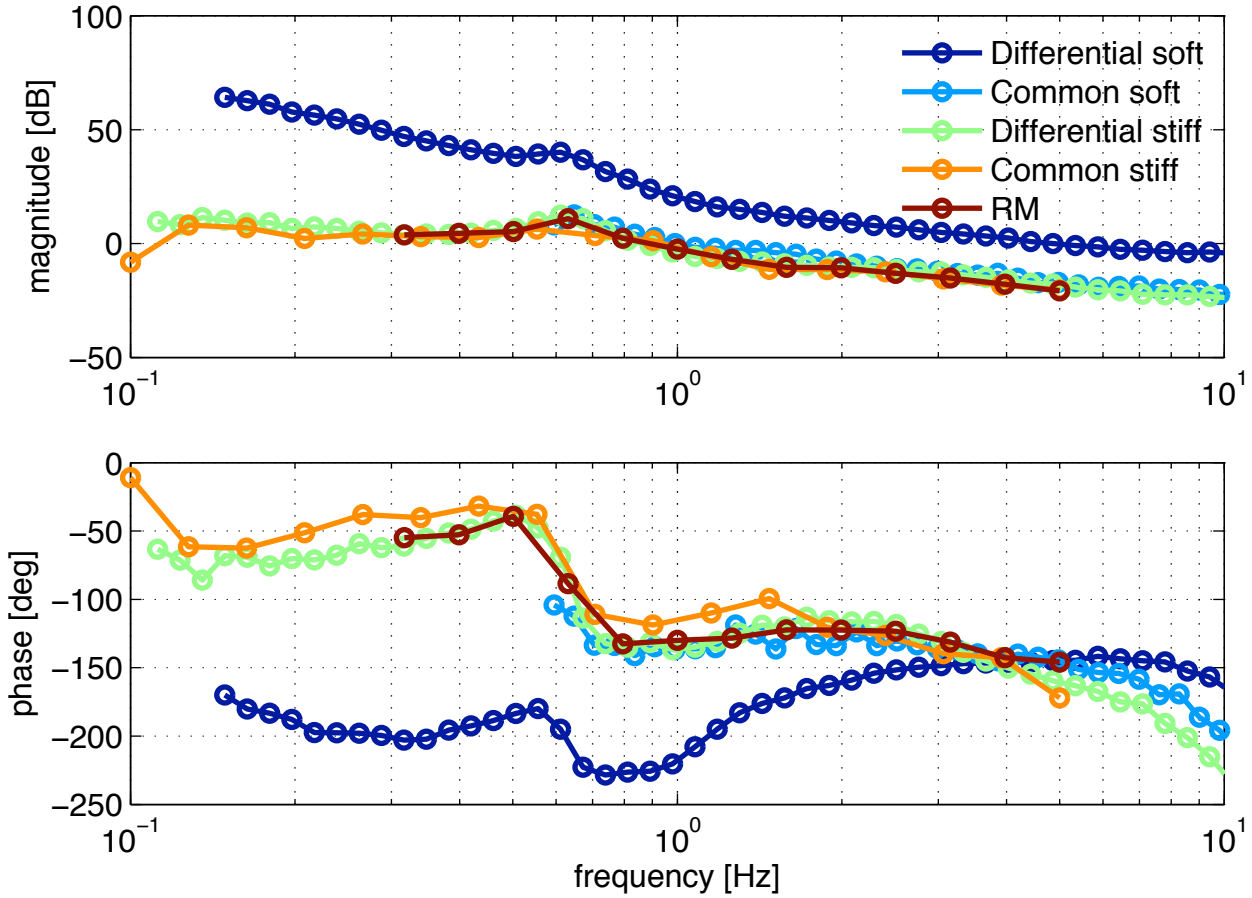

Figure 6-5. Open loop gains (pitch) of the 5 WFS loops as measured with 6 W input power.

each of the wavefront sensor loops as measured during a 6 W lock. As anticipated from the large differential soft signal seen by WFS1 in the sensing matrix measurement (Table 6-1), that is the mode for which we can and do provide the strongest suppression. However, it is also conditionally stable (due to a pole at zero that is engaged after the servo is turned on), as seen by the phase dropping below -180° at two different frequencies. The unity gain frequency (UGF) must be held between roughly 2 Hz and 10 Hz. As shown here, the dSoft UGF is at 5 Hz. For the other degrees of freedom, the UGFs are all between 0.8 and 1 Hz. Each loop has a phase margin of 30° to 40°.

Figure 6-6 shows measurements of how the dSoft (WFS1) OLG changes with power. The 6 W measurement data is the same as that presented in Figure 6-5. From 1 W to 14 W input power the suppression provided by the control loop remains constant, but the stability of the loop

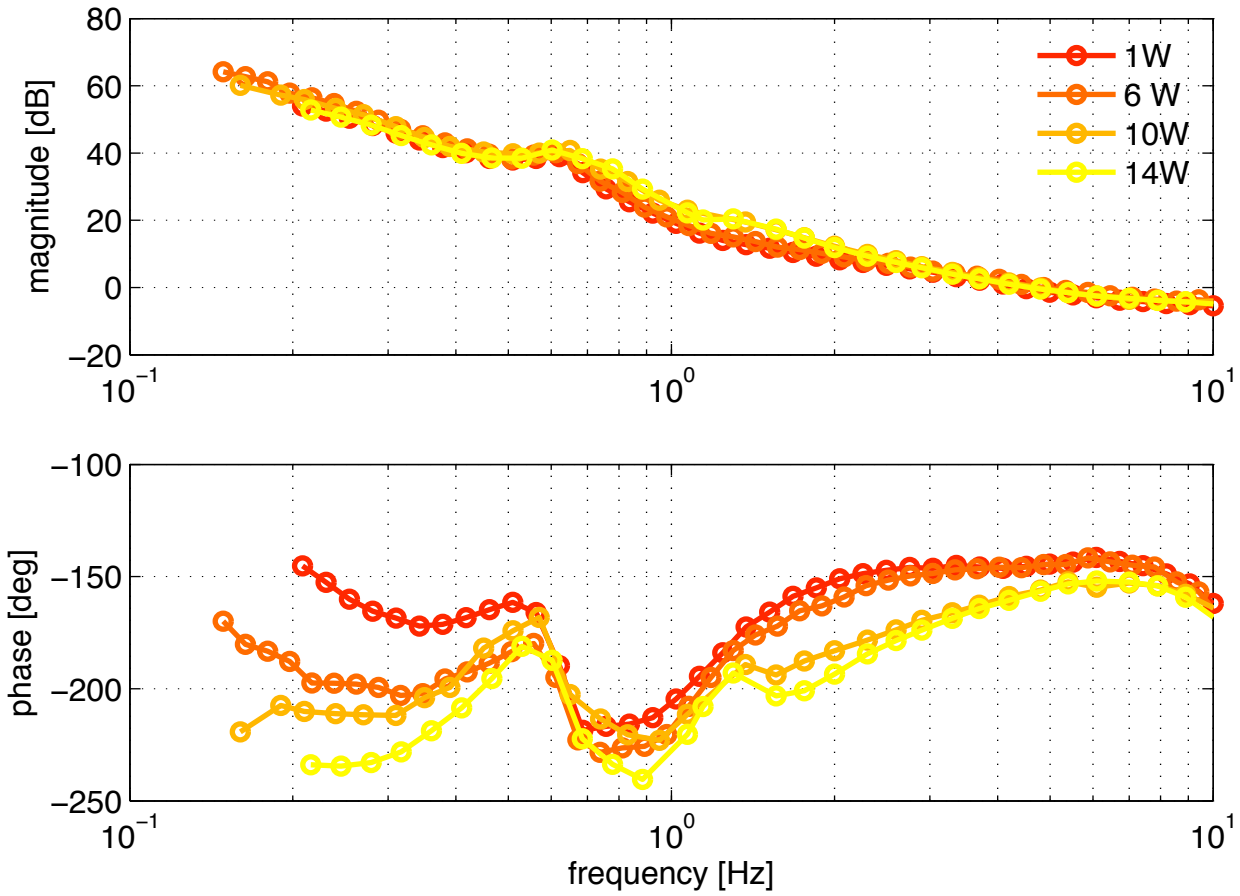

Figure 6-6. Open loop gains (pitch) of the differential soft (WFS1) loop as measured at four different powers.

changes. The phase margin at the UGF decreases by about 10° and a third −180° phase crossing appears at about 500 mHz. The digital filters are fixed with power, so the observed changes are due to a changing plant. This is precisely the result of the radiation pressure angular spring. I present the direct measurement of the plant's modified transfer function in Section 6.9.

6.6 Residual Angular Motion

The residual beam spot motion on the test masses is shown in Figure 6-7. We see the rms beam spot motion on the ETMs is 1 mm and on the ITMs it is 0.8 mm which meets the requirements of Section 4.1. The beam spot motion calibration method and results are presented in Appendix B.1.

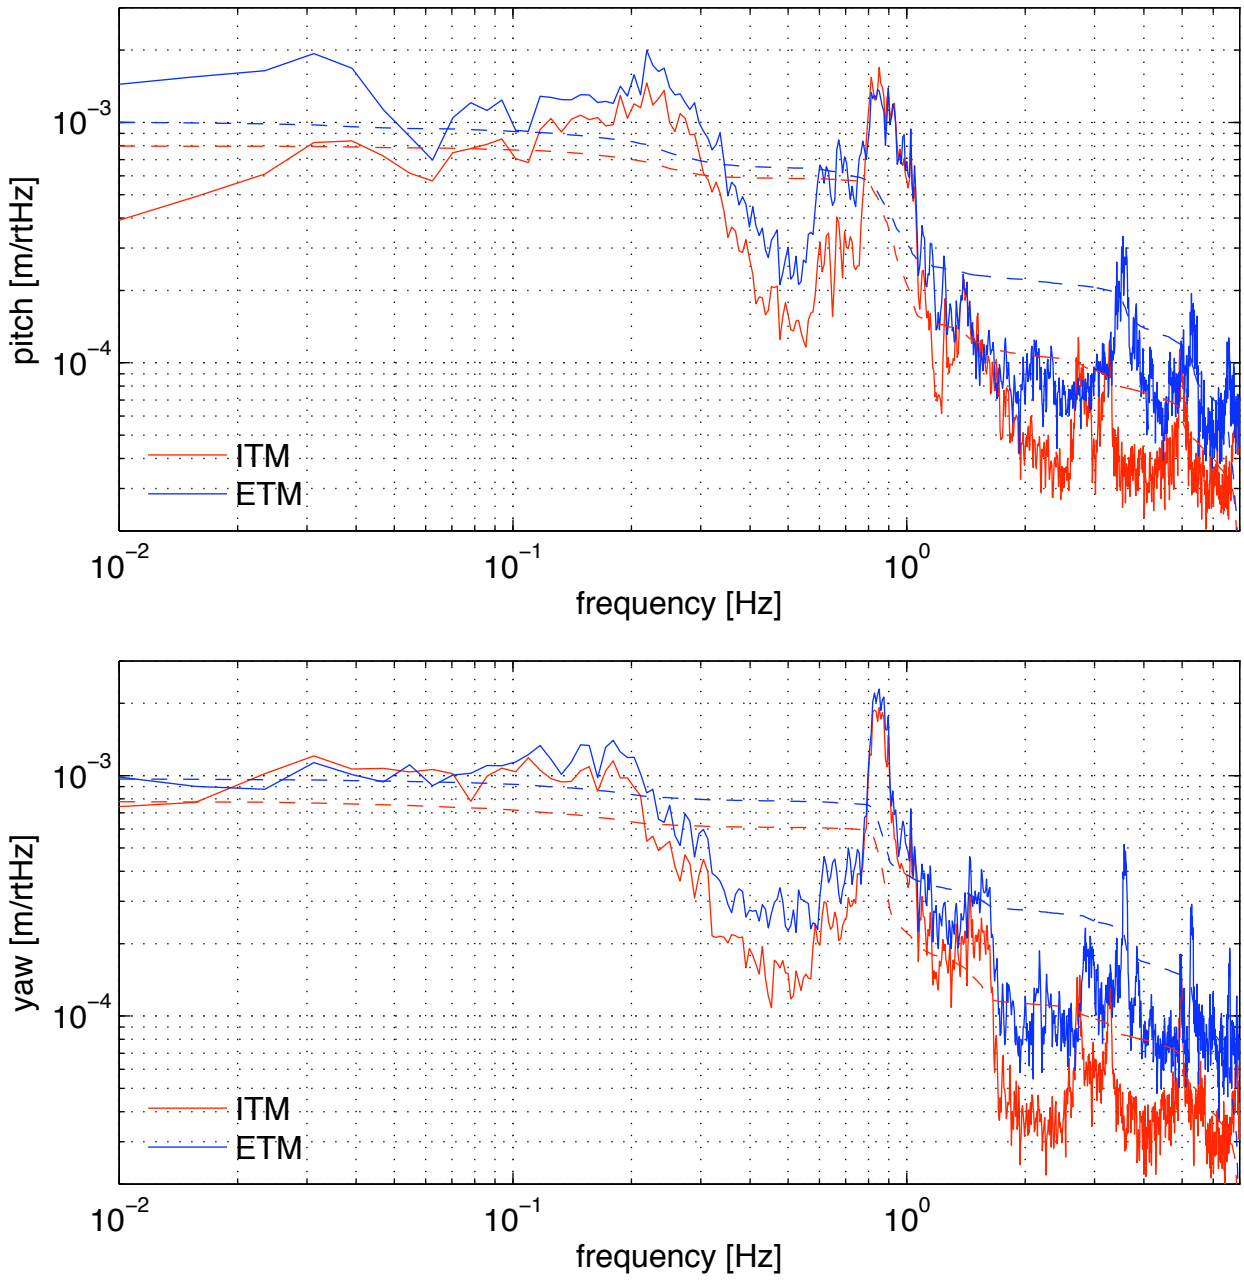

Figure 6-7. Beam spot motion on the ITMs and ETMs during a 16 W lock at night. Ground motion at the time of this measurement is shown in Figure C-7

We quantify the effect of the ASC on the radiation pressure eigenbasis degrees of freedom by comparing spectra of the residual eigenbasis motion during lock to the equivalent eigenbasis motion as witnessed by the optical levers out of lock. The comparisons cannot be perfect because the spectra are necessarily taken at different times and therefore with different seismic noise conditions. However, the effect of the ASC is stark at low frequencies where gain is high, as is seen in Figure 6-8. At 0.01 Hz angular motion of all degrees of freedom is suppressed by at least one order of magnitude. The typical residual rms angular motion is 10^{-7} rad/$\sqrt{\text{Hz}}$. The effect of the high gain for the differential soft degree of freedom, in particular, is seen here, where order of magnitude suppression is seen already at 1 Hz.

The magnitudes of the beam spot motion and the residual mirror motion are consistent and reasonable. For example, for 10^{-7} rad of soft or hard mode motion in one arm, the maximum cavity tilt and displacement are 0.12 μrad and 1.02 mm, respectively, as found in Table 5-1.

Figure 6-9 shows the same data as Figure 6-8 except in the mirror basis instead of the radiation pressure eigenmode basis. The ASC on/off comparison of mirror motion is interesting because it shows that the mirrors actually move more with respect to the ground when they are controlled by the ASC than when they are not controlled by the ASC. This is to be expected because the ground motion is different at each mirror and the job of the ASC is to control the motions of the mirrors with respect to each other, not with respect to the ground.

6.7 ASC to DARM Noise Budget

One of the most important figures of merit for the ASC is how much noise it contributes to DARM. As introduced in Section 4.1, the combination of beam spot motion on the mirror together with the angular motion of the mirror creates an angle to length (A2L) coupling. As long as the length displacement due to this coupling is well below the desired displacement sensitivity, the ASC has done a good enough job. However, being of the most complex of interferometer systems, the ASC does in fact limit the strain sensitivity.

To measure the coupling, a broadband noise injection rather than a typical swept-sine injection is necessary because of the non-linearity of the A2L process (in particular, it is the

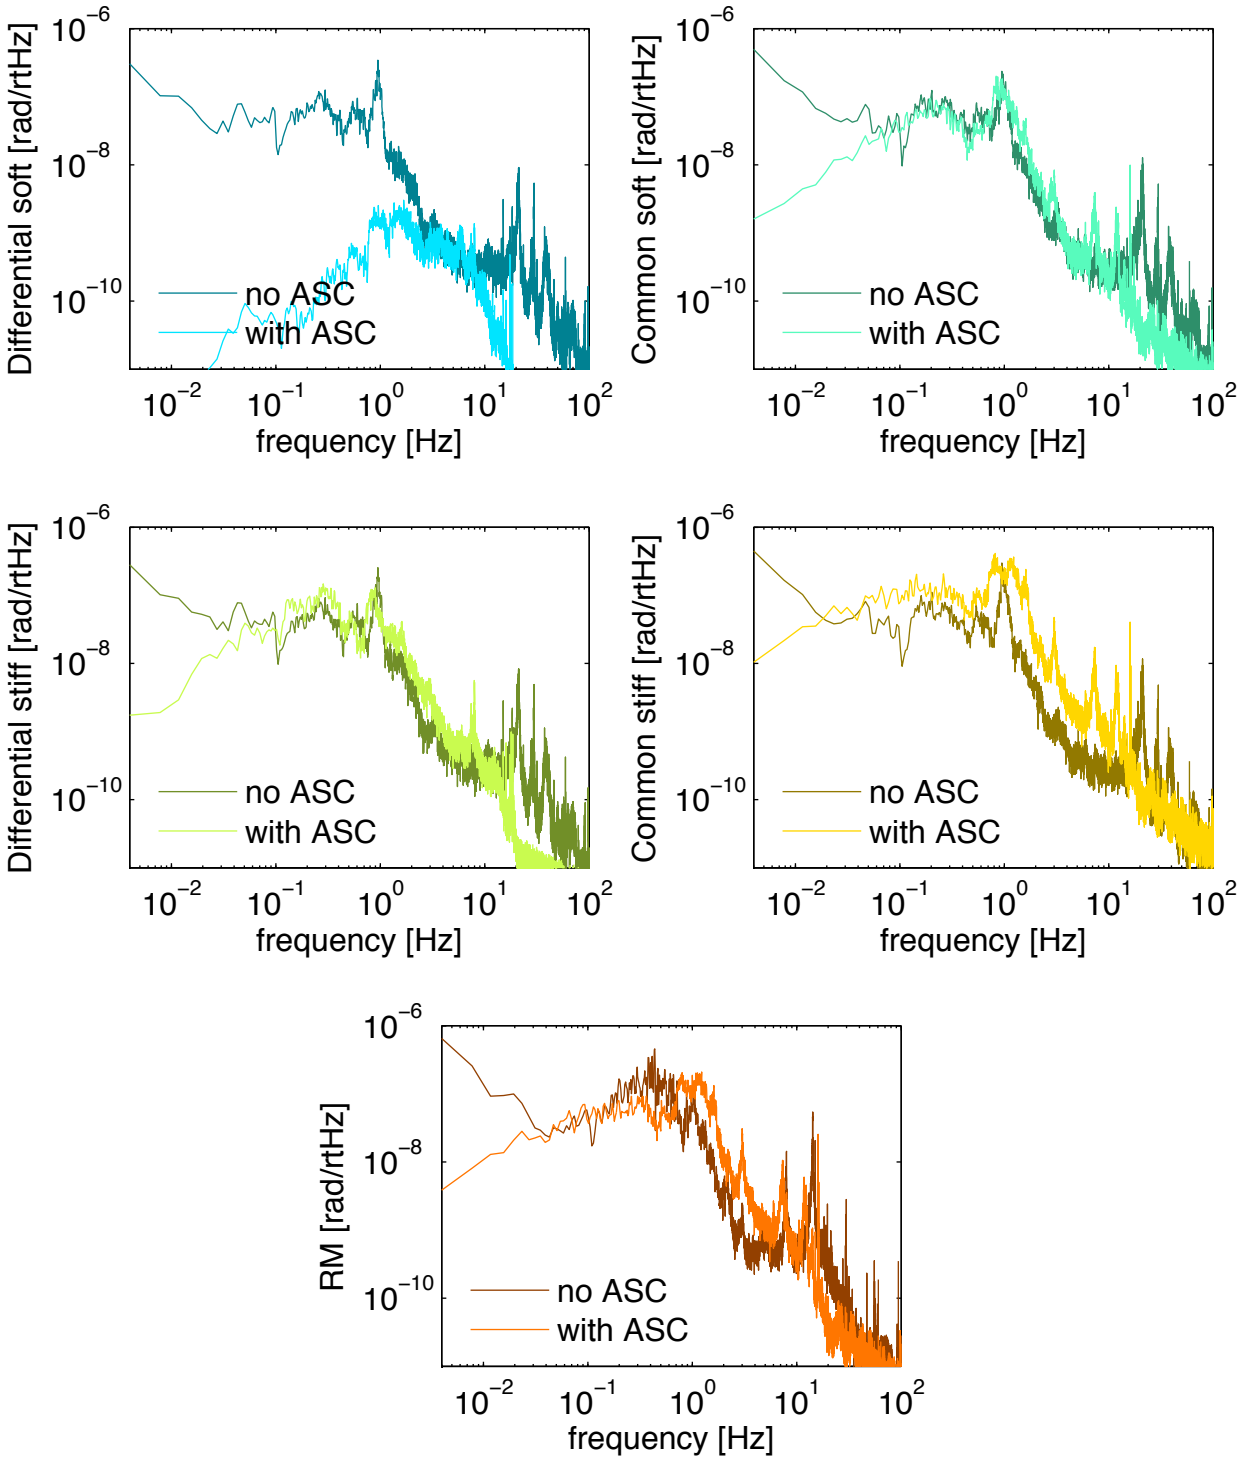

Figure 6-8. Demonstration of angular motion suppression down to 4 mHz due to the ASC. The background motion ("no ASC") is the RP eigenbasis reconstruction of optical lever signals when the interferometer is not locked. Data are taken 45 minutes apart, and the ground motions are shown in Figures C-4 and C-5. The differences in ground motion explains the discrepancies between 1 Hz and 10 Hz.

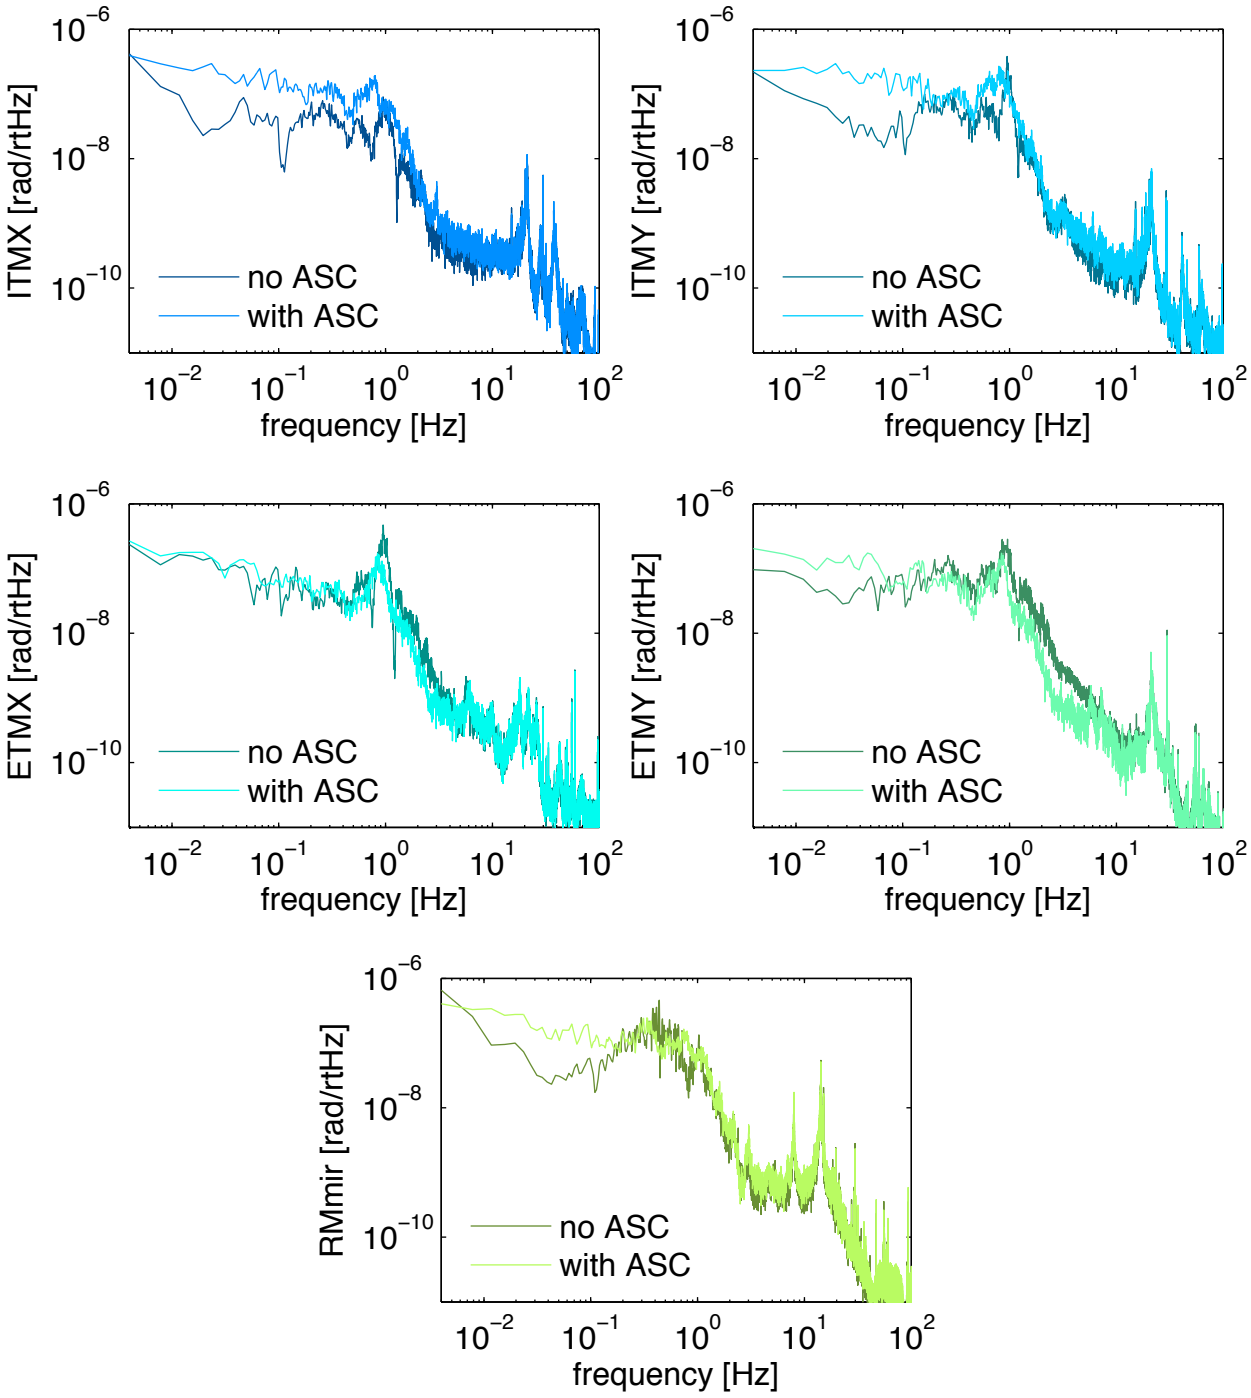

Figure 6-9. Individual mirror motion with and without the ASC. The mirrors move more with respect to their local grounds when the interferometer is controlled than when they're on their own. Data are taken 45 minutes apart, and the ground motions are shown in Figures C-4 and C-5.

convolution of the beam spot and mirror angle spectra). In addition, for the measurement itself to be linear, the amplitude of the noise injection must be only large enough for an effect to be observed; it must not overwhelmingly dominate. Therefore, for both DARM and the appropriate ASC channel, I subtracted quadratically a quiet, no-noise spectrum, q, taken immediately prior to the measurement from the spectrum with noise, n, so as to not assume that the entirety of the observed effect is due to the excitation:

$$N = \sqrt{n^2 - q^2}. \tag{6–10}$$

The transfer function from the ASC to DARM is then:

$$\text{ASC} \rightarrow \text{DARM} = \frac{N_{\text{DARM}}}{N_{\text{ASC}}} \tag{6–11}$$

which can be multiplied by an ASC signal at any time to determine a noise budget.

There are two places in the ASC loop from where I made transfer function measurements by putting in a 40 to 110 Hz broadband excitation: the suspension angular control point,[3] and the WFS error point.[4] Each provides ultimately the same information, but allows the ASC noises to be broken down in different bases. The view using the WFS (radiation pressure eigenmode) basis is useful for commissioning purposes, whereas the optic basis is better suited for including the ASC noise in full interferometer noise budget plots. The primary reasons are that the suspension control channel is recorded to disk and it also serves as the transfer function for the optical lever noise to DARM. Figure 6-10 shows the set of WFS to DARM measured transfer functions, as an example. The units are DARM meters per WFS error point digital counts, so that the digital WFS error signal may be multiplied by the transfer function at any time.

The noise budget for both pitch and yaw at a time when the interferometer was locked with 14 W input power is shown in Figure 6-11 for the optics basis and in Figure 6-12 for the WFS

[3] i.e. `L1:SUS-ETMX_ASCPIT_EXC`

[4] i.e. `L1:ASC-WFS1_PIT_EXC`

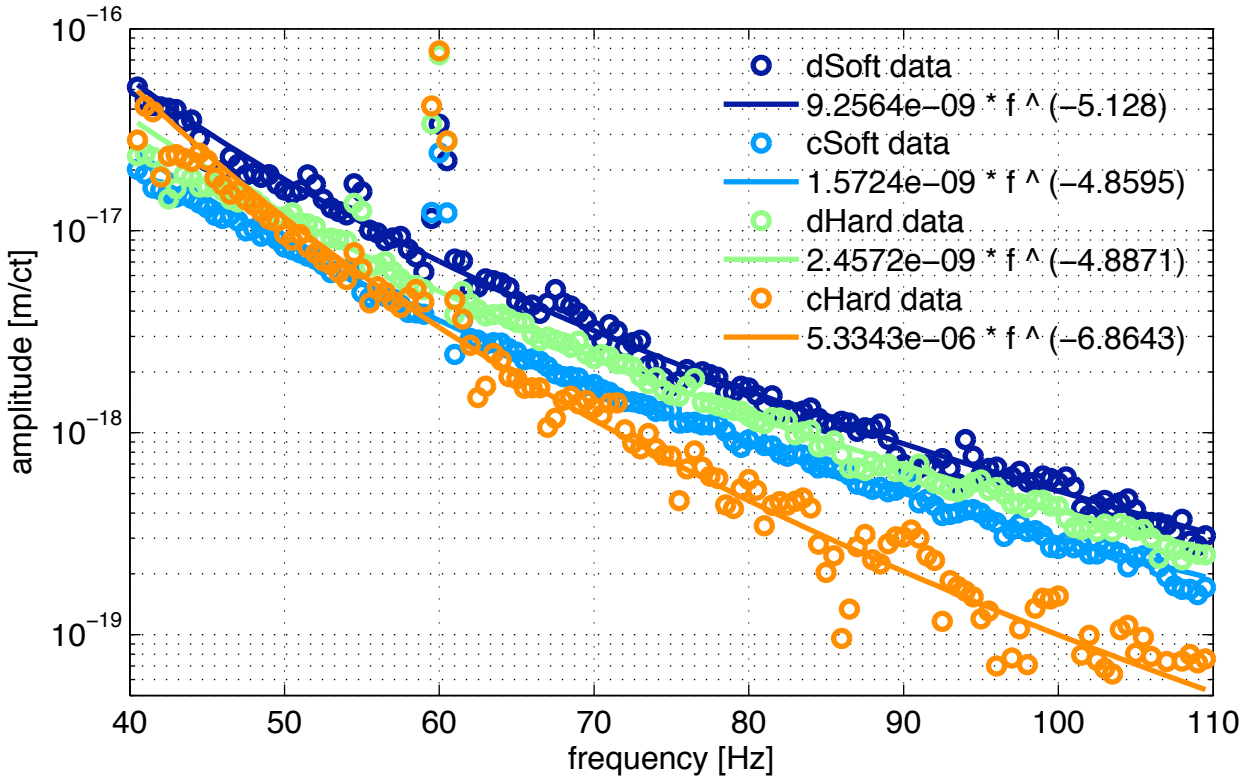

Figure 6-10. ASC to DARM transfer function for four of the five wavefront sensor loops. The RM to DARM transfer function could not be measured because the contribution is too small. The fitted curves can be multiplied by the WFS error signals at any time to calculate the ASC noise contribution to DARM.

basis. Each optic's contribution to DARM is the same within about a factor of two, except for the RM, which is not included in these plots. We were not able to measure the transfer function for RM motion to DARM because so large of an excitation was required to see an effect that the interferometer would lose lock. I proved this indication that the RM's contribution to DARM is indeed insignificant by observing no change in DARM upon turning off the WFS4 cut-off filters. In the WFS basis, the soft modes contribute more to DARM than the hard modes.

Because two different angular control signals, the WFS and the optical levers, independently contribute to the suspension angular control, we can separate their contributions in the noise budget. Furthermore, we compute the quadrature sum of the pitch and yaw contributions, assuming these two degrees of freedom are de-coupled, thus creating an upper limit for the ASC

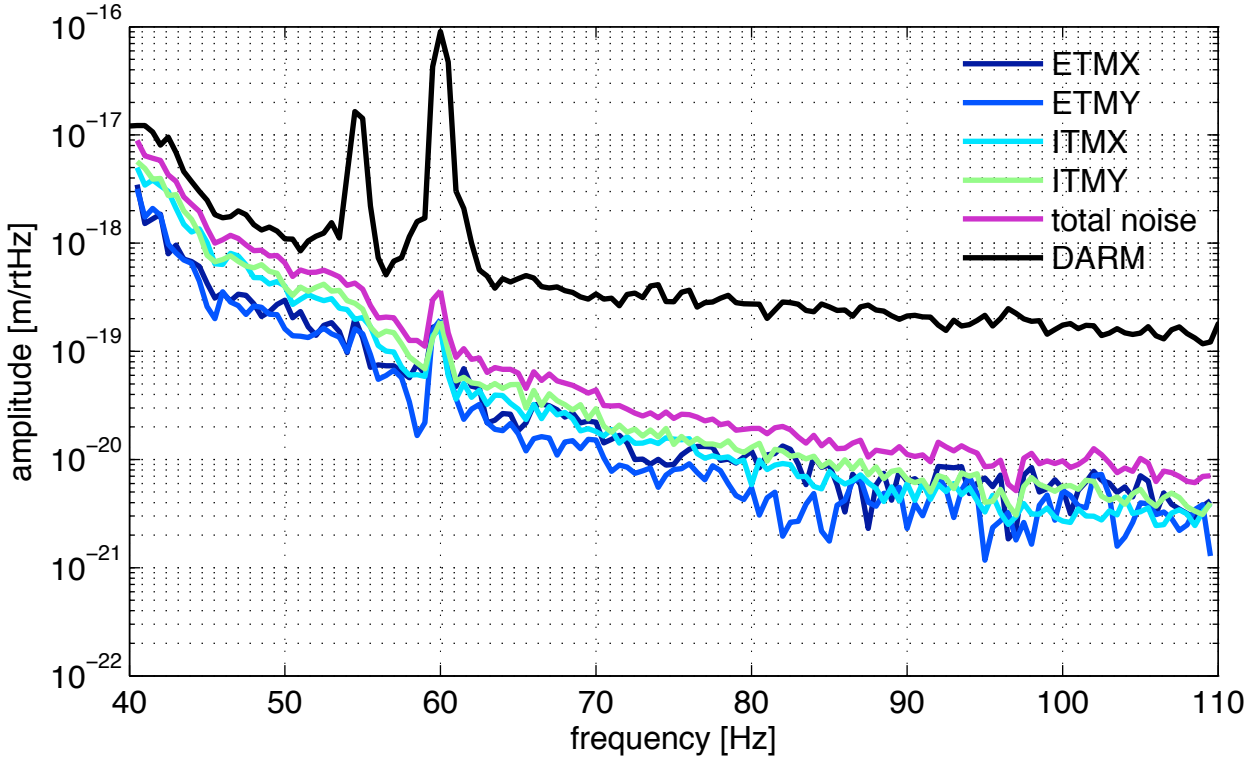

A

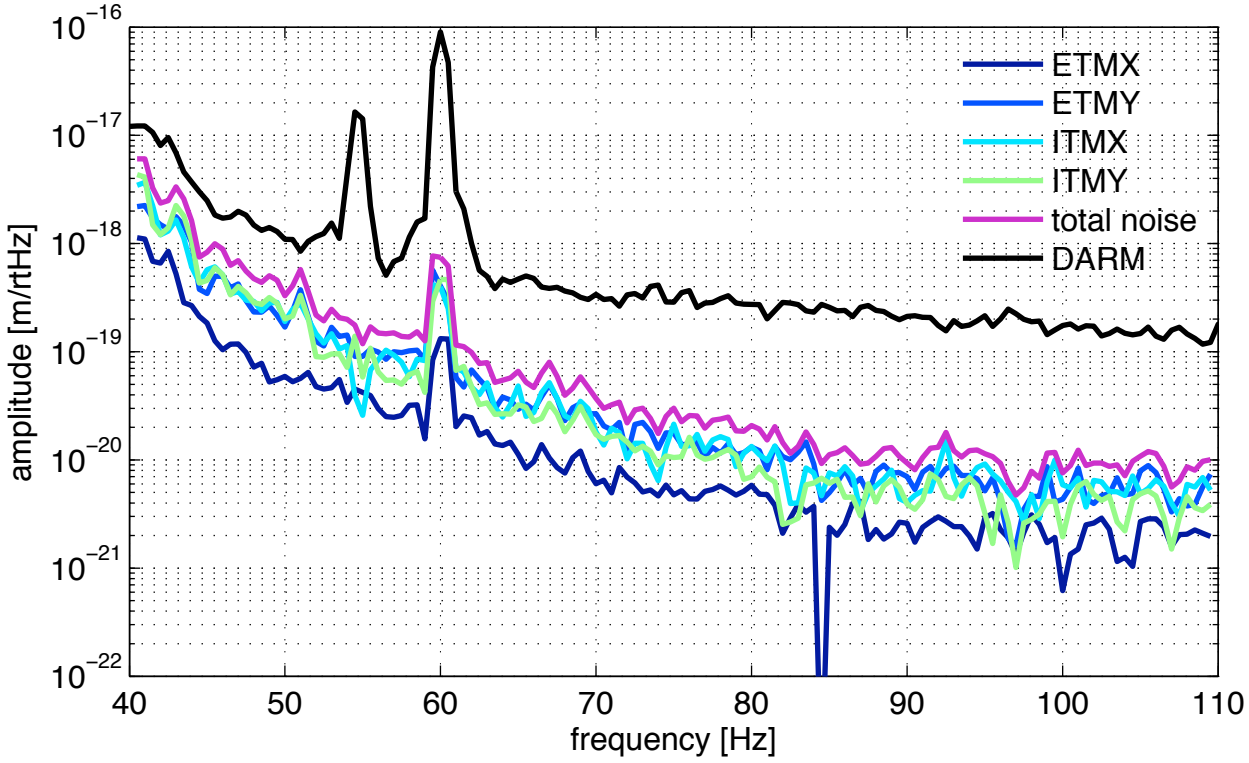

B

Figure 6-11. Optic to DARM noise budget during a 14 W lock. A) Pitch. B) Yaw.

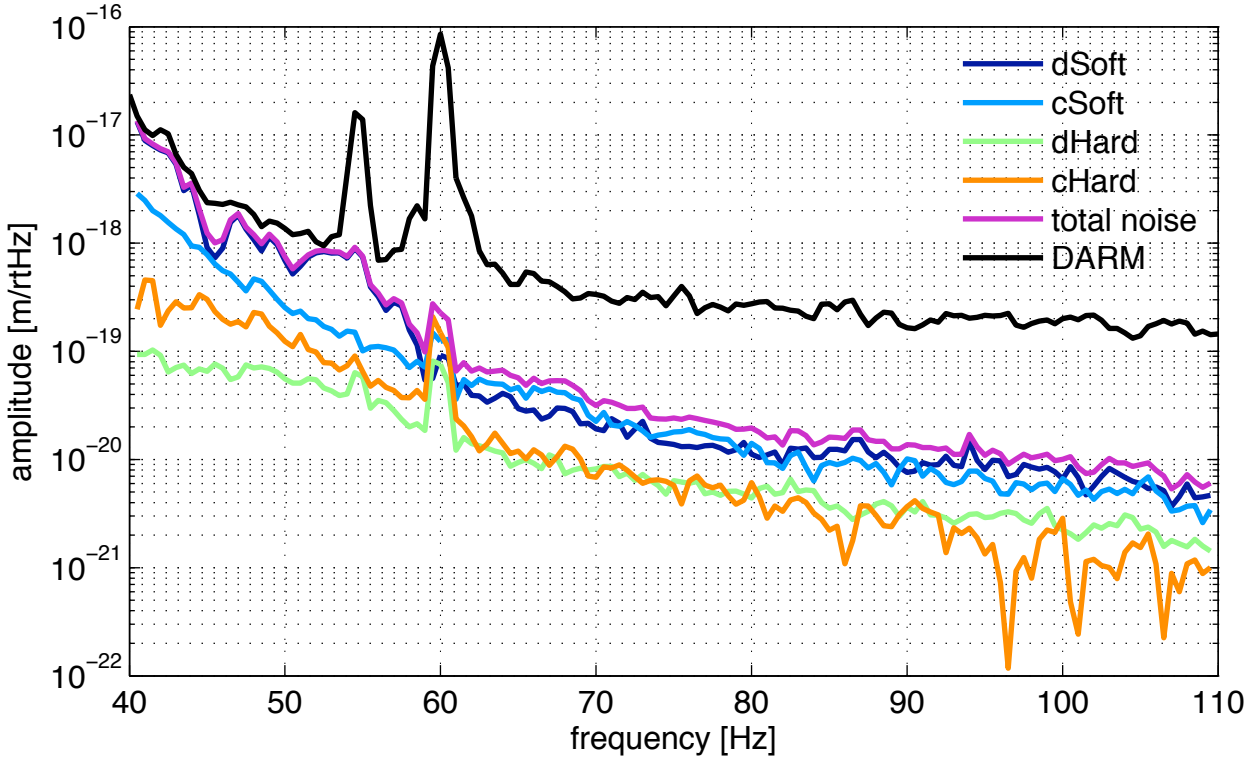

A

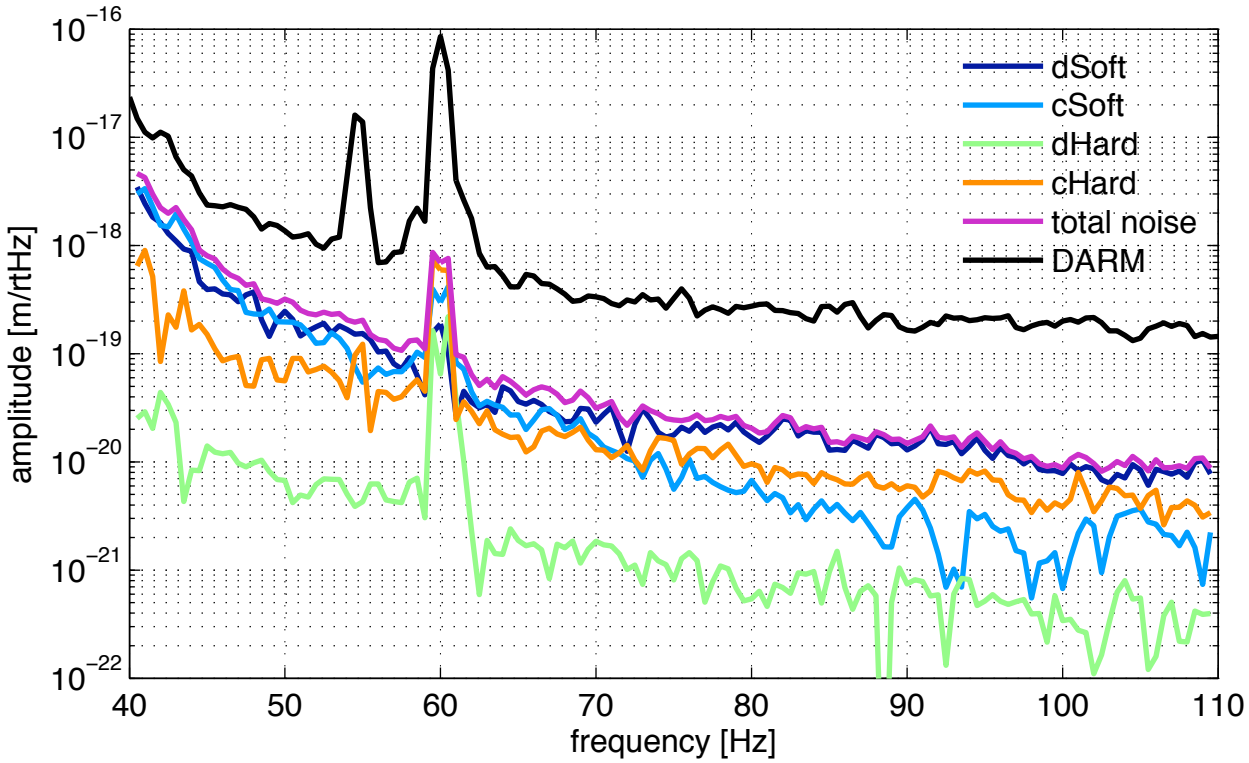

B

Figure 6-12. Wavefront sensor to DARM noise budget during a 14 W lock. A) Pitch. B) Yaw.

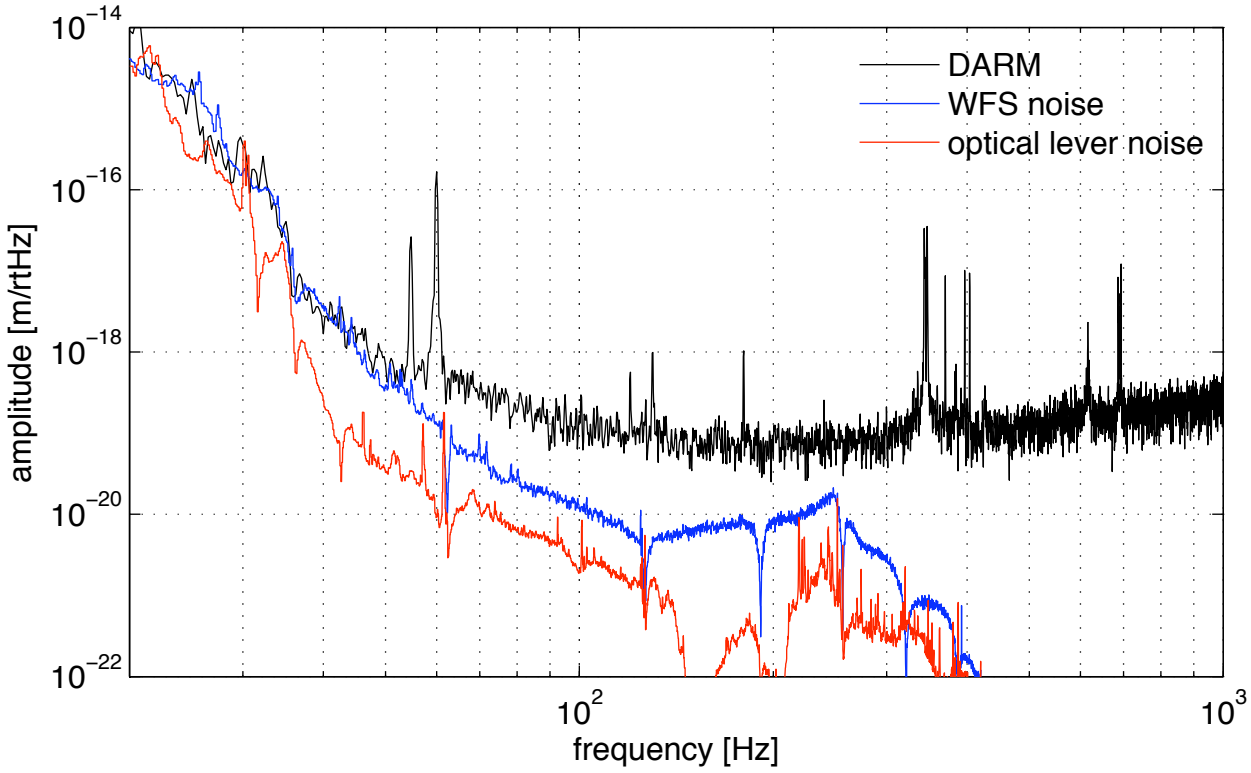

Figure 6-13. Total WFS and optical lever noise contribution to DARM during a 16 W lock at night. Pitch and yaw contributions are added in quadrature under the assumption they are de-coupled. Seismic spectra at the time of this measurement are found in Figure C-8.

contribution to the DARM noise budget. Figure 6-13 shows the final summary of ASC noise in DARM for a 16 W lock at night.

The important message is that the angular sensing and control is, in fact, a limiting noise source for frequencies between 20 and 55 Hz. The ASC becomes less and less of a primary noise source as frequency increases, and by 100 Hz the ASC noise floor is a factor of 10 below DARM. The specific structure of the noise contributions, including the apparent notches, is a direct result of the shape of the control filters. Imperfections in the estimate of DARM below 50 Hz arise because the transfer function is not perfectly stable in time. The seismic noise contribution to DARM (not shown) does in fact sit just below the ASC floor, so the interferometer sensitivity to GWs is not dramatically hindered by the ASC. An example, however, of how to reduce the ASC noise floor, is found through evaluating the WFS control filters.

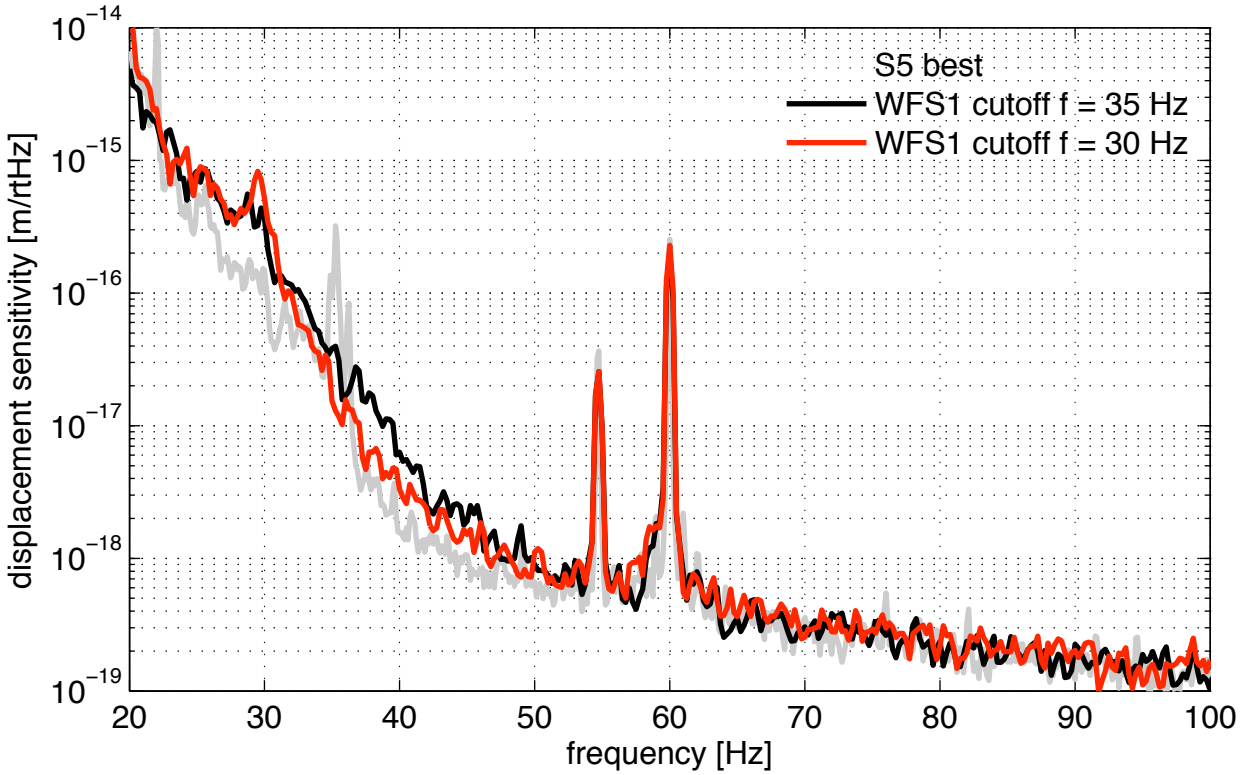

Figure 6-14. Effect of the WFS1 lowpass filter cutoff frequency on strain sensitivity.

The cut-off frequency of the lowpass filters for the WFS control are of particular importance in the DARM noise budget. The lowpass filter is necessary for suppressing the impression of sensing noise on suspension control signals. Steepening the cut-off frequency results in less sensing noise impression, but each pole used to achieve the steeper drop-off introduces an extra 90° of phase loss. Likewise, lowering the cut-off frequency reduces noise impression, but it pushes the phase loss to lower frequencies. The effect is a decrease in phase margin of the WFS loops for a particular UGF, which leads to gain peaking and a greater likelihood of loop instability. A fine balance must therefore be found between loop stability and noise impression. Because WFS1 (differential soft) is the dominant contibutor to the noise budget, we put forth effort to tune its low pass filter's frequency. Figure 6-14 demonstrates the effect on DARM of decreasing the WFS1 cut-off filter frequency from 35 Hz to 30 Hz.

6.8 Seismic Feed-forward to the ASC

We show in Section 6.7 that angular control limits the strain sensitivity at low frequencies and we see from Figure 1-2 that Advanced LIGO aims to improve the low frequency sensitivity by up to two orders of magnitude. Although much of the improvement comes from better seismic isolation, it is prudent to develop methods to reduce the angular control noise for future interferometers. A promising technique that I investigated using Enhanced LIGO data is seismic feed-forward to the ASC.

The concept is to use Wiener filtering [57] to predict the ASC signals from seismometer signals and feed-forward the filtered sesimometer data in real time. The Wiener filter coefficients are chosen to minimize the mean square difference between the WFS error signals and the filtered seismometer signals. For this to work, there must be a correlation between the seismic and ASC signals. I designed FIR Wiener filters from simultaneous seismometer and ASC error signal time series and found that with a sufficiently long Wiener filter (on the order of 1024 taps) and 30 minutes worth of data, we can accurately predict the ASC error signal from 0.1 to 20 Hz. Figure 6-15 shows the comparison of the nominal dSoft (WFS1) error signal with the reduced error signal that one can hope to achieve through feed-forward. Called the residual, this reduced error signal is the result of subtracting the seismometer-predicted ASC error signal from the original error signal. We see that if implemented, feed-forward may reduce the rms angular mirror motion by a factor of two. Further studies may reveal even a greater reduction can be achieved. One aspect of feed-forward that still requires investigation is the stationarity of the Wiener filter over time.

There are several ways in which feed-forward might be accomplished. The two primary options are to send the seismometer-predicted ASC drive to either the mirror coils or to the hydraulic external pre-isolator (HEPI) seismic isolation tables. The benefit of both of these methods is that they would physically reduce the mirror motion and therefore reduce the need for ASC feedback. This would allow for more flexible loop shape design such as lower cut-off filters and result in less impression of ASC noise in DARM. The simplest of these two options is

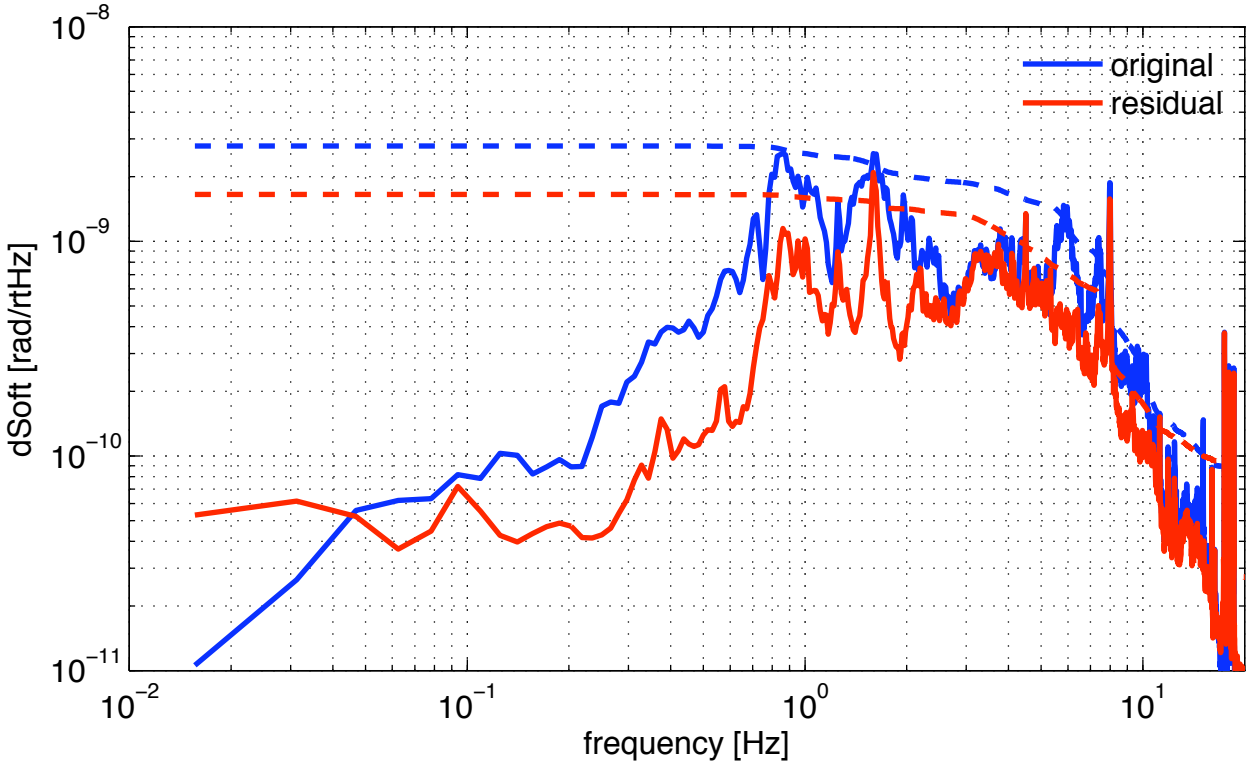

Figure 6-15. Demonstration of potential reduction of WFS error signals using seismic feed-forward.

sending the signal to the coils because the transfer function from the coils to the mirror motion is well known. A disadvantage is that the total coil currents would not be reduced, meaning the noise from magnetic domain flipping, Barkhausen noise, would not be reduced. Feeding-forward to HEPI would reduce the Barkhausen noise, but would require a carefully measured transfer function from HEPI to angular mirror motion. Seismic feed-forward to HEPI to reduce LSC signals was demonstrated during Enhanced LIGO [66].

An alternative method to reducing the ASC noise in DARM does not involve the seismometers or feed-forward to the ASC, but feed-forward to DARM itself. From a carefully measured ASC to DARM transfer function, the predicted contribution of ASC noise in DARM can be subtracted from DARM in real time. This type of feed-forward was implemented in Initial and Enhanced LIGO for the Michelson (MICH) and power recycling cavity (PRC) loops [67, Ch. 2].

Any of these methods may be implemented in post-processing of the DARM data, but there are benefits to doing it in real-time. The primary advantage is the actual reduction of mirror

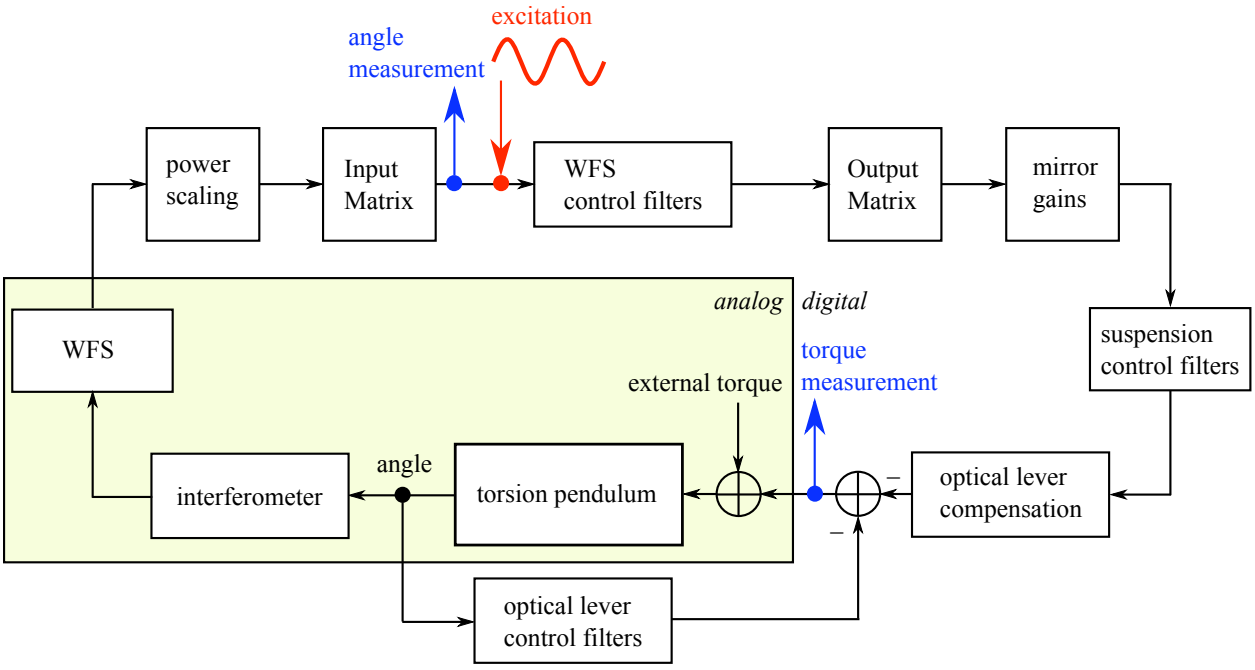

Figure 6-16. Demonstration of radiation pressure eigenbasis torque to angle transfer function measurement. Through a proper choice of measurement locations within the ASC servo, the plant's transfer function can be singled out.

motion which would make for an all-around more stable interferometer. Second, eliminating known noises from the strain spectrum is very useful for commissioning. It allows us to then see other noises in real-time and to be more efficient at the endless game of noise hunting.

6.9 Experimental Measurement of the Radiation Pressure Angular Spring

I directly measured the expected radiation-pressure-modified torque to angle transfer function of the LIGO arm cavities that I first introduced in Chapter 5. Through a clever measurement method, I witness both the stable and unstable modes without any post-data-taking manipulation.

The digital control system in which the angular control feedback system is implemented provides a convenient milieu in which to measure the response of the optomechanical system. By injecting a disturbance somewhere in the loop and measuring the response at selected points in the loop, we can produce a measurement of the optomechanical system that is not sensitive to the details of the control system. Here we use this system to produce measurements of the opto-mechanical plant at several different operating powers, demonstrating the modifications due

to radiation pressure, i.e. the soft and hard modes, sometimes also referred to as the Sidles-Sigg effect.

The various elements of the plant and the control system are depicted in Figure 6-16. For this measurement, I took transfer functions from the torque input to the resulting angular displacement (as measured by the WFS), both in the radiation pressure eigenbasis. Simultaneously, I injected an excitation into the control leg of the servo loop, as is done for the sensing matrix mesaurement described in Section 6.1.1. The resulting measurement reproduces the transfer function of the opto-mechanical plant, independent of the control system.

Results are shown in Figures 6-17 (hard mode) and 6-18 (soft mode); least-squares fits of second-order transfer functions are made to the data. In the hard mode plot, we can clearly see the increase of the resonant frequency with power, from $\sim$ 0.65 Hz at 1 W input power to $\sim$0.95 Hz at 10 W input power. Simultaneously, in the soft mode plot, we see the resonance decrease in frequency as the power is increased from 1 W to 6 W. When the input power is increased to 10 W and beyond, the resonance disappears; the plant has become statically unstable.

These measurements show a clear confirmation of the Sidles-Sigg theory and demonstrate a successful power-independent diagonalization of the sensing and control of the opto-mechanical plant.

6.10 Summary

The ASC performed as it needed to for Enhanced LIGO to operate at higher powers without introducing an excess of control noise. The ASC does limit strain sensitivity at the lowest end of the detection band, but we offer possible solutions to decrease that noise in future detectors. We directly measured the radiation pressure angular spring, confirming our theoretcial understanding of the basic physics that drives the ASC design. Although the ASC in neither base was limiting us, the experience gained for Advanced LIGO is priceless. Advanced LIGO will have heavier mirrors with a different geometry such that radiation pressure torque will not play so large a role [68] and its power recycling cavity will be stable [18].

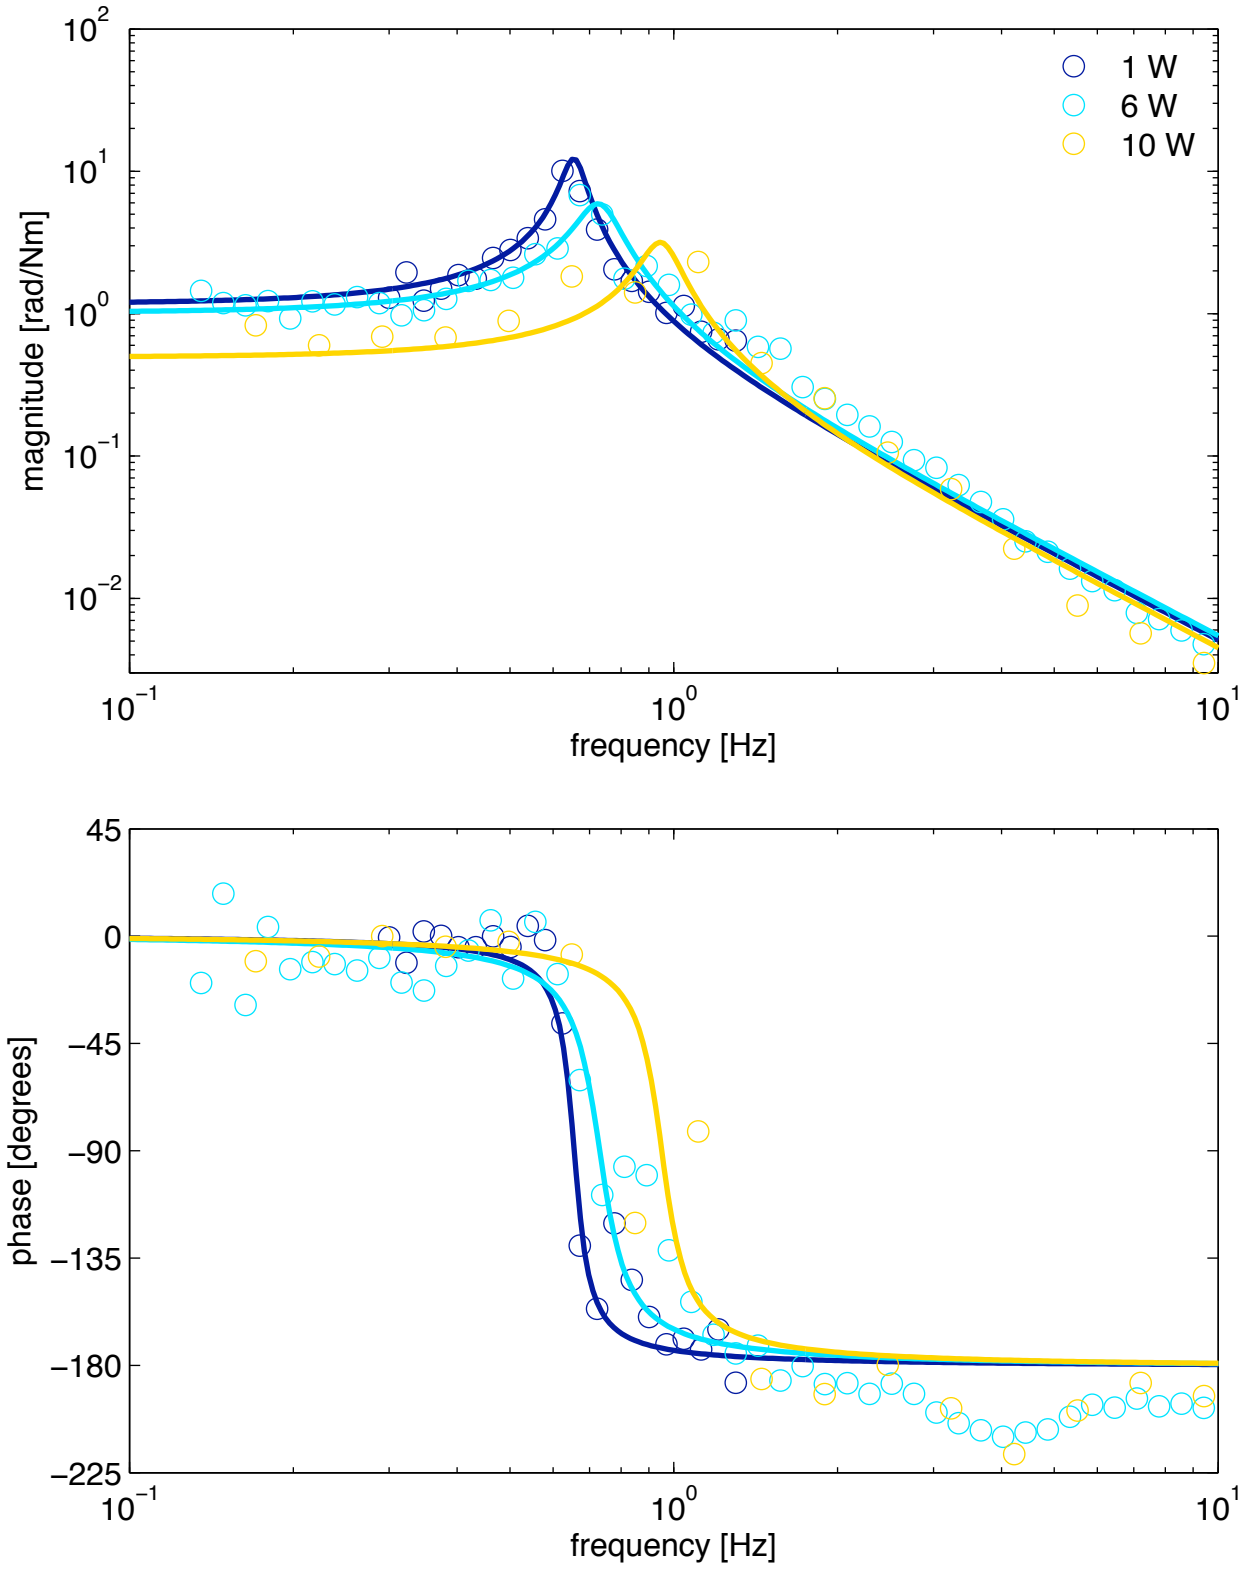

Figure 6-17. Hard opto-mechanical mode measurement and fit for several powers.

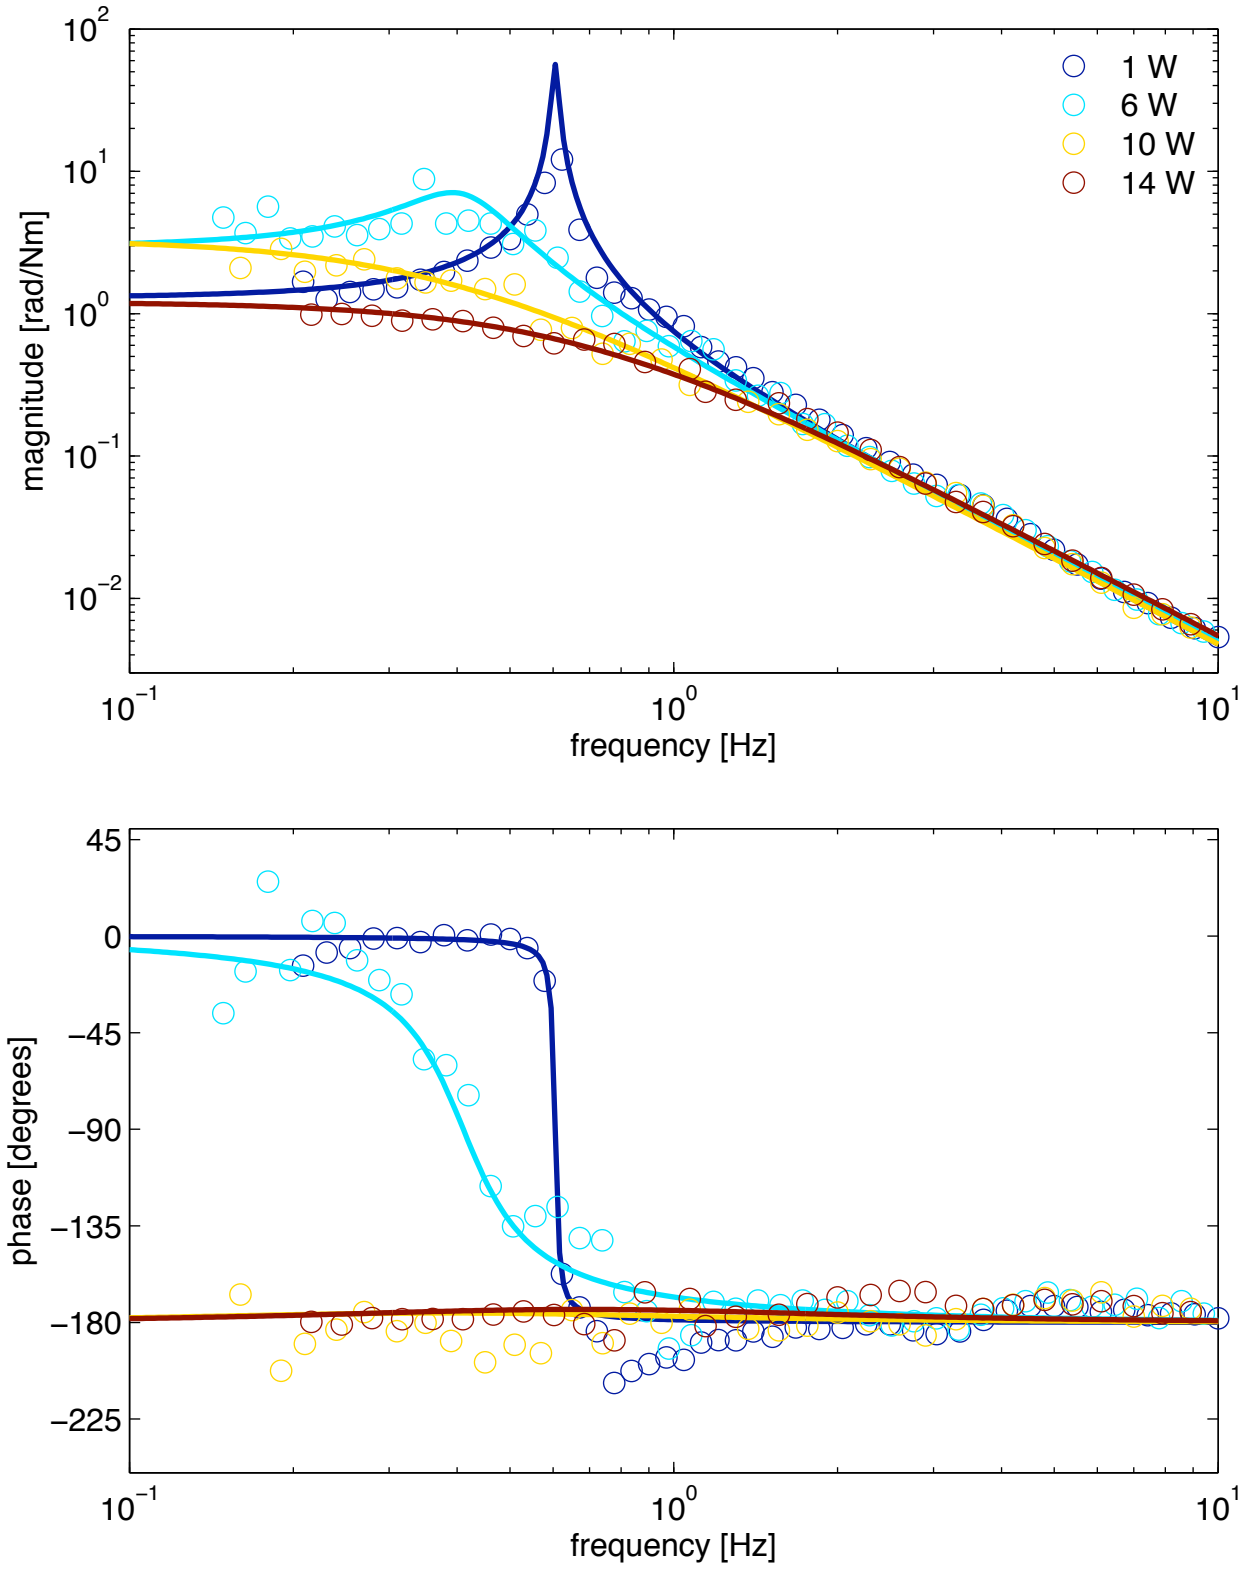

Figure 6-18. Soft opto-mechanical mode measurement and fit for several powers.

CHAPTER 7
CONCLUSIONS

7.1 Higher power in Enhanced LIGO

The broad result of the Input Optics and Angular Sensing and Control upgrades along with the work of many other people on many other subsystems was successful operation of both the Hanford and Livingston interferometers at the highest of laser powers ever used. The typical input powers used during Initial LIGO were between 5 and 7 W and during Enhanced LIGO typical input powers were between 8 and 20 W, as shown by the histograms in Figure 7-1.

A snapshot of how the interferometer reacts to increases in laser power is found in Figure 7-2, which shows the time series of several channels during a typical lock loss and re-lock of the interferometer followed by an increase in input power to 14 W. The increase in the fluctuations of NSPOB, the sideband power in the PRC, prior to the lock loss at $t = 250$ sec indicates PRC instability. The MC regains lock almost immediately and the power is increased to 1 W. The dense region of NSPOB signal shows the flashes of resonance as the interferometer tries to acquire lock, which it ultimately does at around $t = 500$ sec. The laser power is subsequently increased to 8 W, and pulses of TCS power turned on to hasten the ITM cooling process in preparation for the laser power increase to 14 W. All along, the power in the arm cavities increases as well as that reflected from the interferometer. NSPOB is a bit more "hairy" at 14 W than it is at 8 W, but this lock is otherwise stable and provides a neutron star binary inspiral range of 15 Mpc and lasts for 5 hours until an earthquake kills it.

The improvement in strain sensitivity in the shot-noise-limited region due to operating at higher laser power is seen in Figure 7-3. There is a factor of two improvement from Initial LIGO to Enhanced LIGO from about 300 Hz on. The corresponding best neutron star binary range increased from 15 to 20 Mpc. Although more power improves the shot-noise-limited region of the strain spectrum, it does introduce many complications for interferometer operation that must be carefully addressed.

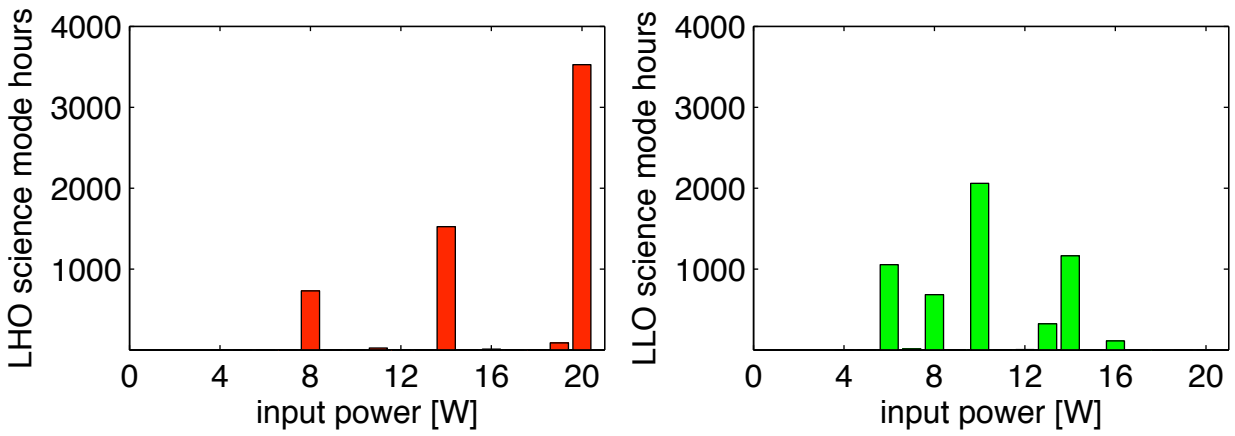

Figure 7-1. Histogram of input powers used during S6 at each site.

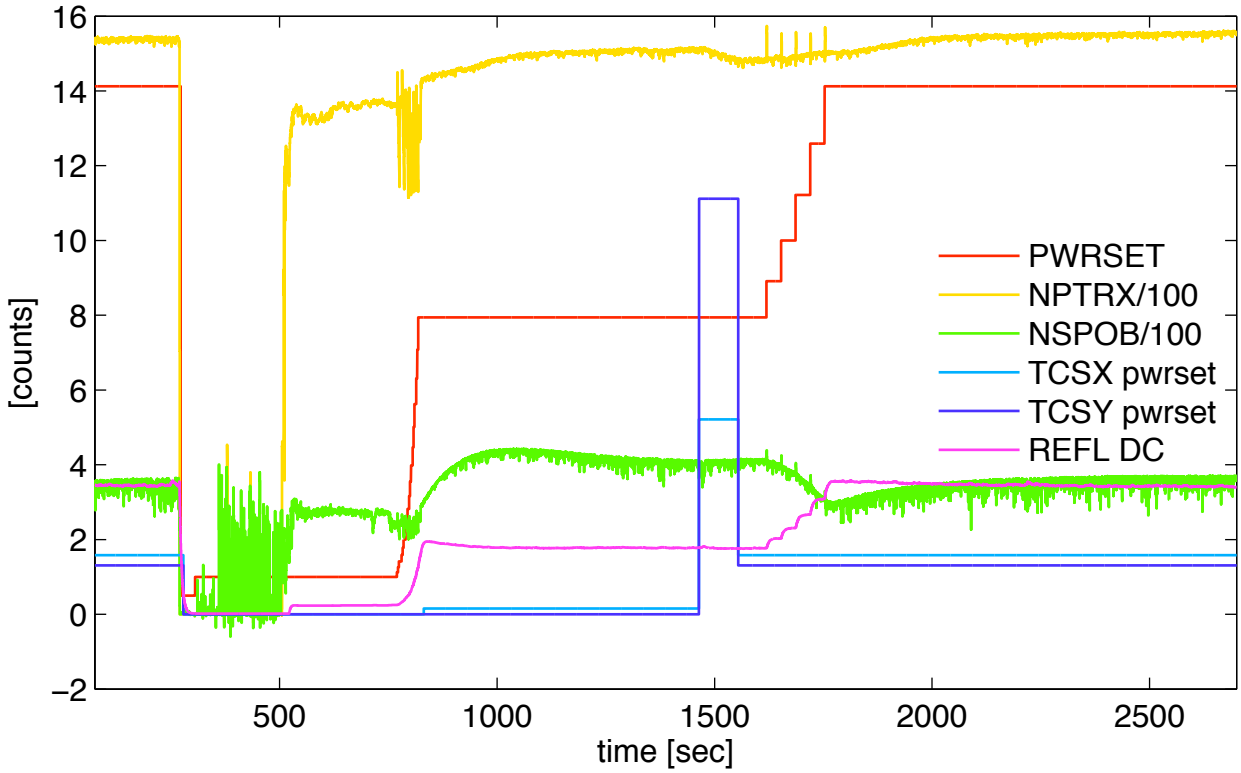

Figure 7-2. Time series of interferometer signals showing a typical lock loss (at about 250 sec) and re-lock (at about 500 sec) followed by an increase of power to 14 W. The y-axis is in units of Watts for the input power, TCSX, and TCSY traces. The other traces are displayed in digital counts. Data is from July 23, 2009.

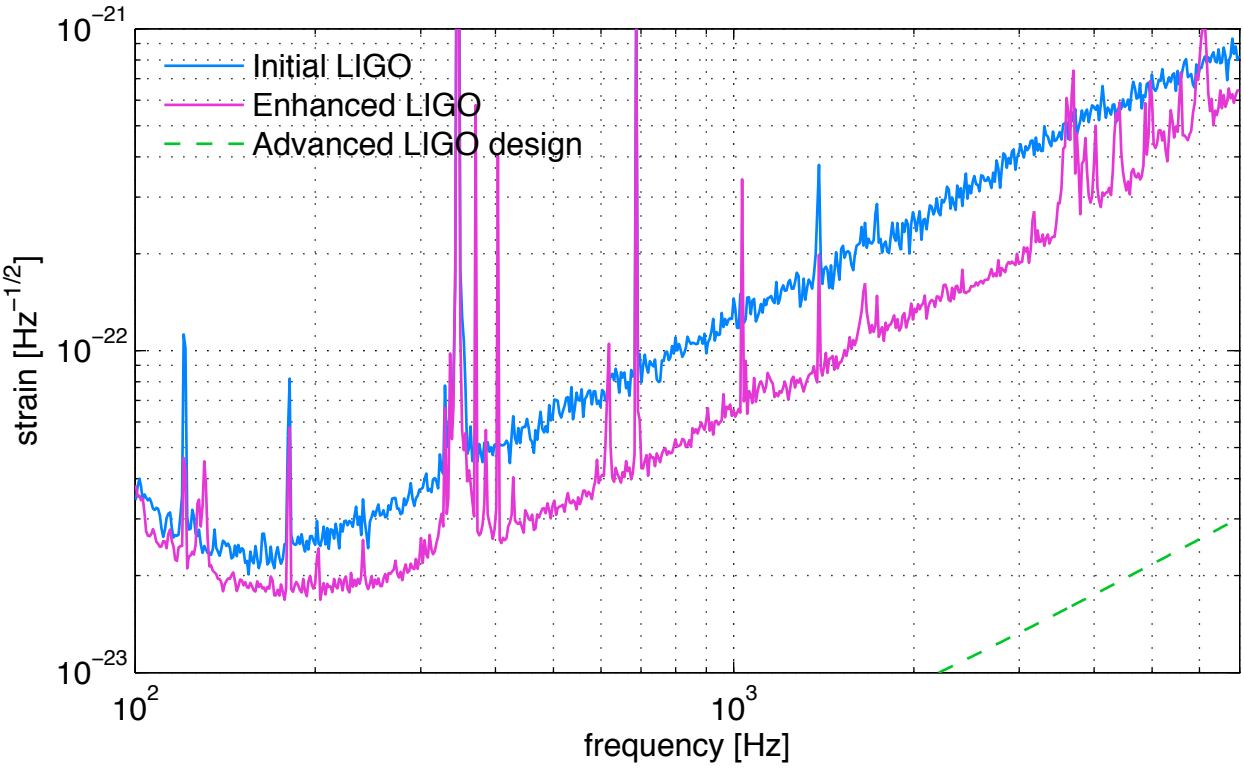

Figure 7-3. Zoom of the shot-noise-limited noise floors of the Initial LIGO and Enhanced LIGO detectors. The improved Input Optics and Angular Sensing and Control enabled the increase in power for Enhanced LIGO. The shot-noise-limited strain sensitivity improved by a factor of two.

7.2 Summary

We described the design of the Enhanced LIGO Input Optics and Angular Sensing and Control subsystems and presented measurements characterizing the systems and their performances when operating with record laser powers. Upgrades to the two systems were necessary for allowing higher laser powers, for improving the efficiency of sending light into the interferometer, and for keeping light in the interferometer once it is there. Higher power in the interferometer improves the shot-noise-limited noise floor, and we succeeded at operating the interferometers with more than twice the highest of powers achievable during Initial LIGO. The Enhanced LIGO shot-noise-limited sensitivity did indeed reach record levels, improving the chances of gravitational-wave detection.

In addition, we directly measured the stable and unstable opto-mechanical modes of the Fabry-Perot arm cavities. We witness the expected effect of radiation pressure torque, demonstrating a clear understanding of the physics that will affect future generations of laser interferometers for gravitational-wave detection. Furthermore, we successfully controlled the stable and unstable opto-mechanical modes without contaminating the gravitational-wave readout in the frequency range of interest.

APPENDIX A
INPUT OPTICS SUPPORTING MATERIAL

A.1 Phase Modulation

Phase modulation multiplies carrier light with field $E_0 e^{i\omega t}$ by $e^{i\Gamma \sin(\Omega t)}$, where Γ is the modulation index and Ω is the frequency of the phase modulation. Using the Jacobi-Anger expression,

$$e^{iz\sin\theta} = \sum_{n=-\infty}^{\infty} J_n(z) e^{in\theta}, \tag{A–1}$$

where J_n are the Bessel functions, we can write the first few terms (n = 0, 1, -1) of the phase-modulated field:

$$E_{\text{modulated}} = E_0 J_0(\Gamma) e^{i\omega t} + E_0 J_1(\Gamma) e^{i(\omega+\Omega)t} + E_0 J_{-1}(\Gamma) e^{i(\omega-\Omega)t} + \ldots \tag{A–2}$$

We see that both an upper and lower primary sideband are created, with frequencies $\omega + \Omega$ and $\omega - \Omega$. Phase modulation does produce an infinite number of sidebands, yet the amplitude of the Bessel function decays rapidly with higher $|n|$, so only this first set of sidebands are significant.

A.2 Mode Cleaner Pole

Optical cavities act as low pass filters for intensity variations of the light sent into them. The model for an intensity noise transfer function of a cavity is that of a single pole:

$$\frac{E_{after}}{E_{before}} = \frac{1}{1 + s/s_0} = \frac{s_0}{s_0 - s} \tag{A–3}$$

where s is a complex parameter. However, we are interested in only purely sinusoidal variations in intensity so we let s be purely imaginary, $s = i\omega$, where ω is an angular frequency.

We measured the intensity noise transfer function of the Livingston mode cleaner upon completion of the Enhanced LIGO Input Optics upgrade. We modulated the intensity of the laser light going into the MC by injecting a swept-sine excitation in `L1:PSL-ISS_EXC` and measured the power variation of the light in two places: before and after the mode cleaner. We used a single photodetector (PDA55) in order to eliminate the PD response, and therefore made the measurement twice. We ensured there was 1 V DC on the PD in both locations.

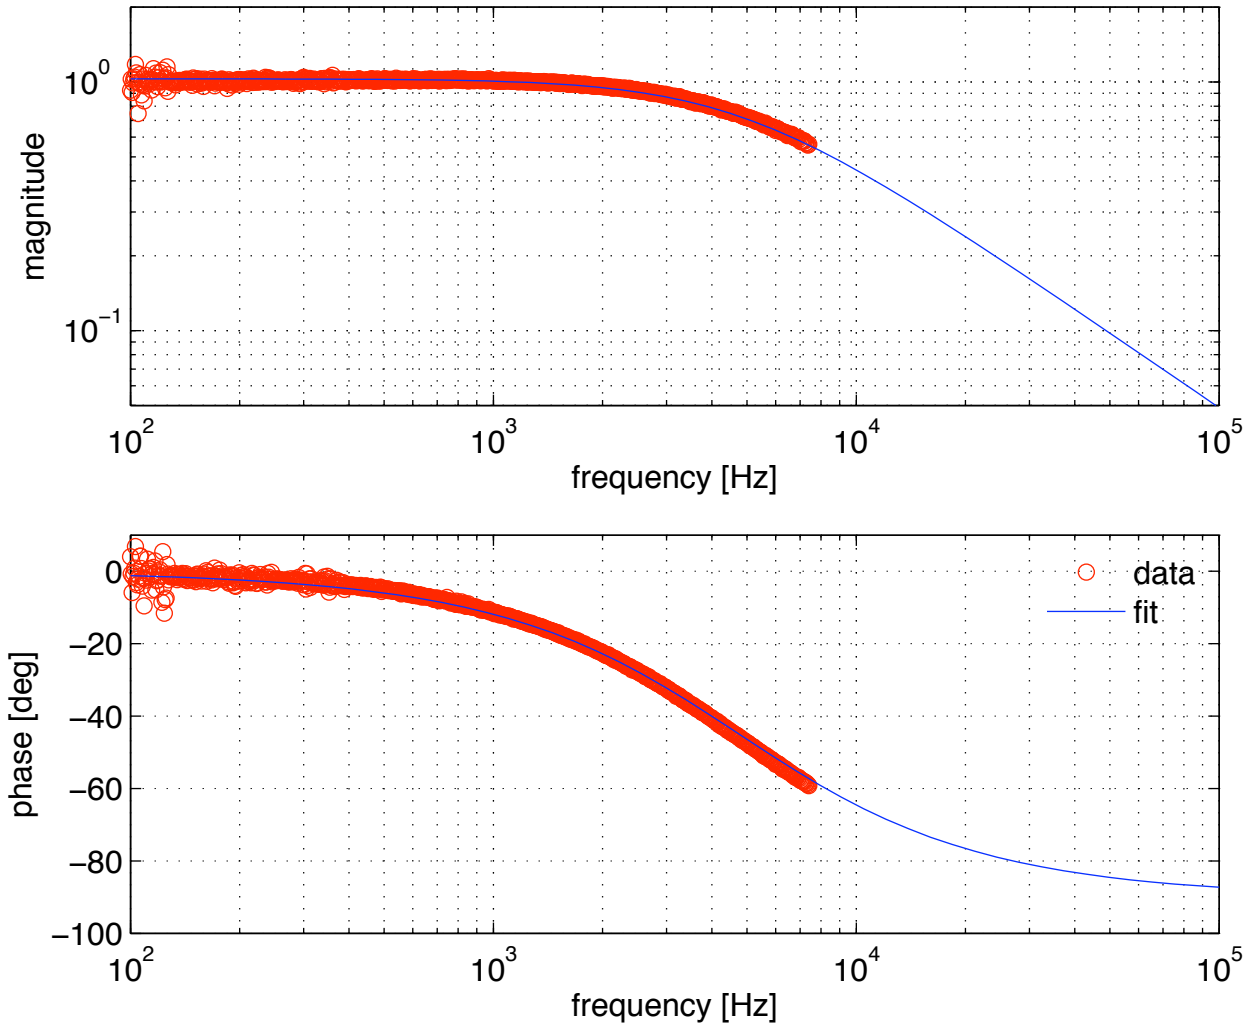

Figure A-1. Livingston mode cleaner intensity noise transfer function. Red open circles are data; solid blue line is a single pole fit. Relating fit parameters to the model, the pole frequency is $f_p = 4762$ Hz.

Figure A-1 shows the transfer function data and the fit (to both magnitude and phase simultaneously). The fit has a pole frequency of $f_p = 4762$ Hz. The $1/e$ ringdown time of the mode cleaner is therefore $\tau = 1/4\pi f_{MC} = 16.7$ µs and the finesse is $\mathscr{F} = \mathrm{FSR}/2f_p = 1282$.

A.3 Gaussian Beam on a Split Photodetector

The power per area of a Gaussian beam traveling along the z-axis is

$$p(x,y) = \frac{2P_0}{\pi w^2}\exp\left[\frac{-2x^2}{w^2}\right]\exp\left[\frac{-2y^2}{w^2}\right] \qquad \text{(A–4)}$$

where w is the beam radius at z. This has been normalized such that $\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(x,y)\,dx\,dy = P_0$. Then, for a beam displaced by x_0 from the center of a split photodetector, the power on the left side is

$$\begin{aligned} P_{left} &= \frac{2P_0}{\pi w^2}\int_{-\infty}^{x_0}\exp\left[\frac{-2x^2}{w^2}\right]dx\int_{-\infty}^{\infty}\exp\left[\frac{-2y^2}{w^2}\right]dy && \text{(A–5)}\\ &= \sqrt{\frac{2}{\pi}}\frac{P_0}{w}\left[\int_{-\infty}^{0}\exp\left[\frac{-2x^2}{w^2}\right]dx+\int_{0}^{x_0}\exp\left[\frac{-2x^2}{w^2}\right]dy\right] && \text{(A–6)}\\ &= \sqrt{\frac{2}{\pi}}\frac{P_0}{w}\left[\frac{w}{2}\sqrt{\frac{\pi}{2}}+\frac{w}{2}\sqrt{\frac{\pi}{2}}\int_0^{\sqrt{2}x_0/w}\exp\left[-t^2\right]dt\right] && \text{(A–7)}\\ &= \frac{P_0}{2}\left[1+\operatorname{erf}\left[\frac{\sqrt{2}x_0}{w}\right]\right]. && \text{(A–8)} \end{aligned}$$

where $\operatorname{erf}(t_0) \equiv \frac{2}{\sqrt{\pi}}\int_0^{t_0}\exp\left[-t^2\right]dt$. The power on the right side is:

$$P_{right} = \frac{P_0}{2}\left[1-\operatorname{erf}\left[\frac{\sqrt{2}x_0}{w}\right]\right]. \qquad \text{(A–9)}$$

We create a normalized yaw as

$$\text{YAW} = \frac{P_{left}-P_{right}}{P_0} = \operatorname{erf}\left[\frac{\sqrt{2}x_0}{w}\right]. \qquad \text{(A–10)}$$

Using the Taylor series expansion of the error function, we have a first order estimate for the relationship between normalized yaw and beam displacement x_0 for a beam of radius w:

$$\frac{x_0}{\text{YAW}} \approx \frac{w}{2}\sqrt{\frac{\pi}{2}}. \qquad \text{(A–11)}$$

The same equation holds true for pitch.

A.4 Beam Propagation Formalism

For the input beam model and for the Input Optics beam drift calibrations, the ABCD matrix formalism is a useful tool to propagate a Gaussian beam. I choose to ignore the fact that the MC beam passes through the substrate of MC3 on its way to the Faraday. I also treat the beam splitter as a flat mirror and ignore the presence of its substrate. I use the thickness of the large optic substrates, $t = 0.01$ m, and account for index of refraction effects when passing through optics.

Signs of radii of curvature are defined per the front face of the optic; for example, all main LIGO optics have a positive R.

For a beam that strikes a flat interface and exits at a curved interface (ie. forward-going transmission through RM):

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ (n_2 - n_1)/Rn_1 & n_2/n_1 \end{bmatrix} \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & n_1/n_2 \end{bmatrix} \tag{A–12}$$

For a beam that strikes a curved interface and exits at a flat interface (ie. transmission through ETM):

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & n_2/n_1 \end{bmatrix} \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ (n_1 - n_2)/-Rn_2 & n_1/n_2 \end{bmatrix} \tag{A–13}$$

For a beam that strikes a flat interface, travels through the substrate, reflects off the back of a curved interface, travels through substrate and exits at the original flat interface (ie. single bounce off RM):

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & n_2/n_1 \end{bmatrix} \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 2/R & 1 \end{bmatrix} \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & n_1/n_2 \end{bmatrix} \tag{A–14}$$

Finally, for prompt reflection off a curved interface (ie. reflection off MMTs):

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -\frac{2}{R} & 1 \end{bmatrix} \tag{A–15}$$

and for propagation a distance d through vacuum:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix}. \tag{A–16}$$

For each of these $n_1 = 1$ is the index of refraction of vacuum and $n_2 = 1.44963$ is the index of refraction of the fused silica used for the optics. Table A-1 shows the radii of curvature of each of the optics for both sites.

Table A-1. Mirror radii of curvatures.

optic	L1 ROC [km]	H1 ROC [km]
MC1	0.01725	
MC2	80	
MC3	0.01725	
MMT1	6.76	6.77
MMT2	3.16	3.17
MMT3	25.16	25.04
RM	15.78	14.40
BS	-189	-336
ITMX	14.760	13.910
ITMY	14.520	13.600
ETMX	8.730	7.260
ETMY	8.720	7.320

A.5 Beam Drift Calibration

We use the ABCD matrix formulation to convert pitch and yaw data of WFS3 and WFS4 into a position and angle at the Faraday isolator. The basic relationship between beam displacement and angle at one location to displacement and angle at another location is given by:

$$\begin{bmatrix} x_1 \\ x_1' \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} x_0 \\ x_0' \end{bmatrix} \tag{A–17}$$

For this application, we want to relate the beam positions on the WFS, x_3 and x_4, to the beam position and angle, x_{FI} and x'_{FI}, at the Faraday isolator. Using only the top equation of Eq. A–17 since the WFS are sensitive to beam position only and not angle, we can write a new relation

$$\begin{bmatrix} x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} A_3 & B_3 \\ A_4 & B_4 \end{bmatrix} \begin{bmatrix} x_{FI} \\ x'_{FI} \end{bmatrix} \tag{A–18}$$

where A_3, B_3 and A_4, B_4 are the A and B ABCD matrix elements for the beam paths from the Faraday isolator to WFS3 and WFS4, respectively. Taking the inverse and writing x_3 and x_4 as a function of the pitch and yaw recorded by the WFS (see Appendix A.3), the useful equation is

$$\begin{bmatrix} x_{FI} \\ x'_{FI} \end{bmatrix} = \begin{bmatrix} A_3 & B_3 \\ A_4 & B_4 \end{bmatrix}^{-1} \begin{bmatrix} w_3 & 0 \\ 0 & w_4 \end{bmatrix} \frac{1}{2}\sqrt{\frac{\pi}{2}} \begin{bmatrix} DOF_3 \\ DOF_4 \end{bmatrix} \tag{A–19}$$

where w_3 and w_4 are the radii of the beam at each WFS and DOF can mean PIT or YAW.

A.6 Carrier Mode-matching into the Interferometer

When a cavity is locked to an input beam, all of the light impinging the cavity is coupled into it if the cavity is impedance matched and if the input beam and the cavity are mode-matched. When these requirements are not met, then light is reflected from the cavity. The cavity visibility is a quantity that summarizes the compound effect of these sources of reflected light.[1] By measuring the interferometer visibility and by the measuring the impedance mismatch, we can deduce the carrier mode-matching into the interferometer.

A.6.1 Interferometer Visibility

The visibility is a measure of how much carrier light is reflected from the locked interferometer compared to how much carrier light is sent in. To measure the visibility, we need to know only a few numbers. We must have a measure of how much light is sent to the interferometer for normalization purposes and we must have a measure of the DC reflected power when the interferometer is both locked and not locked (all light is reflected off of the RM). The visibility is then given by:

$$\text{visibility} = 1 - \frac{P_{REFL_{locked}}}{P_{REFL_{unlocked}}} \frac{P_{IN_{unlocked}}}{P_{IN_{locked}}} \tag{A–20}$$

We have two measures of how much light is being sent into the interferometer (a pick-off of the light before it enters the vacuum and a pick-off of the MC transmitted light) and several of how much light is reflected. An example showing some of these signals for LLO is shown in Figure A-2. The lock stretch ends at $t = 5$ min. Note that the amount of reflected light is increasing up to the end of the lock as the interferometer is losing stability. When lock is lost, the common mode servo kicks the mode cleaner out of lock too, and the MC trans power drops to 0. About 15 seconds later the MC relocks and then the power into it increases. About 15 seconds

[1] When there are sidebands, as is the case for LIGO, the sidebands are also a source of reflected light. We measured that 6% to 8% of the power in the reflected beam during lock is sideband content, and that the 25 MHz sideband visibility is 87%.

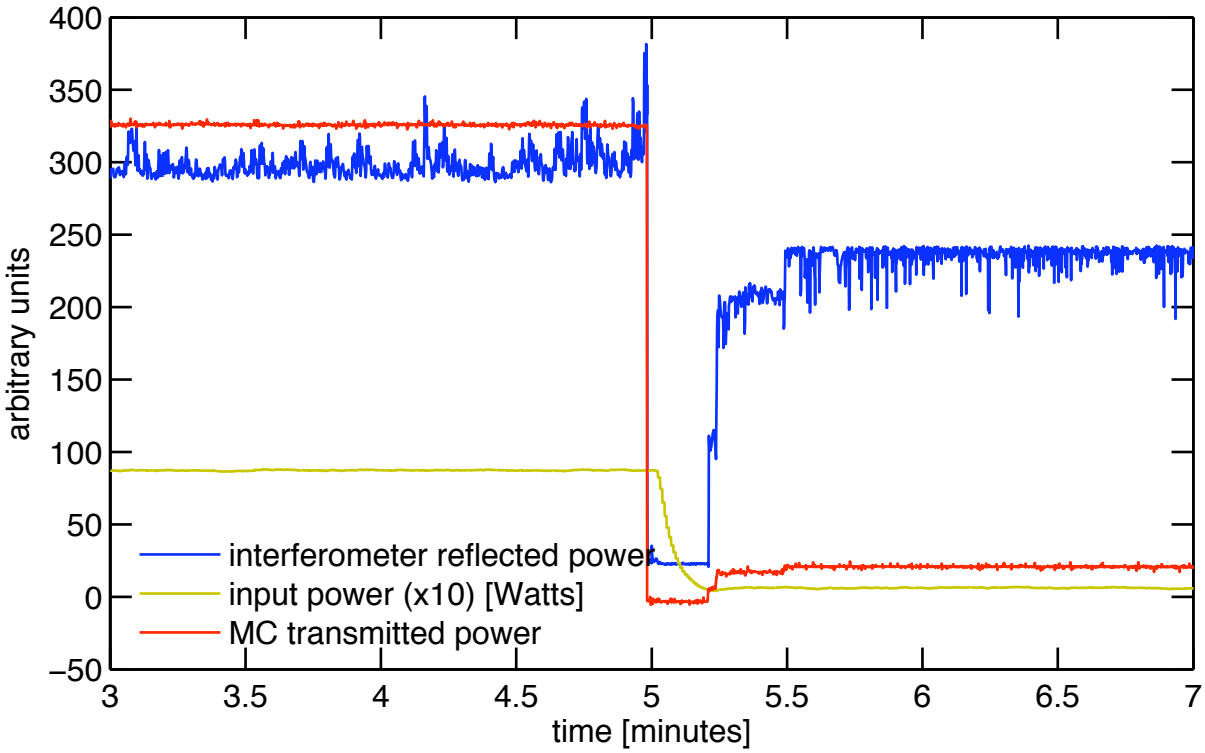

Figure A-2. End of an 8.7 W lock at Livingston on Feb. 23, 2010. The mode cleaner re-locks at 0.5 W about 15 seconds after lock loss and then the power is increased to 2 W. Comparison of the minimum intereferometer reflected power during lock and the maximum reflected power out of lock provides a measure of interferometer mode matching.

after that, at 5.5 minutes, the MC WFS turn on, improving the alignment of the MC to the input beam and we see another step in the power getting through the MC. The interferometer is still not locked, so all light (except for 2.7%) is reflected off of the recycling mirror. The downward spikes in the reflected light trace are the result of interferometer flashes, instances of all mirrors lining up correctly to let some light in.

For this particular example, the visibility is 92.1%. When evaluated for a sampling of lock losses throughout the Enhanced LIGO run, the average visibility is 91.84% $\pm$ 0.07%.

A.6.2 Impedance Matching

A cavity is impedance matched when the input and output couplers have the same reflectivity. If there is a difference between the reflectivities of the two mirrors, the cavity is over- or under-coupled and light will be reflected. Treating the interferometer arms as a single mirror

that forms a cavity with the RM, we want the RM transmission to match the transmission of the arms. This must include all losses in the interferometer such as absorption, scattering, and ETM transmission. Design estimates resulted in an RM power transmission of 2.7%. If losses do not equal 2.7%, then there is an impedance mismatch and we will see light at the reflected port.

The amplitude reflectivity of the interferometer is:

$$r_{ifo} = \frac{r_{arms} - r_{rm}}{1 - r_{rm} r_{arms}}. \tag{A–21}$$

The composite arm cavity amplitude reflectivity is r_{arms} and the RM amplitude reflectivity is $r_{rm} = \sqrt{0.973}$. It is not so simple to know what r_{arms} is in practice. A precise measure of all losses in the arms would be needed. Therefore, we turn to writing r_{arms} in terms of a quantity that we can measure, the power recycled Michelson carrier gain G_{cr}:

$$G_{cr} = g_{cr}^2 = \left[\frac{t_{rm}}{1 - r_{rm} r_{arms}} \right]^2. \tag{A–22}$$

Experimentally, the recycling gain is measured as

$$G_{cr} = T_{rm} \frac{\mathrm{NPTRX} + \mathrm{NPTRY}}{2}, \tag{A–23}$$

where NPTRX and NPTRY are channels recording the amount of light transmitted through the ETMs, normalized such that NPTRX=NPTRY=1 during a single arm lock. For Livingston, $G_{cr} = 39$.

Figure A-3 shows $R_{ifo} = r_{ifo}^2$ as a function of G_{cr}. Curves for a couple different RM reflectivities are shown to give an idea of how the interferometer reflectivity would change for minor mis-approximations of the RM reflectivity. We find that the impedance mismatch for Livingston is only 0.07%.

A.6.3 Mode-matching

Any difference between the interferometer visibility and what is expected from impedance mismatch can be contributed to imperfect mode matching. For Livingston, the interferometer is

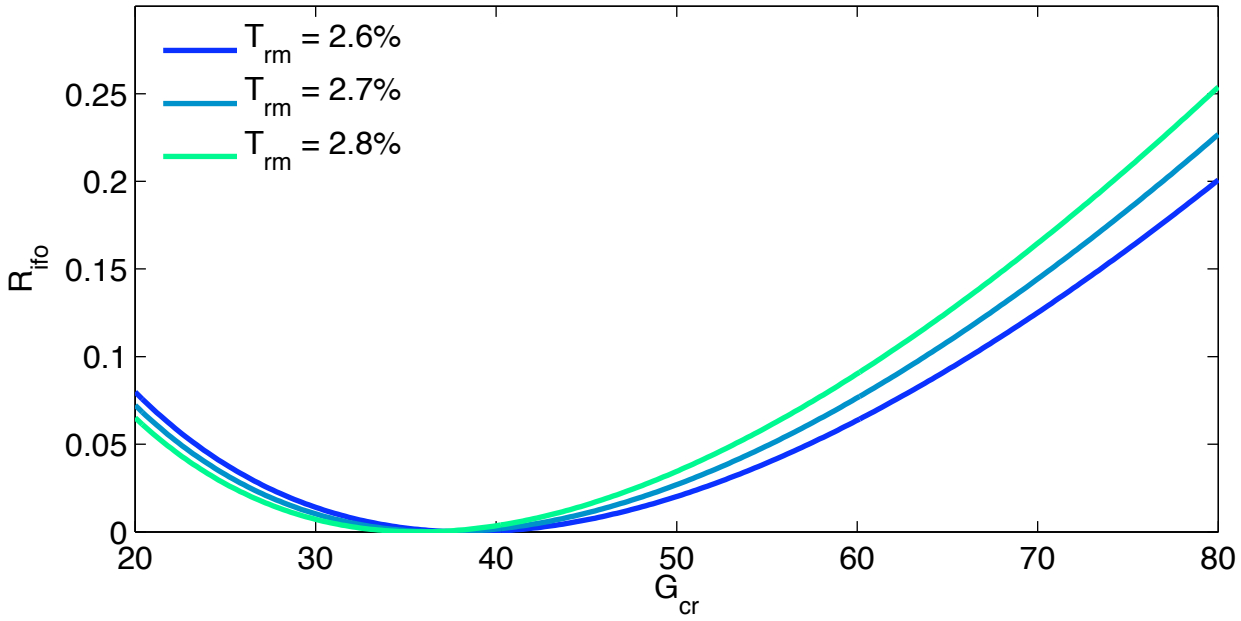

Figure A-3. Interferometer reflectivity due to impedance mismatch. The percentage of power incident on the RM that is reflected by the interferometer is a function of carrier recycling gain and RM reflectivity. The carrier recycling gain is 39 for Livingston. The recycling mirror power transmission is nominally 2.7%.

nearly perfectly impedance matched, so all light at the reflected port is due to imperfect mode matching. Therefore, the LLO mode mismatch during S6 was 8%.

A.7 Overlap Integrals

A measure of the mode matching can be given by the amount of power coupled from one mode into another. This is calculated as the square of the overlap integral of two fields, ψ_1 and ψ_2, for a particular z-axis (propagation direction) cross-section:

$$P = \left|\langle \psi \mid \psi' \rangle\right|^2 = \left[\int_{-\infty}^{\infty} \psi^\star(x)\,\psi'(x)\,dx\right]^2 \tag{A–24}$$

We are interested in the lowest order Hermite-Gaussian mode:

$$\psi(x,z) = u_0(x,z) = \left[\frac{2}{\pi w_0^2}\right]^{1/4}\left[\frac{q_0}{q(z)}\right]^{1/2}\exp\left[\frac{-ikx^2}{2q(z)}\right] \tag{A–25}$$

which can be rewritten as a function of x and q using $q_0 = i\mathrm{Im}(q)$, and $w_0^2 = -2q_0 i/k$:

$$u_0(x,q) = \left[\frac{-k\mathrm{Im}(q)}{\pi q^2}\right]^{1/4}\exp\left[\frac{-ikx^2}{2q}\right] \tag{A–26}$$

Eq. (A–26) is normalized such that $\langle u_0 \mid u_0 \rangle = 1$.

We want to know the square of the overlap integral for two fields given by different q parameters at one location z. First, the overlap integral:

$$\langle q_1 \mid q_2 \rangle = \int_{-\infty}^{\infty} u_0(x,q_1) u_0(x,q_2)\, dx \tag{A–27}$$

$$= \left[\frac{k^2 \mathrm{Im}(q_1^\star)\mathrm{Im}(q_2)}{\pi^2 q_1^{\star 2} q_2^2}\right]^{1/4} \int_{-\infty}^{\infty} \exp\left[-\left[\frac{-ik}{2q_1^\star} + \frac{ik}{2q_2}\right] x^2\right] dx \tag{A–28}$$

$$= \left[-\frac{k^2 \mathrm{Im}(q_1)\mathrm{Im}(q_2)}{\pi^2 q_1^{\star 2} q_2^2}\right]^{1/4} \sqrt{\frac{\pi}{\left[\frac{-ik}{2q_1^\star} + \frac{ik}{2q_2}\right]}} \tag{A–29}$$

$$= [\mathrm{Im}(q_1)\mathrm{Im}(q_2)]^{1/4} \sqrt{\frac{2}{q_1^\star q_2 \left(1/q_2 - 1/q_1^\star\right)}} \tag{A–30}$$

Then, the power is given by:

$$|\langle q_1 \mid q_2 \rangle|^2 = \frac{2\sqrt{\mathrm{Im}(q_1)\mathrm{Im}(q_2)}}{\sqrt{q_1^\star q_1 q_2^\star q_2 \left[\frac{1}{q_2} - \frac{1}{q_1^\star}\right]\left[\frac{1}{q_2^\star} - \frac{1}{q_1}\right]}} \tag{A–31}$$

$$= \frac{2\sqrt{\mathrm{Im}(q_1)\mathrm{Im}(q_2)}}{|q_2 - q_1^\star|} \tag{A–32}$$

Note that Eq. (A–32) simplifies to 1 when $q_1 = q_2$ as expected and that this whole formulation assumes that the beams are propagating along the same z-axis.

APPENDIX B
ANGULAR SENSING AND CONTROL CALIBRATIONS

The typical method of calibrating a digital channel is to inject a signal of known amplitude into the system and take the ratio with the amplitude of the digital measurement of the signal. I describe in this appendix the calibrations I made of some of the angular sensing channels.

B.1 Beam Spot Motion

A quantity of interest is how much the beam moves on the ITMs and ETMs. It is this beam spot motion which, together with the mirror angular motion, creates a length signal that contributes noise to DARM. An elegant way of following the motion of the beam on the test masses is to track pickoffs of the light transmitted or reflected from the mirrors. We have such signals naturally available for the ETMs and ITMs from the QPDs which are otherwise used for ASC sensing. For example, QPDX and QPDY see the light transmitted through each of the ETMs and WFS2 sees the pickoff of light from the wedge of ITMX.

To calibrate the counts of the QPD and WFS2 pitch and yaw error signals,[1] I moved the beam a known distance on the test mass, Δx, and recorded the corresponding Δy of the QPD and WFS2 readback. The ratio $\Delta x/\Delta y$ is the calibration from counts to meters. The details of the procedure are described below.

B.1.1 Moving the Beam

Moving the beam on the mirrors in a controlled fashion is straightforward because of the ASC system. All that we need to do is introduce an offset to the setpoint of the of the beam centering aspect of the ASC servo. For the ETMs we put a DC offset in the `L1:ASC-QPD{X,Y}_{PIT,YAW}_{OFFSET}` channel and for the ITMs we changed the X and Y targets of the beam splitter beam centering servo.

[1] `L1:ASC-QPDY_{PIT,YAW}_IN1` and `L1:ASC-WFS2_DC{Pitch, Yaw}Mon`

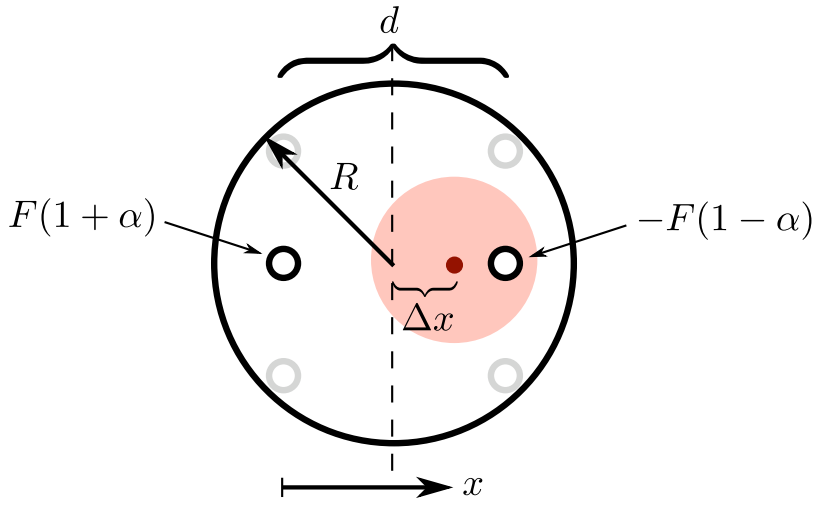

Figure B-1. Geometry of OSEMs and mirror as used for calculating the location of the axis of rotation when the torques are unequal.

B.1.2 Measuring How Much the Beam Has Moved

The more difficult task is measuring just how much the beam has moved. For this, we make use of the lever arm mechanism of angle to length coupling as explained in Section 4.1. The concept of the measurement is to move the axis of rotation of the mirror so that it passes through the center of the beam. We use the OSEMs to change the location of the axis of rotation, and we use DARM to determine when the axis is aligned with the beam center. For example, if we drive the top two OSEMs more than the bottom two OSEMs, we've created an axis of rotation that sits below the center of mass. The result of such tuning is an effective rebalance of the center of mass of the mirror so that it is aligned with the center of the beam. The procedure is:

- Shake the mirror at some frequency f (we use 39.5 Hz) during a full lock
- Demodulate DARM at f for several different sets of OSEM gains
- Fit a quadratic to the demodulated data to pinpoint the OSEM gains that minimize the coupling to DARM

Relating the OSEM gains to absolute beam position on the mirror requires only the geometry of the mirror and OSEM setup as sketched in Figure B-1. We estimate the OSEM locations as being on the edge of the mirror such that the length d of one side of the square that they form is given by $d = \sqrt{2}R$, where $R = 12.5$ cm is the radius of the mirror. Then, collapsing the four OSEMs into a representative two at the centers of two opposite sides of the square and assigning

Table B-1. Calibrations to be used with the QPDX, QPDY, and WFS2 DC pitch and yaw error signals for a measure of beam spot motion.

	ETMX	ETMY	ITMs
pitch	1.03×10^{-5} m/ct	1.21×10^{-5} m/ct	5.52×10^{-2} m/ct
yaw	0.88×10^{-5} m/ct	0.80×10^{-5} m/ct	4.79×10^{-2} m/ct

them gains of $1+\alpha$ and $-(1-\alpha)$ for a force F, we can evaluate where the pivot point x is located by setting the sum of the torques equal to zero:

$$F[1+\alpha]x = F[1-\alpha][d-x]. \tag{B–1}$$

Therefore, the beam location relative to center, Δx, is

$$\Delta x := \frac{d}{2} - x = \alpha\frac{d}{2}, \tag{B–2}$$

and for a change in a pitch or yaw coil gain, the change in beam position, Δx, is:

$$|\Delta x| = \frac{|\Delta\text{gain}|R}{\sqrt{2}}. \tag{B–3}$$

The final calibrations of these channels are shown in Table B-1. [2]

B.2 Angular Mirror Motion

The optical levers provide a straightfoward measure of individual mirror motion. The channels I calibrated were of the form `L1:SUS-ETMX_OPLEV_{P,Y}ERROR`, the optical lever error signals for each of the large optics. I made use of the dependence of power in a misaligned cavity (refer to Appendix C.3) to calibrate the ETM and ITM optical levers, and used a less precise, rudimentary method to calibrate the RM, BS, and MMT3 optical levers.

[2] A minor technicality is that since there are no filters between the QPD error signals and the offset channel, their units are exactly the same. Thus, calculating meters of beam spot motion as a function of offset serves to calibrate the error point. For convenience, this is what I did.

B.2.1 ETM and ITM Optical Levers

I calibrated the arm cavity optical levers by tracking the power loss in the locked arm as one of its mirrors is tilted. The closed form expression for cavity power as a function of mirror tilt is derived in Appendix C.3. All that is needed is a quadratic fit to the data collected. From the fit parameters, I can determine the factor, $\Delta\theta/\Delta y$, which converts the digital counts of the optical lever channel, y, to units of radians.

To make the measurement, I locked a single arm and maximized the power build up. Then I slowly stepped the pitch or yaw pointing of one of the mirrors away to one side of resonance, and then back and to the other side, repeating this several times. All the while, I recorded the optical lever error signal of the mirror whose angle I was changing, and the power in the arm as determined from the amount of light transmitted through the ETMs.[3]

From Eq. C–31, we see that the power in the arm, P, is a function of the form

$$P = P_{max} \exp\left[-b(y-y_0)^2\right], \tag{B–4}$$

where y_0 is the DC offset of the optical lever channel and b is related to physical cavity axis displacement a and tilt α by $by^2 = (a/w_0)^2 + (\alpha/\theta_0)^2$. In order to relate the optical lever signal, y, to physical cavity parameters, we divide by $\Delta\theta^2$ and rearrange to get:

$$\frac{\Delta\theta}{\Delta y} = \sqrt{b}\left[\left[\frac{\Delta a/\Delta\theta}{w_0}\right]^2 + \left[\frac{\Delta\alpha/\Delta\theta}{\theta_0}\right]^2\right]^{-\frac{1}{2}}. \tag{B–5}$$

The terms in the numerators on the right hand side are fixed constants based on the cavity geometry and can be calculated using Eq. C–4. The measurement data and fits are shown for both pitch and yaw in Figure B-2. The ETM optical levers make use of a broader range of optical lever signal than do the ITMs. (Also note that the maximum power in the y-arm is about 10% less than

[3] `L1:LSC-NPTR{X,Y}_OUT16`

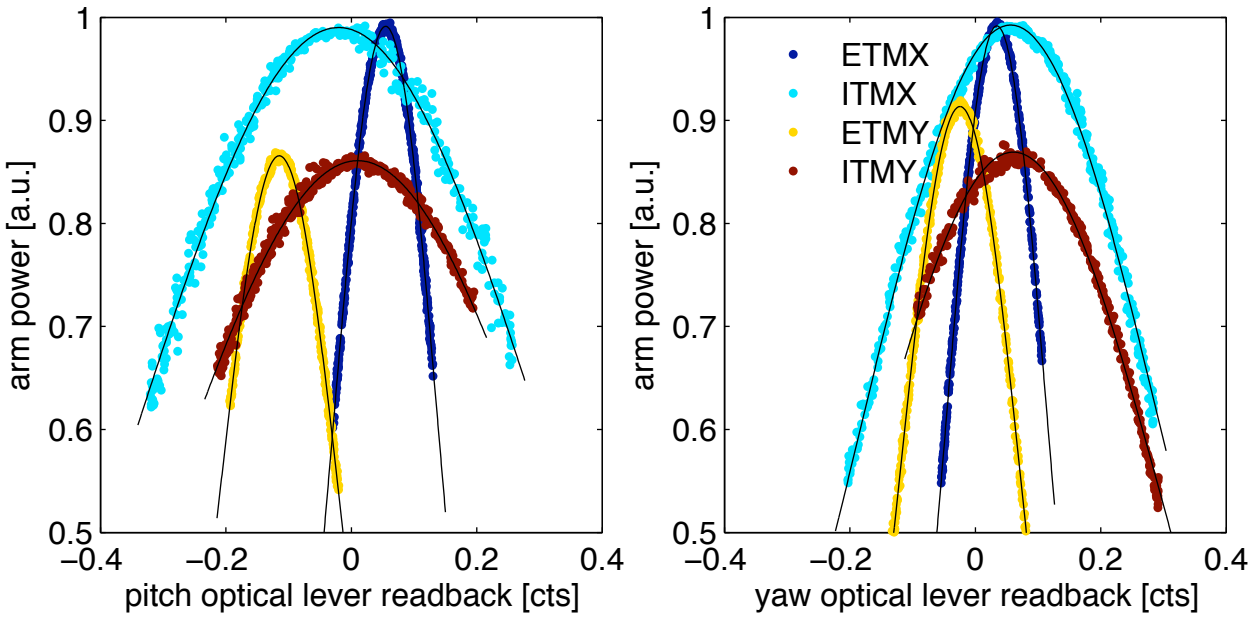

Figure B-2. Optical lever calibration data and fits to Eq. B–4.

that in the x-arm. This is true at both Hanford and Livingston, and is due to the priority given to the x-arm in the alignment scheme, as explained in Appendix C.4.)

B.2.2 RM, BS, and MMT3 Optical Levers

To calibrate the RM, BS and MMT3 optical levers, I used my own eyes and the camera images and known dimensions of the ETM beam cages. With the interferometer unlocked, I moved the optics in pitch and yaw, tracking the beam's movement on the ETM cages. In order to calibrate the RM, I used the reflection off ITMY to the RM and onto the ETMY cage. The BS and MMT3 required only straight shots to the ETMs. For yaw, I moved the mirrors until the beam was centered on each vertical suspension post, and for pitch I moved the beam from the center of the mirror to the top of the cage. The beam moves by $x = 2\theta$ on across the cage when the mirror moves by θ, so with small angle approximations, the mirror angle is simply $x/2L$ where L is the distance from mirror to ETM cage.

The final $\Delta\theta/\Delta x$ calibrations of all optical levers are in Table B-2.

Table B-2. Calibrations to be used with the optical lever error signals for a measure of angular mirror motion. Units are µrad/ct.

	ETMX	ETMY	ITMX	ITMY	RM	BS	MMT3
pitch	49.4	43.0	14.9	15.6	61.9	47.3	57.4
yaw	50.7	43.3	20.1	20.2	42.5	63.5	55.5

Table B-3. Demodulation chain calibration for each quadrant of each WFS. Units are µV/count.

	q1	q2	q3	q4	average
WFS1	0.35	0.32	0.34	0.35	0.34
WFS2	8.8	8.6	8.7	8.5	8.7
WFS3	6.4	5.8	5.8	5.7	5.9
WFS4	6.3	5.4	5.3	7.4	6.1

B.3 WFS Error Signals

The WFS error signals[4] are physically Watts of power at the detectors, which the WFS electronics convert into a voltage. To turn WFS counts into voltage of signal at the output of the detector, we must backtrack through the electronics and calibrate the WFS demodulation chain.

The analog to digital RF chain for the WFS includes a demodulation board, a whitening board, an anti-alias board, and the ADC. I calibrated this chain by injecting a sine wave of the same frequency as a typical WFS signal, yet of known voltage into the WFS demodulation board. Comparing the peak to peak voltage of this input sine wave to the peak to peak amplitude of the resulting digital counts signal provides the Volts per count conversion. The calibrations are presented in Table B-3.

It should be noted that the demodulation chain calibration numbers for all quadrants of a particular WFS differ no more than 20% from the average. The demodulation chain does not significantly distort the error signals.

B.4 Angular Optical Gain

The calibrated WFS sensing matrix (using the calibration presented in Sec. B.3) can be used in conjuction with the simultaneously measured drive matrix (as described in Sec. 6.1.3)

[4] i.e. `L1:ASC-WFS1_QP`

to calculate the angular optical gain. Each row of the sensing matrix must be divided by the measured amplitude of that row's DOF excitation, which is found as the diagonal elements of Table 6-4. The result of doing so gives the angular optical gain of the interferometer in terms of WFS Volts per radian.

APPENDIX C
ANGULAR SENSING AND CONTROL SUPPORTING MATERIAL

C.1 Optical Lever Open Loop Transfer Function

The optical lever open loop transfer function for ETMX is shown in Figure C-1. The measurement was made when the interferometer was unlocked and only the optical lever and OSEM damping loops on. Because of the two unity gain crossings, one at 2.2 Hz and one at 0.2 Hz, the optical lever loop provides only AC (velocity) damping. There is no DC control.

The model uses the pendulum torque to angle transfer function as the plant and the filters used during Enhanced LIGO as the control. I use the meausurement to tune the pedulum parameters in the model to best determine the actual damping coefficient and resonant frequency. For ETMX, we find the pitch resonant frequency to be 0.65 Hz, 8% different from the theoretical 0.6 Hz. Also, we measure the ETMX damping coefficient to be 0.02.

C.2 Misaligned Cavity Axis

Here I provide the geometric argument that shows how to calculate the tilt a and displacement α of a cavity as a function of mirror misalignment. Cavity tilt is defined by the angle formed between the line that connects the two beam spots (as given by Eq. 5–1) and the line joining the centers of the mirrors. Cavity displacement uses the same two lines, yet is defined by the distance between them at the location of the waist of the resonant spatial mode. Based on pure geometry, the cavity displacement and tilt are:

$$\begin{bmatrix} a \\ \alpha \end{bmatrix} = \frac{1}{L} \begin{bmatrix} z_2 & z_1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \tag{C–1}$$

where z_i is the distance to the waist from mirror i calculated as:

$$z_1 = \frac{g_2(1-g_1)L}{g_1+g_2-2g_1g_2} \tag{C–2}$$

$$z_2 = L - z_1. \tag{C–3}$$

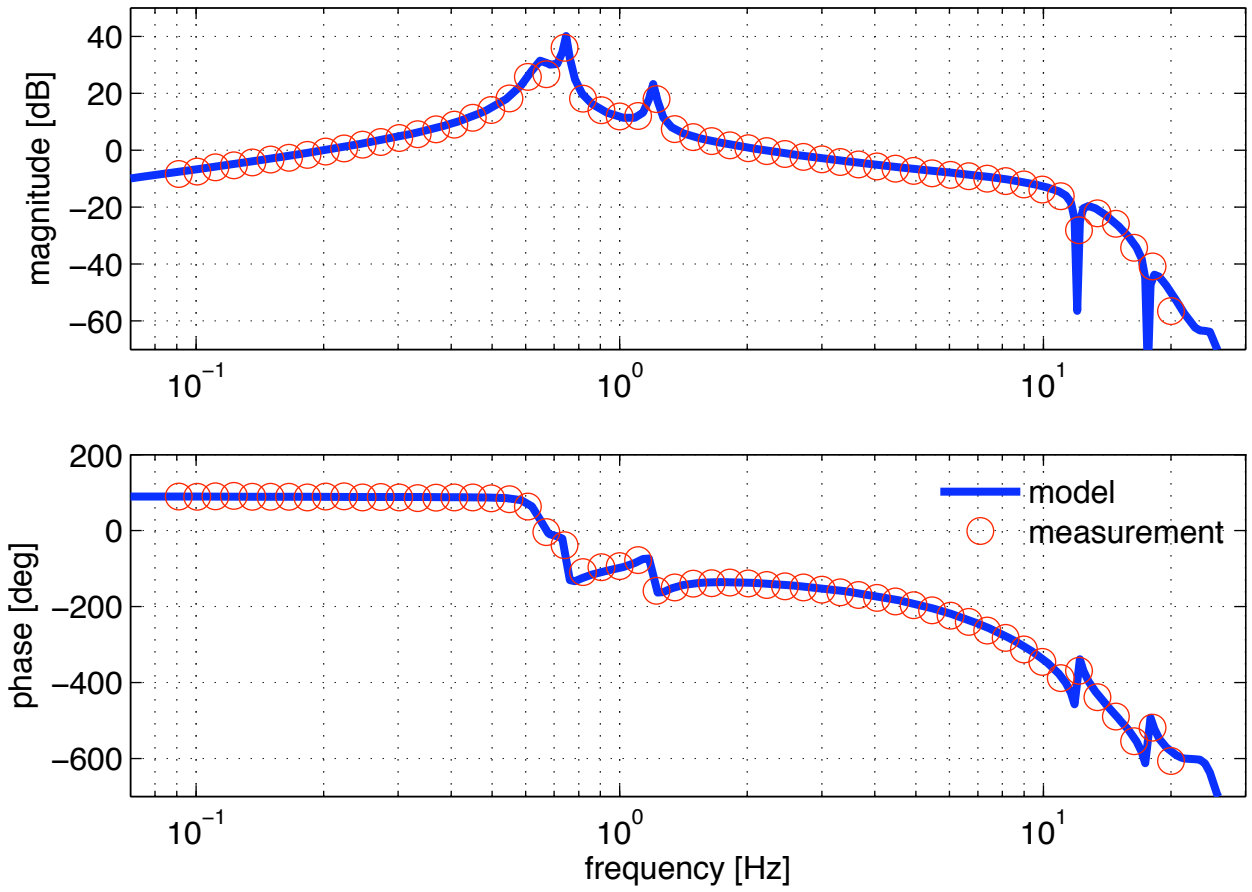

Figure C-1. ETMX pitch optical lever open loop transfer function. The model of the plant is tuned to match the data, resulting in a pitch resonance of 0.65 Hz and a damping factor of $\gamma = 0.02$. The UGF is at 2.2 Hz and the phase margin is 38°.

Clearly, we can combine Eqs. (5–1) and (C–1) to arrive at an equation directly relating mirror tilt to cavity displacement and tilt:

$$\begin{bmatrix} a \\ \alpha \end{bmatrix} = \frac{1}{1 - g_1 g_2} \begin{bmatrix} g_2 z_2 + z_1 & z_2 + g_1 z_1 \\ -g_2 + 1 & -1 + g_1 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}. \tag{C–4}$$

C.3 Power in a Misaligned Cavity

I'll show how to calculate the power in a cavity as a function of cavity axis displacement and tilt. Combined with the results of Eq. C–4 we determine how the power build-up in a cavity depends on a single mirror's angular displacement.

The field of a lowest-order Gaussian laser beam along one axis at the beam waist is:

$$\psi(x) = U_0(x) = \left[\frac{2}{\pi w_0^2}\right]^{1/4} \exp\left[-\left[\frac{x}{w_0}\right]^2\right] \tag{C–5}$$

where w_0 is the beam waist radius and U_0 is the lowest-order Hermite polynomial. The Hermite polynomials are orthonormal, ie. $\langle U_i \mid U_j \rangle = \delta_{ij}$. For example, the next to lowest order polynomial is:

$$U_1(x) = \left(\frac{2}{\pi w_0^2}\right)^{1/4} \frac{2x}{w_0} \exp\left[-(x/w_0)^2\right] = \frac{2x}{w_0} U_0(x) \tag{C–6}$$

C.3.1 Displaced Cavity

The field of a cavity with a *displaced* z-axis at the cavity waist is:

$$\begin{aligned} \psi\prime(x) =& \psi(x-a) && \text{(C–7)} \\ =& U_0(x-a) && \text{(C–8)} \\ =& c_0 U_0(x) + c_1 U_1(x) + c_2 U_2(x) + \ldots && \text{(C–9)} \end{aligned}$$

where a is the displacement of the axis and c_i are constants.

We want to know c_0, the projection of the displaced cavity field onto the beam field:

$$\begin{aligned} c_0 &= \langle \psi \mid \psi\prime \rangle && \text{(C–10)} \\ &= \int_{-\infty}^{\infty} \psi(x)\, \psi\prime(x)\, dx && \text{(C–11)} \\ &= \exp\left[-a^2/2w_0^2\right] && \text{(C–12)} \end{aligned}$$

The power in this mode is the square of the overlap of the two fields:

$$\begin{aligned} P_0 &= |\langle \psi \mid \psi\prime \rangle|^2 && \text{(C–13)} \\ &= \exp\left[-[a/w_0]^2\right] && \text{(C–14)} \end{aligned}$$

For the purpose of wavefront sensing, we need to know the amplitude, c_1, of the first order U_1 field. This can be approximated as demonstrated in Anderson [62] using the Taylor series expansion of the exponential in $\psi\prime(x) = U_0(x-a)$, assuming a displacement a that's small

compared to waist size w_0.

$$\begin{aligned}
\psi\prime(x) &= \left[\frac{2}{\pi w_0^2}\right]^{1/4} \exp\left[-\left[\frac{x-a}{w_0}\right]^2\right] && \text{(C–15)}\\
&= \left[\frac{2}{\pi w_0^2}\right]^{1/4} \left[1-\left[\frac{x-a}{w_0}\right]^2+\mathscr{O}(a^4)\right] && \text{(C–16)}\\
&= \left[\frac{2}{\pi w_0^2}\right]^{1/4} \left[\frac{2xa}{w_0^2}\left[1-\frac{x^2}{w_0^2}+...\right]+\left[1-\frac{x^2}{w_0^2}+\frac{1}{2}\frac{x^4}{w_0^4}-...\right]+\mathscr{O}(a^2)\right] && \text{(C–17)}\\
&= \left[\frac{2}{\pi w_0^2}\right]^{1/4} \left[1+\frac{2xa}{w_0^2}+\mathscr{O}(a^2)\right]\exp\left[-\left[\frac{x}{w_0}\right]^2\right] && \text{(C–18)}\\
&= U_0(x)+\frac{a}{w_0}U_1(x)+... && \text{(C–19)}
\end{aligned}$$

Notice that here we find $c_0 = 1$, which is consistent with the exact result of Eq. C–12 when we apply our $a^2 \approx 0$ approximation. We find that the amplitude of the first order Hermite-Gauss field for a displaced cavity is

$$c_1 = a/w_0. \tag{C–20}$$

C.3.2 Tilted Cavity

The field of a cavity with a *tilted* z-axis at the cavity waist is a tad more complex to derive. We assume the tilt, α, is small such that $\sin\alpha \approx \alpha$ and $\cos\alpha \approx 1$. Also, we assume the beam divergence angle, $\theta_0 = \lambda/\pi w_0$, is small such that the wavefronts near the waist can be considered parallel to one another.

Here, the important quantity to consider is the phase of the cavity field at the cross-section of the beam waist. The phase is either advanced or retarded compared to that of the beam:

$$\begin{aligned}
\psi\prime(x) &= \psi(x\prime)\exp[-ikz\prime] && \text{(C–21)}\\
&\approx \psi(x\cos\alpha)\exp[-ikx\sin\alpha] && \text{(C–22)}\\
&\approx \psi(x)\exp[-ikx\alpha] && \text{(C–23)}\\
&= U_0(x)\exp[-ikx\alpha] && \text{(C–24)}
\end{aligned}$$

where $k = 2\pi/\lambda$ and λ is the wavelength of the laser light.

The overlap of the fields of the beam and tilted cavity is $\exp[-\alpha^2/2\theta_0^2]$. Therefore the power is:

$$P_0 = \exp[-(\alpha/\theta_0)^2]. \tag{C–25}$$

An expansion of the exponential in Eq. C–24 for a small tilt α gives:

$$\begin{aligned} \psi\prime(x) &= U_0(x)[1 + ikx\alpha + \mathcal{O}(\alpha^2)] && \text{(C–26)} \\ &= U_0(x) + \frac{ik\alpha w_0}{2} U_1(x) + \mathcal{O}(\alpha^2). && \text{(C–27)} \end{aligned}$$

Therefore, the amplitude of the first order Hermite-Gauss field for a tilted cavity is

$$c_1 = ik\alpha w_0/2. \tag{C–28}$$

C.3.3 Displaced and Tilted Cavity

The most general case, of course, is to have a cavity axis that is both displaced *and* tilted at the beam waist:

$$\psi\prime(x) = \psi(x-a)\exp[-ik(x-a)\alpha]. \tag{C–29}$$

We find:

$$\langle \psi \mid \psi\prime \rangle = \exp\left[-\frac{a^2}{2w_0^2}\right] \exp\left[-\frac{\alpha^2}{2\theta_0^2}\right] \exp\left[-\frac{ia\alpha}{x_0\theta_0}\right] \tag{C–30}$$

and

$$P_0 = \exp\left[-\frac{a^2}{w_0^2}\right] \exp\left[-\frac{\alpha^2}{\theta_0^2}\right]. \tag{C–31}$$

C.4 Initial DC Alignment of the Interferometer

After any kind of in-vacuum work, the DC alignment of the mirrors is usually too poor for the interferometer to lock. A bootstrapping process of tweaking the alignment by hand is necessary, assuming the mirrors start out pointing in generally the right direction, as is usually the case. As pointed out in Sec. 4.4, the QPDs at the end stations are the fixed reference points for the overall alignment, so this process begins with making sure the light reaches them. We then adjust the rest of the mirrors to maximize power build-up in the arms and to maximze spatial overlap of the light reflected from each arm.

An outline of the process is presented here. “Misalign” means to intentionally point a mirror so far away from any known good positions as to eliminate it from the configuration. “Align” and “restore” mean to bring a mirror or configuration to the best known position(s). Centering the beam on a mirror is accomplished by using the suspension cage surrounding the mirror as a reference. Camera images and QPD readback provide the signals used for beam centering.

X-arm

- restore the x-arm (misalign RM, ITMY, and ETMY, align ITMX and ETMX)
- use ITMX to center the beam on QPDX
- use ETMX to center the beam on ITMX
- with x-arm locked, use MMT3 to maximize the x-arm power build-up (NPTRX, can expect about 95%)
- save the MMT3, ITMX, and ETMX alignment settings

Y-arm

- restore the y-arm (misalign ITMX and ETMX, align ITMY and ETMY)
- use ITMY to center the beam on QPDY
- use ETMY to center the beam on ITMY
- with y-arm locked, use BS to maximize the y-arm power build-up (NPTRY, can expect about 90%)
- save the BS, ITMY, and ETMY alignment settings

Relative x-arm and y-arm

- note AS beam position on camera while toggling between x-arm and y-arm locks
- use ETMs to align the two AS beams
- restore the Michelson (misalign ETMs, align ITMs)
- use BS to make AS port as dark as possible
- re-do y-arm alignment if ambitious

Recycling mirror

- restore the PRM (misalign ETMs, align ITMs and RM)
- use RM to center beam on ETMY cage

Restore full interferometer–off you go!

C.5 Photodiodes

The basic photoconductive photodiode is depicted in Figure C-2. A negative voltage applied to the anode, called the bias voltage V_{bias}, does not yield any current across the photodiode until electrons are released by the energy of photons striking the diode. Electrons flow in the direction

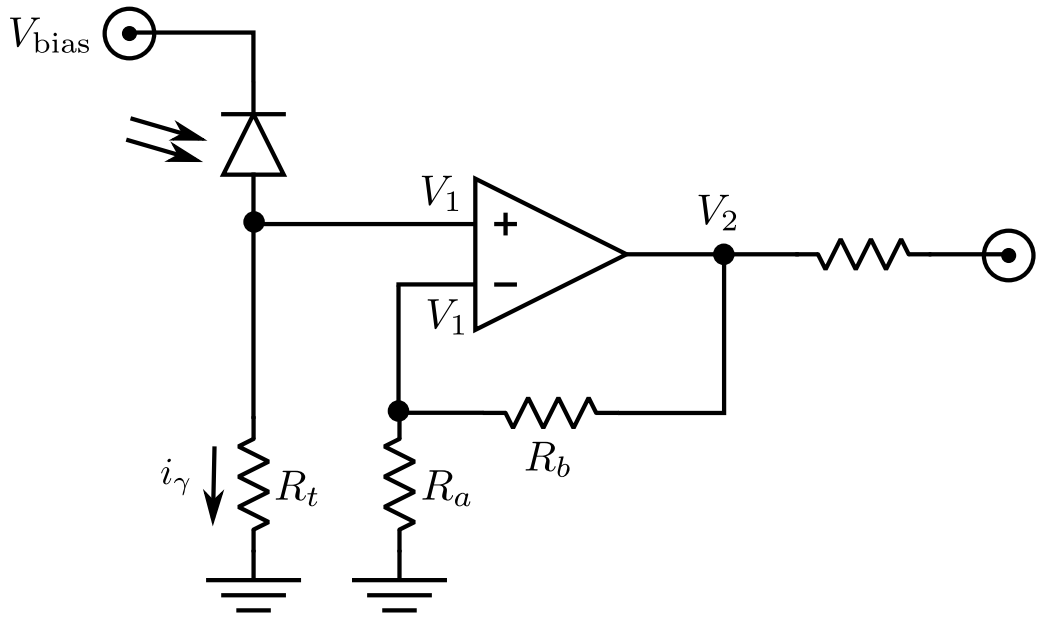

Figure C-2. Schematic of a basic photoconductive photodiode.

of the cathode which is connected to ground, and produce a current

$$i_\gamma = \frac{q_e}{h\nu} P\varepsilon \tag{C–32}$$

where q_e is the electron charge, h is Planck's constant, ν is the frequency of the incident light, P is the power of the light, and ε is the quantum efficiency of the diode. The quantity $q_e/h\nu$ is known as the *responsivity*. For LIGO where $\lambda = 1064$ nm, the responsivity is 0.86 A/W.

To have an accurate measure of the power on the photodiode, we must measure the produced current i_γ. This is accomplished by inserting a resistor in series between the cathode and ground and measuring the voltage drop across the resistor, whose resistance is called the *transimpedance*, R_t. A typical voltmeter, despite having high resistance, will inevitably draw some small amount of current, thereby underestimating the current produced by the light. The clever solution is to make use of an op-amp as a voltmeter. Op-amps have two primary properties:

- the inputs draw no current
- the voltages at the inputs to the op-amp are the same

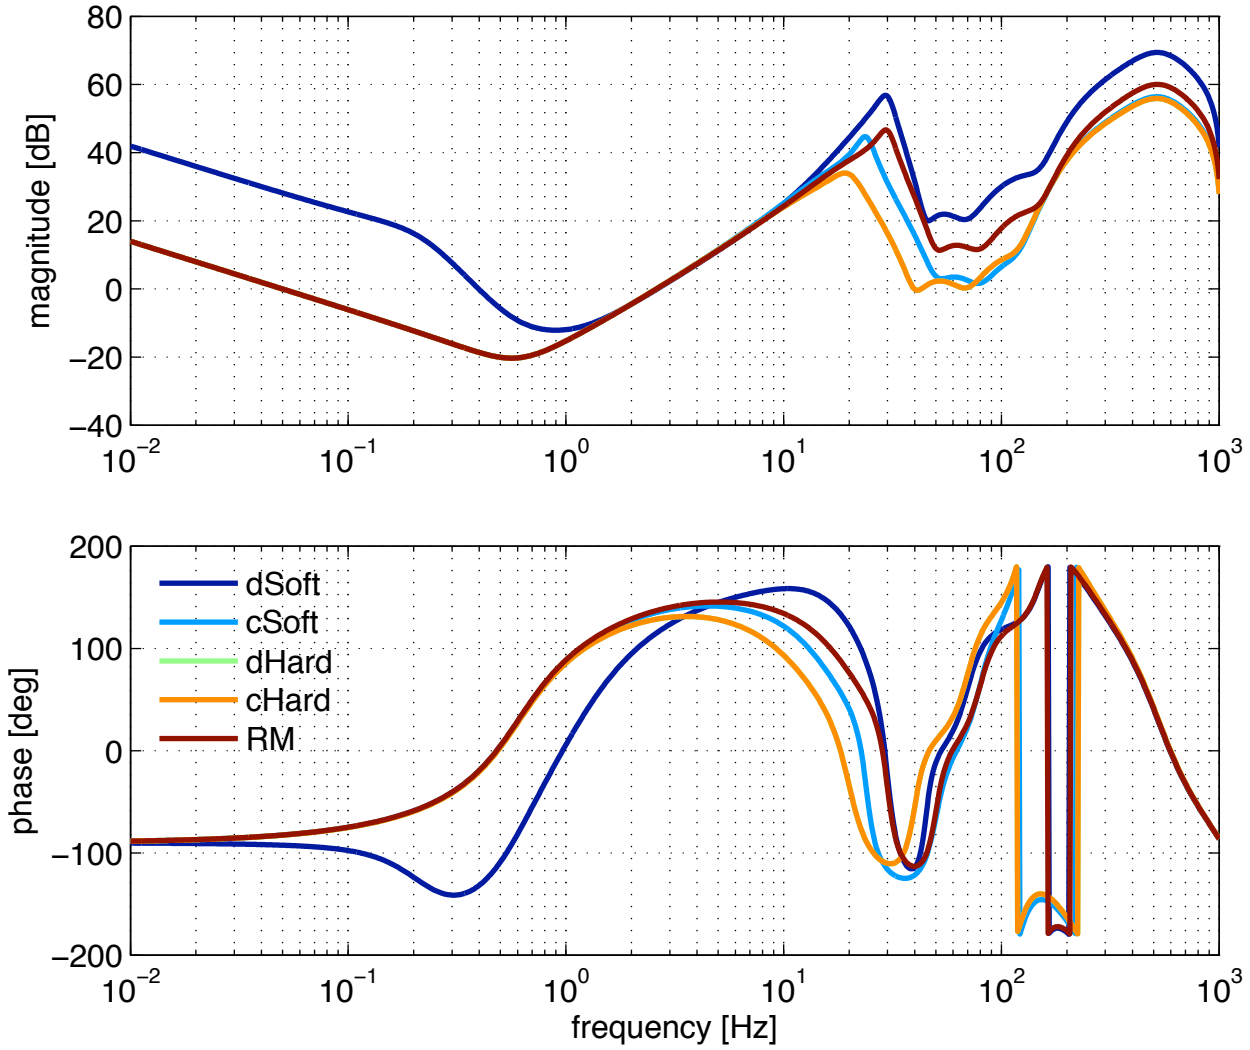

Figure C-3. WFS digital control filters.

The voltage at the input to the op-amp, V_1, is determined by measuring the voltage at the output of the op-amp, V_2, and accounting for the op-amp's gain such that:

$$V_1 = \frac{R_a}{R_a + R_b} V_2. \tag{C–33}$$

The photo current is therefore $i_\gamma = V_1 / R_t$.

C.6 WFS Control Filters

Figure C-3 is a record of the digitial filters used for the WFS control servo during Enhanced LIGO.

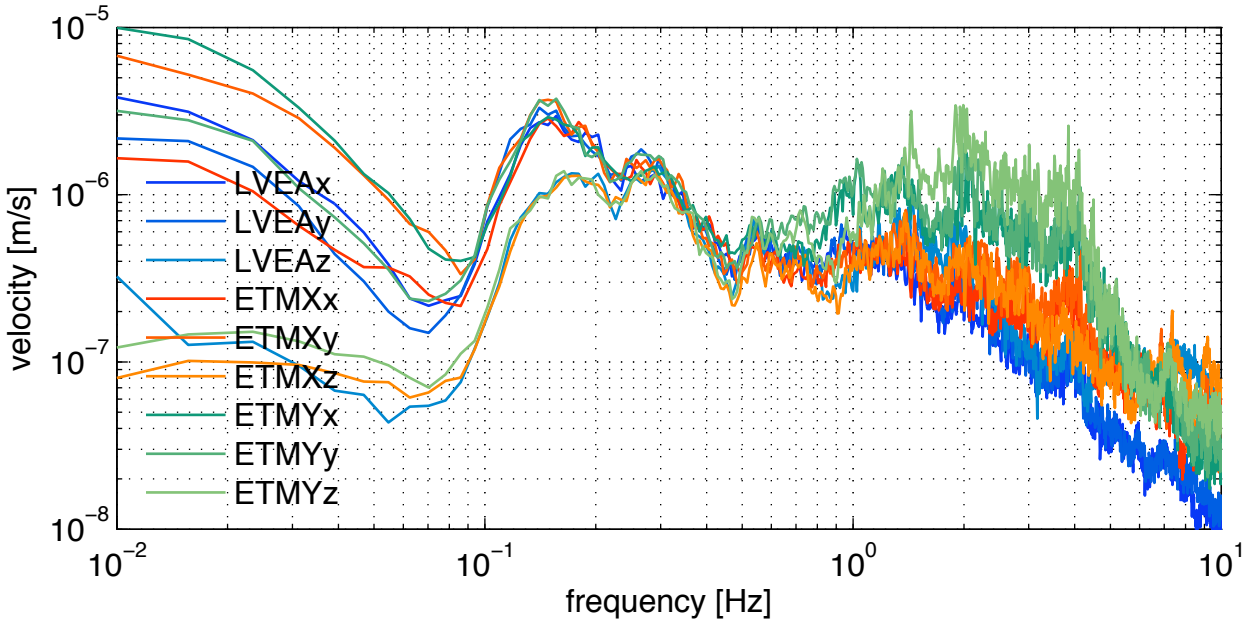

Figure C-4. Ground motion at time of optical lever spectra when the interferometer was unlocked (Figure 4-1) and at time of ASC suppression demonstration (Figure 6-8). GPS time is 956751915 (May 1, 2010 07:25 CDT).

C.7 Seismic Spectra

There are three seismometers at LLO for the purpose of monitoring the ground motion at the corner station (LVEA) and two end stations. The calibration for the digitally collected data is 2.4×10^{-9} m/s/count. As stated in the introduction of Ch. 6, I include here snapshots of the ground motion at the time of ground motion-sensitive measurements. Each spectra represents 30 minutes of data centered around the time of the measurement. There are three degrees of freedom for each seismometer, x, y, and z. They are aligned with the interferometer's x and y arm coordinate system.

Although the interferometer is very sensitive to ground motion, the range of seismic activity for which it can maintain full lock is reasonably large. The microseism (0.1-0.35 Hz) varies by factors of several seasonally (it's worse in the winter), and the 1-3 Hz motion varies by factors of several from day to night. Seismic motion the interferometer typically cannot handle includes events like earthquakes (0.03-0.1 Hz) and heavy activity on site (3-10 Hz). Motion at these frequencies are otherwise constant and at a level so as to not affect interferometer operation.

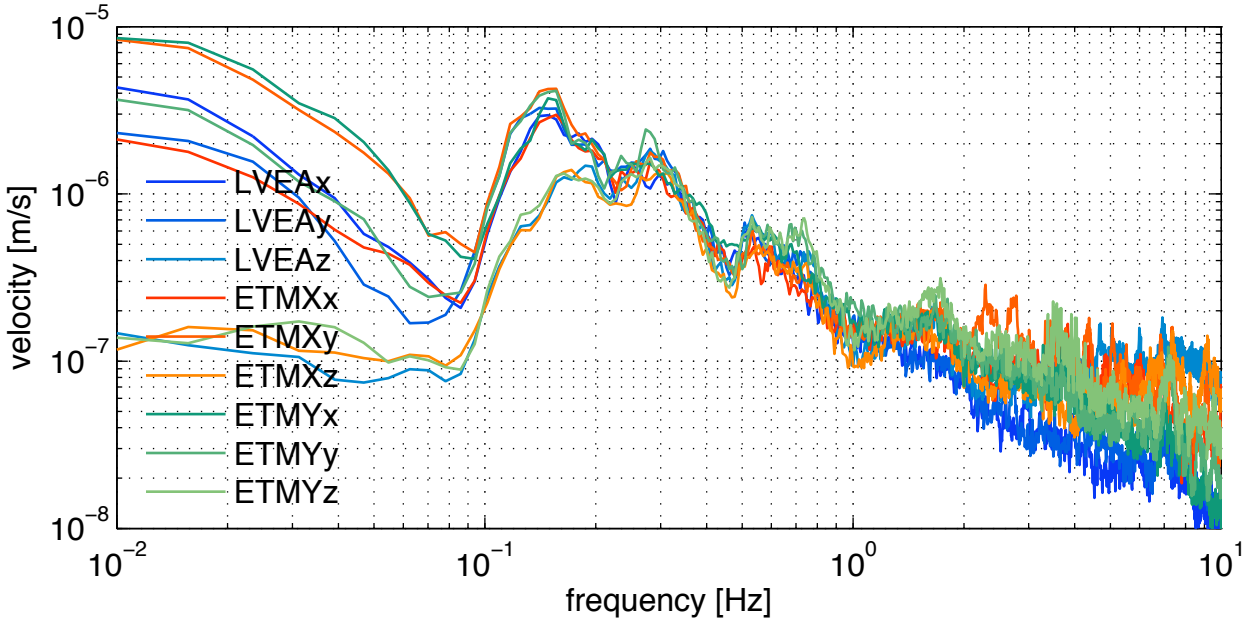

Figure C-5. Ground motion at time of ASC suppression demonstration (Figure 6-8). GPS time is 956751915 (May 1, 2010 08:15 CDT).

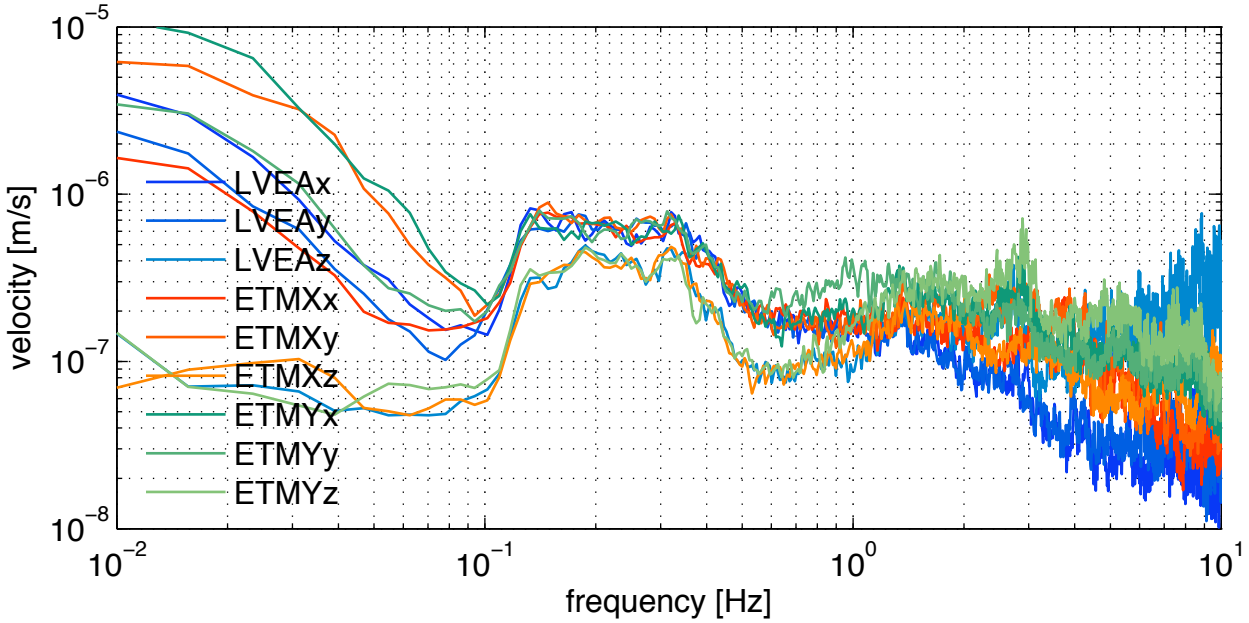

Figure C-6. Ground motion at time of input beam motion impression measurement (Figure 6-1). GPS time is 971128215 (Oct. 14, 2010 16:50 CDT).

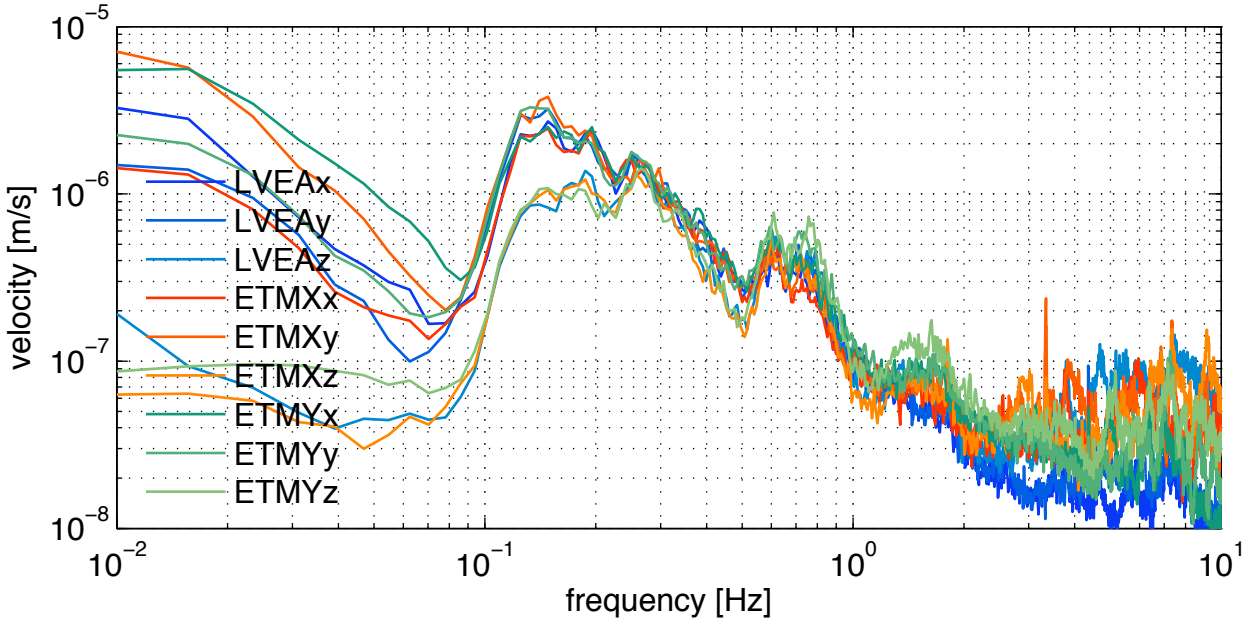

Figure C-7. Ground motion at time of beam spot motion measurement (Figure 6-7). GPS time is 956728935 (May 1, 2010 01:02 CDT).

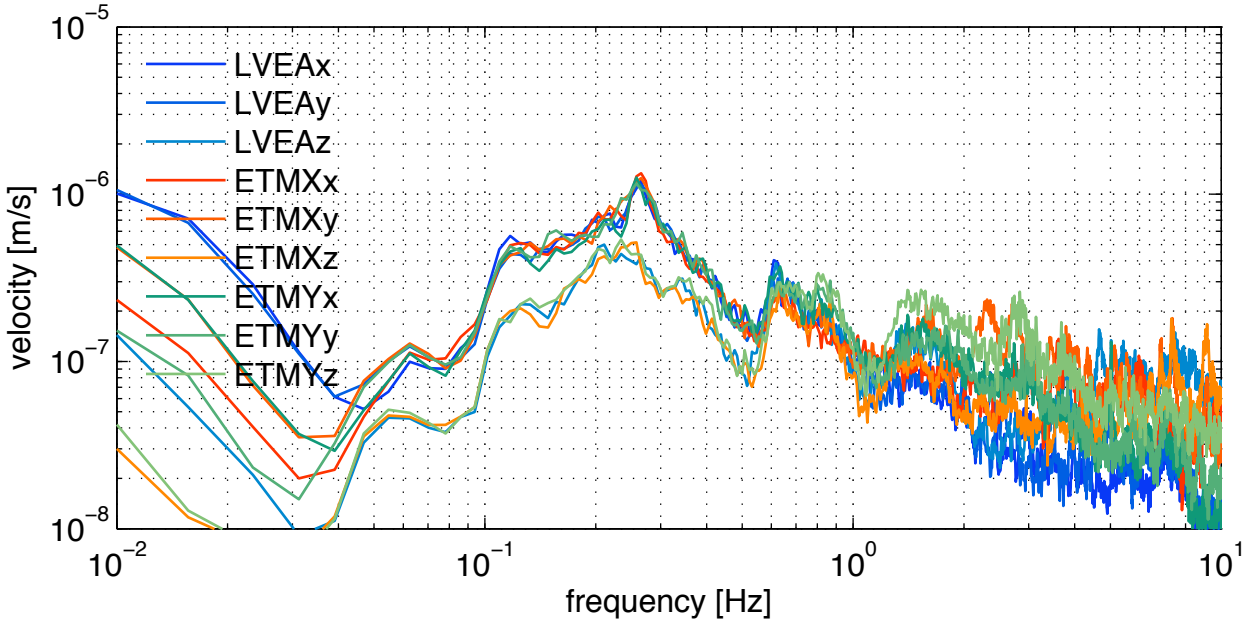

Figure C-8. Ground motion at time of ASC to DARM noisebudget plot (Figure 6-13). GPS time is 958456964 (May 21, 2010 01:02 CDT).

REFERENCES

[1] J. Weber, Physical Review Online Archive (Prola) **117**, 306 (1960).

[2] R. L. Forward, Physical Review D **17**, 379 (1978).

[3] S. M. Carroll (1997).

[4] P. R. Saulson, *Fundamentals of interferometric gravitational wave detectors* (World Scientific, 1994).

[5] B. P. Abbott *et al.*, Physical Review D (Particles, Fields, Gravitation, and Cosmology) **80**, 102001+ (2009).

[6] J. Abadie *et al.*, Physical Review D **83**, 122005+ (2011).

[7] J. Abadie *et al.*, The Astrophysical Journal **722**, 1504+ (2010).

[8] B. Allen and J. D. Romano, Physical Review D **59**, 102001+ (1999).

[9] M. Maggiore, Physics Reports **331**, 283 (2000).

[10] J. Abadie *et al.*, Classical and Quantum Gravity **27**, 173001+ (2010).

[11] R. A. Hulse and J. H. Taylor, Astrophys. J. Lett. **195**, L51 (1975).

[12] J. M. Weisberg and J. H. Taylor, in *Binary Radio Pulsars*, vol. 328 of *Astronomical Society of the Pacific Conference Series*, F. A. Rasio and I. H. Stairs, eds. (2005), vol. 328 of *Astronomical Society of the Pacific Conference Series*.

[13] G. Hobbs *et al.* (2009).

[14] P. Linsay, P. Saulson, R. Weiss, and S. Whitcomb, "A study of a long baseline gravitational wave antenna system", Tech. Rep. T830001, Massachusetts Institute of Technology (1983).

[15] B. P. Abbott *et al.*, Reports on Progress in Physics **72**, 076901+ (2009).

[16] F. Acernese *et al.*, Journal of Optics A: Pure and Applied Optics **10**, 064009+ (2008).

[17] H. Lück *et al.*, Classical and Quantum Gravity **23**, S71 (2006).

[18] Advanced LIGO Systems Group, "Advanced LIGO Systems Design", Tech. Rep. T010075, LIGO Laboratory (2009).

[19] R. Adhikari, P. Fritschel, and S. Waldman, "Enhanced LIGO", Tech. Rep. T060156, LIGO Laboratory (2006).

[20] T. Fricke *et al.*, Classical and Quantum Gravity, accepted for publication (2011).

[21] J. S. Kissel, "Calibrating and improving the sensitivity of the LIGO detectors", Ph.D. thesis, Louisiana State University (2010).

[22] M. Frede, B. Schulz, R. Wilhelm, P. Kwee, F. Seifert, B. Willke, and D. Kracht, Opt. Express **15**, 459 (2007).

[23] P. Willems, A. Brooks, M. Mageswaran, V. Sannibale, C. Vorvick, D. Atkinson, R. Amin, and C. Adams, "Thermal Compensation in Enhanced LIGO", (2009).

[24] H. Lück *et al.*, Journal of Physics: Conference Series **228**, 012012+ (2010).

[25] S. Solimeno, F. Barone, C. de Lisio, L. Di Fiore, L. Milano, and G. Russo, Physical Review A **43**, 6227 (1991).

[26] J. Sidles and D. Sigg, Physics Letters A **354**, 167 (2006).

[27] E. Hirose, K. Kawabe, D. Sigg, R. Adhikari, and P. R. Saulson, Appl. Opt. **49**, 3474 (2010).

[28] L. Barsotti and M. Evans, "Modeling of Alignment Sensing and Control for Enhanced LIGO", Tech. Rep. T080186, LIGO Laboratory (2009).

[29] D. Garfinkle, American Journal of Physics **74**, 196+ (2005).

[30] T. Fricke, "Homodyne Detection for Laser-Interferometric Gravitational Wave Detectors", Ph.D. thesis, Louisiana State University (2011).

[31] R. Adhikari, "Sensitivity and Noise Analysis of Big Interferometers", Ph.D. thesis, MIT (2004).

[32] P. R. Saulson, Physical Review D **42**, 2437 (1990).

[33] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, C. Adams, R. Adhikari, P. Ajith, B. Allen, G. Allen, and E. Amador Ceron, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **624**, 223 (2010).

[34] E. D. Black, American Journal of Physics **69**, 79 (2001).

[35] J. Camp, D. Reitze, and D. Tanner, "Input/Output Optics Conceptual Design", Tech. Rep. T960170, LIGO Laboratory (1996).

[36] J. Camp, D. Reitze, and D. Tanner, "Input Optics Design Requirements Document", Tech. Rep. T960093, LIGO Laboratory (1997).

[37] P. Fritschel, R. Bork, G. González, N. Mavalvala, D. Ouimette, H. Rong, D. Sigg, and M. Zucker, Appl. Opt. **40**, 4988 (2001).

[38] R. Adhikari, A. Bengston, Y. Buchler, T. Delker, D. Reitze, Q.-z. Shu, D. Tanner, and S. Yoshida, "Input Optics Final Design", Tech. Rep. T980009, LIGO Laboratory (1998).

[39] UF LIGO Group and IAP Group, "Upgrading the Input Optics for High Power Operation", Tech. Rep. E060003, LIGO Laboratory (2006).

[40] E. A. Khazanov, O. V. Kulagin, S. Yoshida, D. B. Tanner, and D. H. Reitze, IEEE Journal of Quantum Electronics **35**, 1116 (1999).

[41] A. L. Bullington, B. T. Lantz, M. M. Fejer, and R. L. Byer, Appl. Opt. **47**, 2840 (2008).

[42] M. Arain, "A Note on Substrate Thermal Lensing in Mode Cleaner", Tech. Rep. T070095, LIGO Laboratory (2007).

[43] V. Quetschke, Coherent Optical Technologies and Applications pp. CMC1+ (2008).

[44] F. Raab and S. Whitcomb, "Estimation of Special Optical Properties of a Triangular Ring Cavity", Tech. Rep. T920004, LIGO Laboratory (1992).

[45] E. Khazanov, N. Andreev, A. Babin, A. Kiselev, O. Palashov, and D. H. Reitze, J. Opt. Soc. Am. B **17**, 99 (2000).

[46] G. Mueller, R. S. Amin, D. Guagliardo, D. McFeron, R. Lundock, D. H. Reitze, and D. B. Tanner, Classical and Quantum Gravity **19**, 1793+ (2002).

[47] E. Khazanov *et al.*, IEEE Journal of Quantum Electronics **40**, 1500 (2004).

[48] The VIRGO Collaboration, Appl. Opt. **47**, 5853 (2008).

[49] N. P. Barnes and L. B. Petway, J. Opt. Soc. Am. B **9**, 1912 (1992).

[50] T. Delker, R. Adhikari, S. Yoshida, and D. Reitze, "Design Considerations for LIGO Mode-Matching Telescopes", Tech. Rep. T970143, LIGO Laboratory (1997).

[51] "Component Specification: Substrate, Mode Cleaner Flat Mirror", Tech. Rep. E970148, LIGO Laboratory (1998).

[52] M. Punturo, "The mirror resonant modes method for measuring the optical absorption", Tech. Rep. VIR-001A-07, VIRGO (2007).

[53] P. Kwee, F. Seifert, B. Willke, and K. Danzmann, The Review of scientific instruments **78** (2007).

[54] I. Snetkov, I. Mukhin, O. Palashov, and E. Khazanov, Opt. Express **19**, 6366 (2011).

[55] P. Fritschel and D. Shoemaker, "Alignment Sensing/Control Design Requirements Document", Tech. Rep. T952007, LIGO Laboratory (1997).

[56] ISC Group, "ASC Wavefront Sensing Final Design", Tech. Rep. T980064, LIGO Laboratory (1998).

[57] N. Wiener, *Extrapolation, Inerpolation, and Smoothing of Stationary Time Series* (The M.I.T. Press, 1975).

[58] R. P. Feynman, *Feynman lectures on physics.* (Addison Wesley Longman, 1970).

[59] N. Mavalvala, D. Sigg, and D. Shoemaker, Appl. Opt. **37**, 7743 (1998).

[60] N. Mavalvala, "Alignment issues in laser interferometric gravitational-wave detectors", Ph.D. thesis (1997).

[61] D. Sigg, "Wavefront Sensor", Tech. Rep. T960111, LIGO Laboratory (1996).

[62] D. Z. Anderson, Appl. Opt. **23**, 2944 (1984).

[63] P. Fritschel, N. Mavalvala, D. Shoemaker, D. Sigg, M. Zucker, and G. González, Appl. Opt. **37**, 6734 (1998).

[64] A. E. Siegman, *Lasers* (University Science Books, 55D Gate Five Road, Sausalito, CA 94965, 1986).

[65] L. Barsotti, private communication (2008).

[66] R. DeRosa, J. Driggers, D. Atkinson, H. Miao, V. Frolov, M. Landry, and R. Adhikari, in preparation (2011).

[67] S. Ballmer, "LIGO interferometer operating at design sensitivity with application to gravitational radiometry", Ph.D. thesis, Massachusetts Institute of Technology (2006).

[68] L. Barsotti, M. Evans, and P. Fritschel, Classical and Quantum Gravity **27**, 084026+ (2010).

BIOGRAPHICAL SKETCH

Katherine Laird Dooley was born in Rhinebeck, NY to Janine L Protzman and Alan P Dooley. She grew up in Poughkeepsie, NY with three younger brothers, Brian, Greg, and Tim, and graduated from Spackenkill HS in 2002. She went to Vassar College and graduated in 2006 with a major in physics and minors in French and mathematics. Kate began her physics PhD program at the U. of Florida in the fall of 2006, and moved to Baton Rouge, LA in the fall of 2007 to carry out her graduate research at LIGO Livingston.