High-Power Operation of Interferometric Gravitational-Wave Detectors

by

Aaron Buikema

B.S., Haverford College (2013)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

Author .....
Department of Physics
December 23, 2019

Certified by .....
Nergis Mavalvala
Curtis and Kathleen Marble Professor of Astrophysics
Thesis Supervisor

Accepted by .....
Nergis Mavalvala
Associate Department Head of Physics

High-Power Operation of Interferometric Gravitational-Wave Detectors

by

Aaron Buikema

Submitted to the Department of Physics
on December 23, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

With the conclusion of the first two observing runs of the Advanced LIGO detectors, which saw the first direct detection of gravitational waves, we are firmly in the era of gravitational-wave astronomy. To reach the highest sensitivities, current interferometric gravitational-wave detectors are designed for hundreds of kilowatts of circulating optical power. At these high circulating powers, the sensitivity of the detectors to gravitational waves will be limited by the quantum properties of the light: shot noise at frequencies above $\sim 100\,\mathrm{Hz}$, and quantum radiation pressure noise at lower frequencies. To reach the high powers necessary for achieving the quantum noise limits imposed by the light, it is essential to solve the control problems and understand the additional noise introduced by high power operation. Additionally, development of high-power laser sources that reach the stringent noise and reliability requirements is crucial.

This work comprises three experiments aimed at reaching the radiation-pressure-dominated regime of interferometric gravitational-wave detectors. The first part presents results from a high-power, meter-long Fabry-Pérot Michelson interferometer to probe classical and quantum radiation pressure effects using a gram-scale mechanical oscillator. The second part is an exploration of the effects of electric fields and charging of test masses on the sensitivity of the LIGO detectors, which may limit the ability to observe radiation-pressure effects. Finally, we describe the development and characterization of a high-power, narrow-linewidth ytterbium-doped fiber amplifier for use in future gravitational-wave detectors.

Thesis Supervisor: Nergis Mavalvala
Title: Curtis and Kathleen Marble Professor of Astrophysics

Acknowledgments

I have had the privilege of working with so many wonderful individuals in a number of different settings. This thesis was only possible due to their hard work and support.

First and foremost, thank you to my advisor, Nergis; you are a tireless advocate for all of your students. Thank you to the rest of my committee, Scott Hughes and Kerstin Perez, for your valuable support and feedback during this process.

Thank you to the entire team at Lincoln Labs for welcoming me. A special thanks to Steve Augst for teaching me everything I know about fiber amplifiers and Franklin Jose for working in the trenches and keeping an eye on me. Thank you to Peter Fritschel for supervising the fiber amplifier project and helping whenever needed.

Thank you to Valera and Brian for giving me the opportunity to spend a few months at LIGO Livingston. In particular, thank you to Carl Blair for taking me under your wing and getting me up to speed with the EFM, and to Carl Adams and Michael Laxen, for getting into the weeds with me and for some reason believing I had a plan. Thank you to all of the commissioners and staff for welcoming me to a place very different from home, and most importantly, teaching me how to eat crawfish.

To my roommates Clara and everyone who lived at Café Faraz (Woody, Southwick, Zach, Desk, Marjon, Sedona, Nora, Henri), thank you for taking my mind off physics to teach me about psychology, biology, transportation, and so many other things; for indulging me in weird weekend adventures; and for checking in.

Thanks to my control room σύντροφος (comrade) Antonios Kontos for making the many late-night sessions bearable. Your advice in times of need to "make it better" was always appreciated.

To my fellow LIGO grad students: you provided solace and advice whenever needed (scientific and otherwise), made sure I never took myself too seriously, and believed in me even when I didn't.

Thank you to everyone else in the MIT LIGO Lab, in particular, Marie Woods for her exceptional administrative work, Fred Donovan for IT assistance, and to Myron MacInnis, who taught me the most useful skills to have as an experimentalist. This

is one of the most helpful, knowledgeable, and supportive cast of characters I've ever met, and I am so fortunate to have spent so much of my time with you.

Lastly, thank you to my family. To my parents, thank you for supporting me every step of the way. And to Rachel, thanks for always being there when I needed it.

Contents

List of Figures

List of Tables

Conventions

We will use the following conventions in this thesis:

- All power spectral densities (PSDs) will be single-sided. Following Bendat and Piersol [1], we will denote by $G_{aa}(\Omega)$ the one-sided PSD as a function of angular frequency Ω in units of $[a]^2/\text{Hz}$, where the units of a is denoted by $[a]$. While it is easier mathematically to deal with PSDs, we will most often plot results using amplitude spectral densities (ASDs), defined as $\sqrt{G_{aa}(\Omega)}$ and with units of $[a]/\sqrt{\text{Hz}}$.

- The transmission and reflection of optical elements will usually be given in terms of amplitude, will be real, and will be denoted with lowercase letters t, r. In power, they will be uppercase ($T \equiv t^2$, $R \equiv r^2$).

- So that we don't need to define front and back surfaces of optics, we will use the phase convention where transmission through an optic picks up an additional 90° phase shift and reflections do not pick up any additional phase. Thus, the transformation matrix for a beam splitter, for example, is given by

$$\begin{pmatrix} it & r \\ r & it \end{pmatrix}, \tag{1}$$

  for real r, t. For a discussion of the subtleties that allow us to do this, see [2].

- Following standard convention, we will generally refer to the FWHM of a cavity resonance in frequency units as $\Delta\nu$, whereas the HWHM in units of angular frequency will be written as $\gamma = \pi\Delta\nu$.

- While the effect of a passing gravitational wave is to induce a strain, for most of this thesis we will use length units. The conversion between the two for the LIGO detectors is $\Delta L = Lh$ for strain h and $L = 4\,\text{km}$.

Acronyms

A+ aLIGO Plus

aLIGO Advanced LIGO

AMD acoustic mode damper

AOM acousto-optic modulator

AS anti-symmetric

ASD amplitude spectral density

ASE amplified spontaneous emission

BBH binary black hole

BNS binary neutron star

BS beamsplitter

CARM common arm length

CBC coherent beam combining

CE Cosmic Explorer

CMRR common-mode rejection ratio

CP compensation plate

CW continuous-wave

DAC digital-to-analog converter

DARM differential arm length

DCPD DC photodetector

DOF degree of freedom

EFM electric field meter

EOM electro-optic modulator

ERM end reaction mass

ESD electrostatic drive

ET Einstein Telescope

ETM end test mass

FP Fabry-Pérot

FPMI Fabry-Pérot Michelson interferometer

FSR free spectral range

FWHM full width at half maximum

GW gravitational wave

HC Hänsch and Couillaud

HEPI hydraulic external pre-isolator

HOM higher-order mode

HWHM half width at half maximum

iLIGO Initial LIGO

IMC input mode cleaner

ISI internal seismic isolation

ISS intensity stabilization servo

ITM input test mass

LHO LIGO Hanford Observatory

LIGO Laser Interferometer Gravitational-Wave Observatory

LLO LIGO Livingston Observatory

LMA large-mode-area

LO local oscillator

MFD mode field diameter

MICH Michelson DOF

MOPA master oscillator power amplifier

NPRO non-planar ring oscillator

OMC output mode cleaner

OMIT optomechanically induced transparency

OPO optical parametric oscillator

OSEM optical sensor and electromagnetic motor

PBGF photonic-bandgap fiber

PBS polarizing beam splitter

PC polarization controller

Pcal photon calibrator

PCF photonic crystal fiber

PD photodetector

PDH Pound-Drever-Hall

PI parametric instability

PM polarization-maintaining

PMC pre-mode cleaner

PRM power recycling mirror

PSD power spectral density

PSL pre-stabilized laser

PUM penultimate mass

QPD quadrant photodetector

QRPN quantum radiation pressure noise

RH ring heater

RPN relative power noise

SBS stimulated Brillouin scattering

SNR signal-to-noise ratio

SQL standard quantum limit

SRM signal recycling mirror

TCS thermal compensation system

UGF unity-gain frequency

YDFA ytterbium-doped fiber amplifier

Chapter 1

Introduction

Developments in precision optical interferometry have led to record-setting measurements of force sensing [3], atomic clocks [4], and probes of fundamental physics [5]. In particular, with 100 kW of circulating light, optical interferometry has been used to make sub-attometer displacement measurements for the detection of gravitational waves [6, 7]. Further sensitivity improvements will come with even higher operating powers.

This thesis will describe work carried out to improve the sensitivity of current and future gravitational-wave detectors, specifically for high-power operation.

1.1 Gravitational Wave Detection

It is fitting that the theory of general relativity was in part inspired by the measurements of Michelson and Morley [8], who used optical interferometry to disprove the existence of the aether. More than a century later, optical interferometry continues to play a major role in tests of general relativity in the form of interferometric gravitational-wave detectors.

First predicted by Einstein in 1916 as a consequence of his theory of general relativity, gravitational waves are disturbances in the curvature of spacetime due to an accelerating quadrupole mass distribution. A passing gravitational wave (GW) will introduce a differential strain between a ring of free-masses (figure 1-1), which can

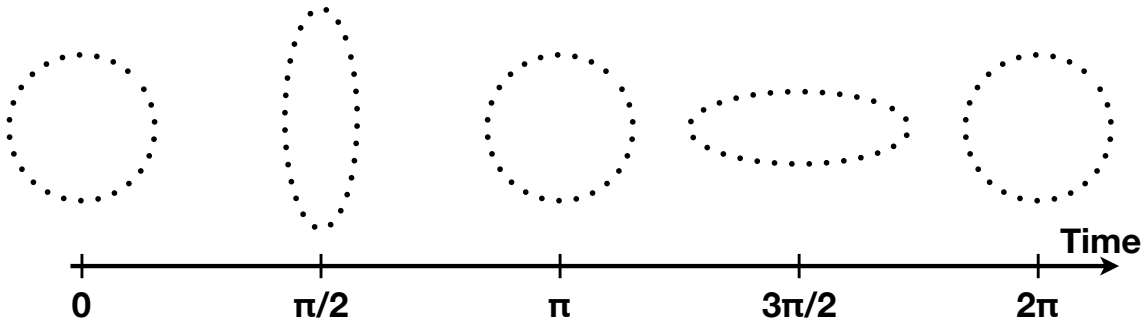

Figure 1-1: The effect of a passing +-polarized gravitational wave traveling into the page on a ring of free masses. This effect has been greatly exaggerated.

be measured by a sufficiently sensitive laser interferometer. The first evidence for the existence of gravitational waves came from measurements of the orbital decay of a binary pulsar system [9]. The first direct detection of GWs was announced [10] soon after the completion of the Advanced LIGO (aLIGO) detectors in 2015. Since then, GWs have been detected from a binary neutron star (BNS) [11], which heralded a comprehensive multi-messenger followup observation campaign across the electromagnetic spectrum [12]. All together, there have been eleven confirmed GW events from the first and second Advanced LIGO (aLIGO) observing runs (O1 and O2) [13]. Together with a number of other detectors in a worldwide network that includes Advanced Virgo [14], GEO 600 [15], and KAGRA [16], many more events are to come in O3 and beyond.

1.2 Laser Interferometry and LIGO

In the 1970's it was proposed that these waves could be measured directly via laser interferometry[1], and in the 1990's, the Laser Interferometer Gravitational-Wave Observatory (LIGO), a pair of 4 km long Fabry-Pérot Michelson interferometers (FPMIs), one located in Hanford, WA, and the other in Livingston, LA, was constructed. The layout of the aLIGO detectors is shown in figure 1-2. The mirrors are all suspended from silica fibers to provide isolation from ground motion and to act as free-falling

[1]See [19] for a history of GW measurements.

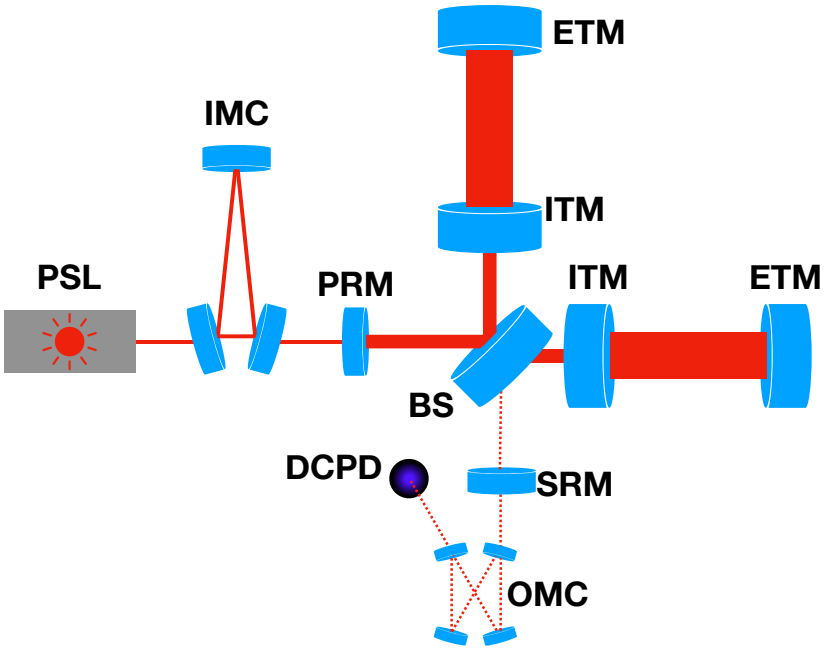

Figure 1-2: Optical layout of the Advanced LIGO detectors. Light from the pre-stabilized laser (PSL) enters the input mode cleaner (IMC) before traveling to the dual-recycled Fabry-Pérot Michelson interferometer. A beamsplitter (BS) splits the light before sending it to two 4 km arm cavities, each consisting of an input test mass (ITM) and end test mass (ETM). An output mode cleaner (OMC) cleans the output beam before the differential arm length (DARM) (GW signal) is read out on a DC photodetector (DCPD). The power recycling mirror (PRM) returns the light reflected from the interferometer to enhance the circulating power, and a signal recycling mirror (SRM) increases the detector bandwidth. Not shown is the squeezer [17].

test masses (figure 1-3). Each test mass is 40 kg to minimize radiation pressure effects. Arm cavities amplify the circulating laser power and increase the effective length of the arms. These detectors must be sensitive enough to measure strains at the level of $h \approx 10^{-21}$ and below.

Initial LIGO (iLIGO) ran from 2002–2007 and consisted of five science runs (S1–S5). Subsequently, the instrument was upgraded to Enhanced LIGO [20, 21] and observed for a sixth science run that concluded in 2010. A number of improvements were installed, many as a test before the aLIGO upgrade, including a hydraulic external pre-isolator (HEPI) for improved seismic isolation [22], a new readout scheme [23], and higher input power [24]. No gravitational waves were detected.

Subsequently, the LIGO instruments were upgraded to Advanced LIGO (aLIGO) [26],

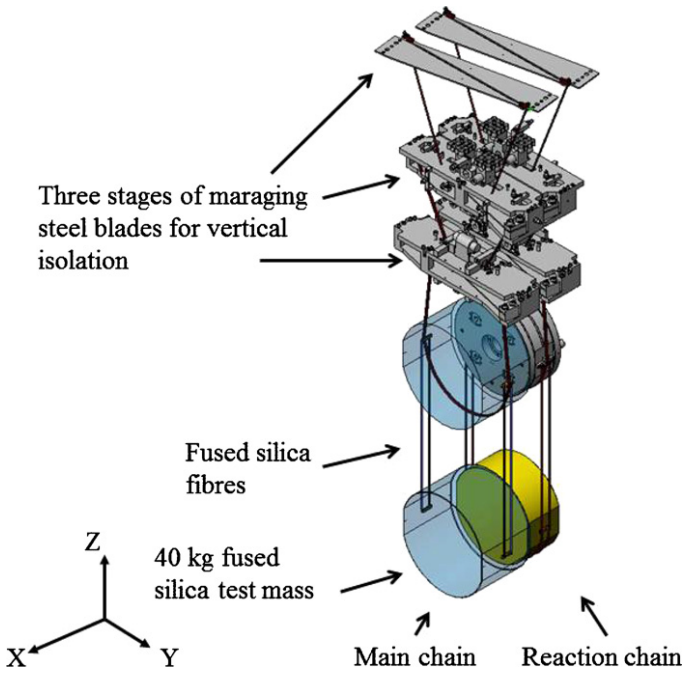

Figure 1-3: aLIGO test mass suspension. Each test mass is 40 kg. Not shown is the cage surrounding the suspension. Adapted from [18].

which began the first observing run (O1) in 2015. This upgrade saw improved suspensions, larger test masses, the addition of a signal recycling mirror, and, most relevant to this thesis, an increase in the available laser power. O1 saw the first direct detection of gravitational radiation [10], and subsequent observing runs have seen even more events.

The steady increase in optical power over these observing runs, as well as the addition of a "squeezer" [17], have produced the best high-frequency sensitivity yet of $5 \times 10^{-24}/\sqrt{\text{Hz}}$ at 100 Hz. The high-frequency region is important for low-mass sources like BNSs, as these coalescences occur in the kHz region. Particularly for BNSs, in the last part of the inspiral ($f_{\text{GW}} \gtrsim 500\,\text{Hz}$) tidal effects become most important [27] and thus improved sensitivity in this region is paramount for understanding the neutron star equation of state.

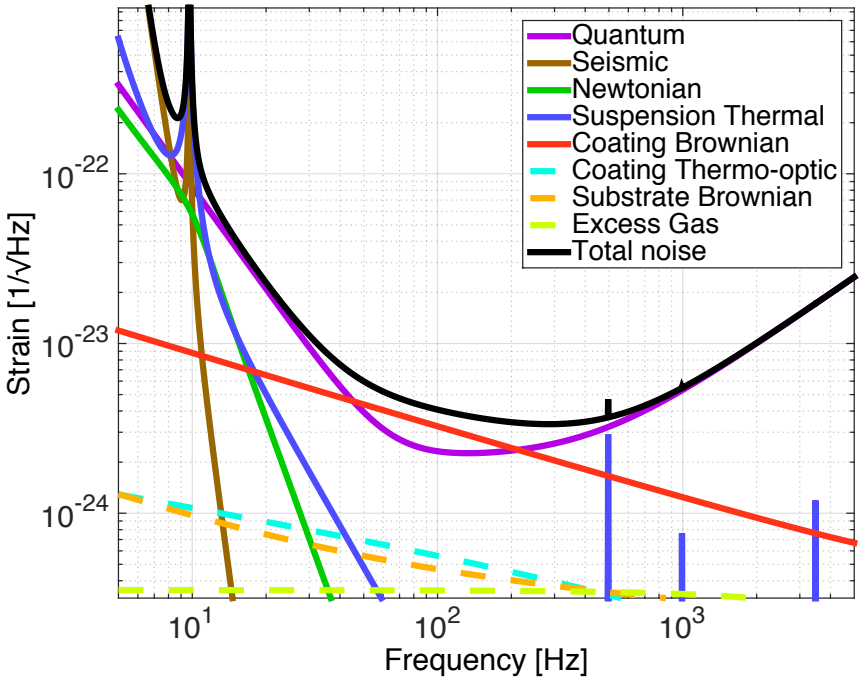

Figure 1-4: Advanced LIGO design noise curve. Adapted from [25].

1.3 High-Power Operation

Figure 1-4 shows the design sensitivity curve for the aLIGO detectors. The sensitivity is limited at most frequencies by coating thermal noise due to the Brownian motion of the dielectric reflective coating on the test mass surfaces, and "quantum noise," which is due to the quantum nature of the light used to probe the position of the test masses. As we will see in chapter 2, the high-frequency portion of this noise term (shot noise) can be improved with increased circulating power. Indeed, the aLIGO detectors are designed to maximize the shot noise sensitivity by using a high-power stabilized laser [24], a power recycling cavity, arm cavities, and, as of O3, squeezed light injection [17]. At design sensitivity, the detectors will each have 750 kW of circulating light [26].

Operating at such high powers comes with a number of challenges. Thermal distortions of the mirror due to finite absorption will change the cavity mode matching. Radiation pressure effects can modify the angular control plant [28], and can lead to parametric instabilities (PIs) due to the coupling of laser light to body modes of the optics [29]. Finally, due to the quantum mechanical nature of the light used to probe

the mirror position, at increasing powers the measurement back action introduces an additional fluctuating force that will limit sensitivity, as we will discuss in section 2.4. While some of these issues have largely been solved through a thermal compensation system (TCS) [30] and acoustic mode dampers (AMDs) that reduce the body mode Q factors [31], a lack of available laser power and other effects have limited further power increases [32].

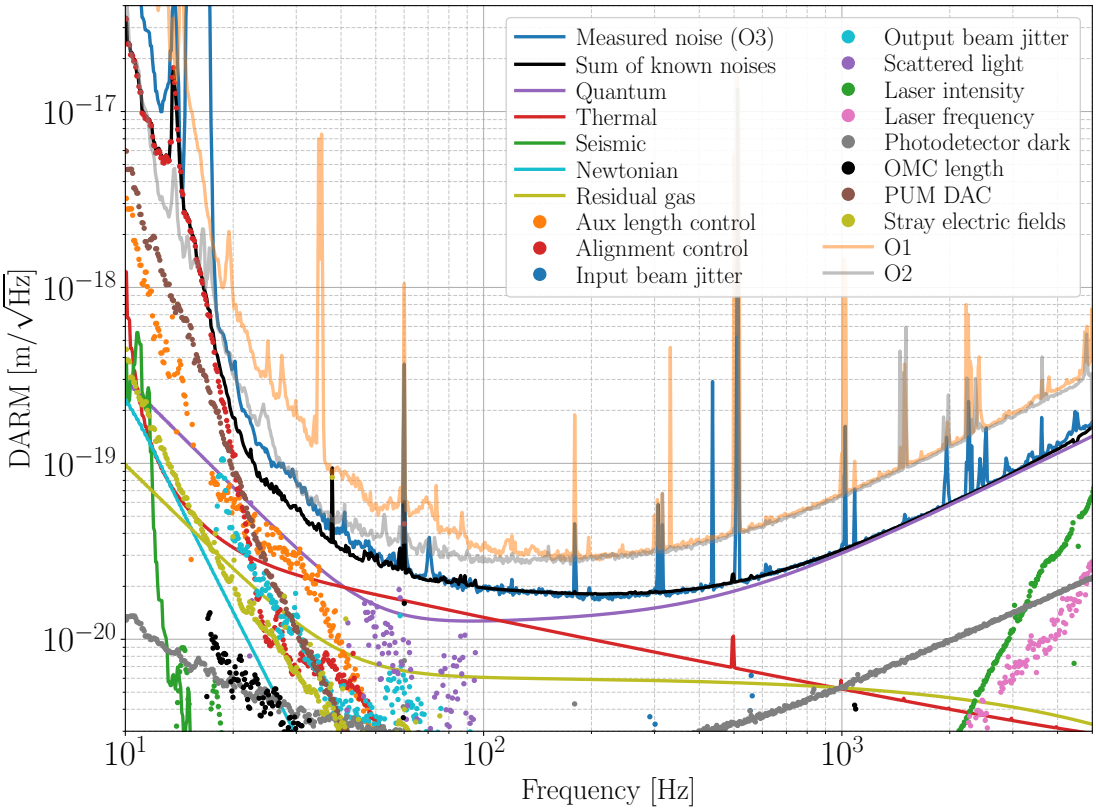

Figure 1-5: Current noise budget for LIGO Livingston Observatory (LLO). While limited by shot noise at high frequencies, below 100 Hz the detector is limited by technical and unknown noise sources. Adapted from [32].

With the recent addition of squeezed light injection [17], the shot noise sensitivity at LIGO Livingston Observatory (LLO) can be improved by 3.2 dB, roughly the equivalent of a 110% increase in circulating power. In principle, increased levels of squeezing will allow for improved sensitivity without the problems of increased laser power. In practice, however, optical losses will limit the amount of injected squeezed light. Additionally, without a filter cavity [33], the improvement in shot

noise sensitivity will come at the expense of excess quantum radiation pressure noise (QRPN) (see section 2.4).

As a result, continued increases in power will be required for additional improvements in the sensitivity of interferometric GW detectors. Finally, additional noise sources (figure 1-5) limit the sensitivity. This thesis will discuss some of the limiting factors to increasing power in such interferometers, as well as work to bring these detectors into a regime dominated by the radiation pressure effects.

1.4 Thesis Overview

This thesis is presented in four parts: Chapter 2 provides an introduction to the precision interferometry of GW detectors, reviews the theory of optical cavities, and provides a brief introduction to quantum optics in the context of optomechanics experiments such as LIGO. Chapter 3 describes a high-power gram-scale Fabry-Pérot Michelson interferometer to study the dynamics and noise properties in a radiation-pressure-dominated regime. Chapter 4 describes work carried out at the LIGO Livingston Observatory (LLO) to measure, characterize, and mitigate the noise contribution of test mass charging and stray electric field coupling, which could mask the effect of radiation pressure noise. Chapter 5 discusses the development and characterization of a narrow-linewidth ytterbium-doped fiber amplifier for use in future gravitational wave detectors. Finally, we conclude with an outlook on the future of high-power operation of these and future gravitational-wave detectors.

Chapter 2

Optical Interferometry and Quantum Optomechanics

We will begin with an introduction to optical interferometry, followed by a discussion of optical resonators and how they can improve the sensitivity of our instruments. We will then describe how to generate a control signal from such cavities and provide examples, including a novel technique that can be used for an arbitrarily birefringent resonator. Finally, we provide an introduction to optomechanics and discuss the limits to measurements of such systems.

As we will see, the sensitivity of these methods is generally improved with increased laser power, which is the motivation of the rest of this thesis.

2.1 Introduction to Interferometry: Michelson Interferometer

As an introduction to the concept of interferometry and the motivation for going to high power, let's consider a simple Michelson interferometer (figure 2-1) with arm lengths l_x and l_y, respectively. Assume the beamsplitter is a 50:50 splitter ($t^2 = r^2 = 0.5$), that the end mirrors are perfect reflectors, and that this system is lossless. If we send in a monochromic light beam $E_{\text{in}} = E_0 e^{i(\omega t - kx)}$ and propagate this

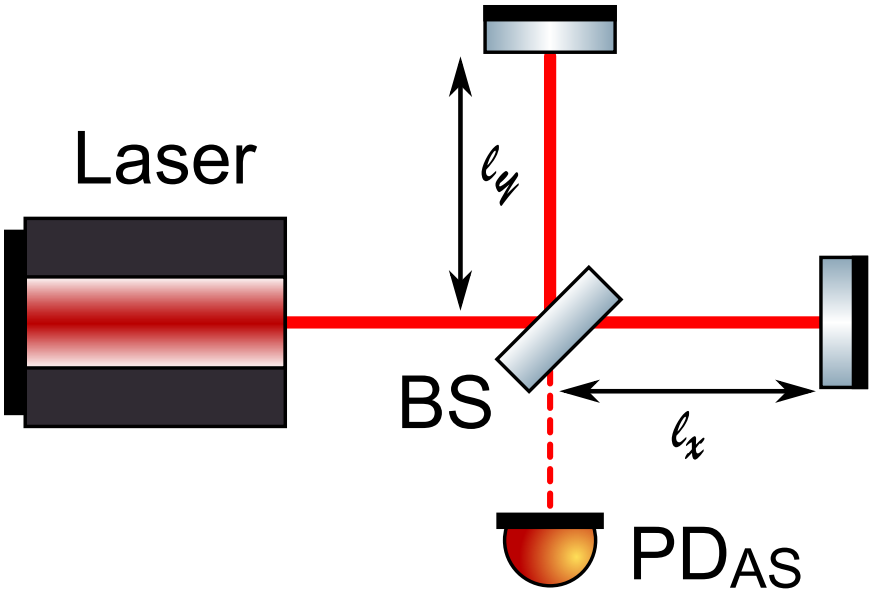

Figure 2-1: Basic Michelson interferometer. A beamsplitter (BS) splits the incoming light down two arms. After reflecting back and accumulating some phase, the two beams interfere and this power is measured at the anti-symmetric (AS) port.

through the system, we find that the reflected and output (anti-symmetric) fields are given by

$$\begin{aligned} E_{\text{refl}} &= itE_x + rE_y \\ &= \left(-t^2 e^{-2ikl_x} + r^2 e^{-2ikl_y}\right) E_{\text{in}} \\ &= \frac{1}{2} e^{-2ikl_x} \left(e^{2ik(l_x - l_y)} - 1\right) E_{\text{in}} \end{aligned} \tag{2.1}$$

$$\begin{aligned} E_{\text{as}} &= rE_x + itE_y \\ &= irt \left(e^{-2ikl_x} + e^{-2ikl_y}\right) E_{\text{in}} \\ &= \frac{i}{2} e^{-2ikl_x} \left(e^{2ik(l_x - l_y)} + 1\right) E_{\text{in}}. \end{aligned} \tag{2.2}$$

The measured power $P \propto |E|^2$ at each port is

$$P_{\text{refl}} = \frac{1}{2}[1 - \cos 2\phi] P_{\text{in}} = P_{\text{in}} \sin^2 \phi \tag{2.3}$$

$$P_{\text{as}} = \frac{1}{2}[1 + \cos 2\phi] P_{\text{in}} = P_{\text{in}} \cos^2 \phi. \tag{2.4}$$

for $\phi \equiv k(l_x - l_y) = \omega(l_x - l_y)/c$. The total power exiting the interferometer is equal to the input power, as must be the case of a lossless system. Thus, the Michelson measures the *difference* in phase accumulated by the light traveling down both arms; this fact allows such geometries to measure the differential strain of a gravitational wave. For small changes in ϕ, we will observe a change in power at the AS port of

$$\frac{dP_{\text{as}}}{d\phi}\delta\phi = -P_{\text{in}} \sin 2\phi\delta\phi = -2P_{\text{in}} \sin\phi\cos\phi\delta\phi. \tag{2.5}$$

In practice, we will want to keep the power at the AS port small, since fluctuations in the input power and frequency will couple to the output; the differential nature of this setup eliminates some of this common-mode noise. However, there is a more fundamental source of noise related to the detection of discrete quanta of light: the photon shot noise. Because the arrival of photons at the photodetector is a random process and each arrival is independent[1], the power spectrum is flat and is given by

$$G_{PP}(\Omega) = 2\hbar\omega P. \tag{2.6}$$

Note that this noise is independent of frequency and is proportional to laser power (in PSD units). This will appear as an apparent phase noise given by

$$\begin{aligned} G_{\phi\phi}(\Omega) &= G_{P_{\text{as}}P_{\text{as}}}(\Omega)\left(\frac{dP_{\text{as}}}{d\phi}\right)^{-2} \\ &= \frac{2\hbar\omega P_{\text{as}}}{4P_{\text{in}}^2 \sin^2\phi\cos^2\phi} = \frac{\hbar\omega}{2P_{\text{in}}\sin^2\phi}. \end{aligned} \tag{2.7}$$

This shot noise limit is minimized for $\sin^2\phi = 1$, which corresponds to no light at the AS port, and for this reason it is often referred to as the "dark port[2]."

Additionally, increased input optical power improves the shot noise sensitivity. As

[1]Note that this assumption breaks down when we use "squeezed light" and can then drop below the usual shot noise sensitivity [34].

[2]Note that this simple interferometer cannot be operated exactly at $\phi = 0$ with this readout scheme since the signal goes to zero. Additionally, other technical noise sources like photodetector noise become important, so there must be a small offset to ϕ. In the LIGO detectors, an analog of the Pound-Drever-Hall (PDH) locking scheme (section 2.3.3) is used and thus we can keep the Michelson at zero offset [35, 36].

we will see in section 2.4, this is not the full story. However, even in the complex aLIGO interferometers it remains true that, all else equal, higher optical power allows for greater shot-noise-limited sensitivity.

As we have shown, the Michelson interferometer provides a way to measure the phase difference between two light beams propagating down the two arms, with the phase sensitivity limited by shot noise. In the next section, we will show one tool that can be used to enhance the phase sensitivity.

2.2 Optical Resonators

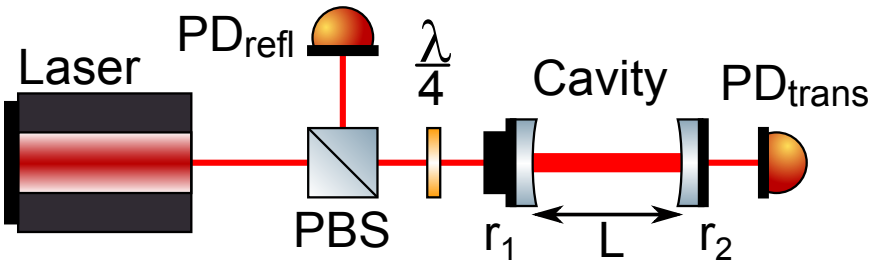

Figure 2-2: A Fabry-Pérot (FP) resonator setup. The cavity consists of input mirror with field reflectivity r_1 and output mirror with field reflectivity r_2. One photodetector measures the transmitted beam power, while the second measures the reflected beam with the use of a quarter-wave plate $\lambda/4$ and a polarizing beam splitter (PBS).

We can enhance the phase sensitivity of our Michelson interferometer by using optical resonators. Figure 2-2 shows the simplest form of optical resonator, a Fabry-Pérot (FP) cavity of length L and mirror field reflectivity and transmission of r_i and t_i for $i = 1, 2$ referring to the input and output mirrors, respectively. In the case of no loss, energy conservation requires $r_i^2 + t_i^2 = 1$. We will solve for the quasistatic solution for the fields around the cavity for a monochromatic input beam, such as from a laser. First let's consider the field immediately after the input mirror E_{circ}. Due to the finite transmission of the input mirror, this field will be a superposition of a leaked input field E_{in} and the circulating field that has traversed the cavity and picked up a complex roundtrip "gain" g_{rt}:

$$E_{\text{circ}} = it_1 E_{\text{inc}} + g_{rt} E_{\text{circ}}. \tag{2.8}$$

In general, this gain factor will depend on the roundtrip length of the cavity and the laser frequency[3]. We can then write the circulating field in terms of the incident field[4]. Continuing in this manner for the cavity reflected and transmitted fields, we can relate all such fields to the input field:

$$g_{\text{circ}} \equiv \frac{E_{\text{circ}}}{E_{\text{in}}} = \frac{it_1}{1-g_{rt}} = \frac{it_1}{1-r_1 r_2 \exp[-i2L\omega/c]} \tag{2.9a}$$

$$t_{\text{cav}} \equiv \frac{E_{\text{trans}}}{E_{\text{in}}} = it_n\alpha' \exp[-i\omega p'/c] g_{\text{circ}} = -\frac{t_1 t_2 \exp[-iL\omega/c]}{1-r_1 r_2 \exp[-i2L\omega/c]} \tag{2.9b}$$

$$r_{\text{cav}} \equiv \frac{E_{\text{r}}}{E_{\text{in}}} = r_1 - \frac{t_1^2}{r_1}\frac{g_{rt}}{1-g_{rt}} = \frac{r_1 - r_2 \exp[-i2L\omega/c]}{1-r_1 r_2 \exp[-i2L\omega/c]}, \tag{2.9c}$$

where we have defined the complex roundtrip gain factor

$$g_{rt} = \alpha r_1 r_2 (r_3 ...) \exp[-i\omega p/c] = r_1 r_2 \exp[-i2L\omega/c], \tag{2.10}$$

for cavity mirror (field) reflectivity r_i, roundtrip power loss $1-\alpha^2$, roundtrip length $p = 2L$, and speed of light in the cavity medium c[5]. The values α' and p' correspond to the value associated with the path from input mirror to mirror n from which the transmitted beam is measured. We have included the expressions for a general optical resonator in terms of g_{rt}, as well as the values for a lossless Fabry-Pérot (FP) cavity. We will focus on the reflected field for much of the following analysis, as it consists of a promptly reflected beam as well as leakage from the cavity. This properly means it will generally provide the best length sensing signal. Note that for the reflected beam, the reflectivity of non-input mirrors is on the same footing as cavity loss; for this reason, many authors will include output transmission as "loss."

A few things stick out from these expressions. First, equation (2.9a) indicates that we can produce a significantly higher circulating power than the input power.

[3]Because we are calculating the quasistatic solution, we have ignored any time delay in g_{rt} due to the finite speed of light. We will return to this point.

[4]This can also be solved by considering the infinite geometric sum consisting of the light that has made n roundtrips to get the same result.

[5]These equations all assume infinite plane wave inputs and do not consider any other effects on the input light, such as polarization shifts. In general there will exist a set of cavity eigenmodes, and these expressions apply to each eigenmode.

Second, this condition occurs for $g_{rt} \approx 1$, which occurs for

$$\omega p/c = 2\pi n \iff \nu p/c = n, n \in \mathbb{Z} \tag{2.11}$$

for optical frequency $\nu = \omega/2\pi$. The separation between consecutive resonance conditions is the free spectral range (FSR):

$$\nu_{\text{FSR}} \equiv \frac{c}{p} = \frac{c}{2L}. \tag{2.12}$$

We also define the *finesse* $\mathcal{F}$:

$$\mathcal{F} \equiv \frac{\Delta\nu}{\nu_{\text{FSR}}} \approx \frac{\pi\sqrt{|g_{rt}|}}{1-|g_{rt}|} = \frac{\pi\sqrt{r_1 r_2}}{1-r_1 r_2}, \tag{2.13}$$

where $\Delta\nu$ is the cavity full width at half maximum (FWHM)[6], and we assume $|g_{rt}| \approx 1$. The last expression is the finesse for a lossless FP cavity.

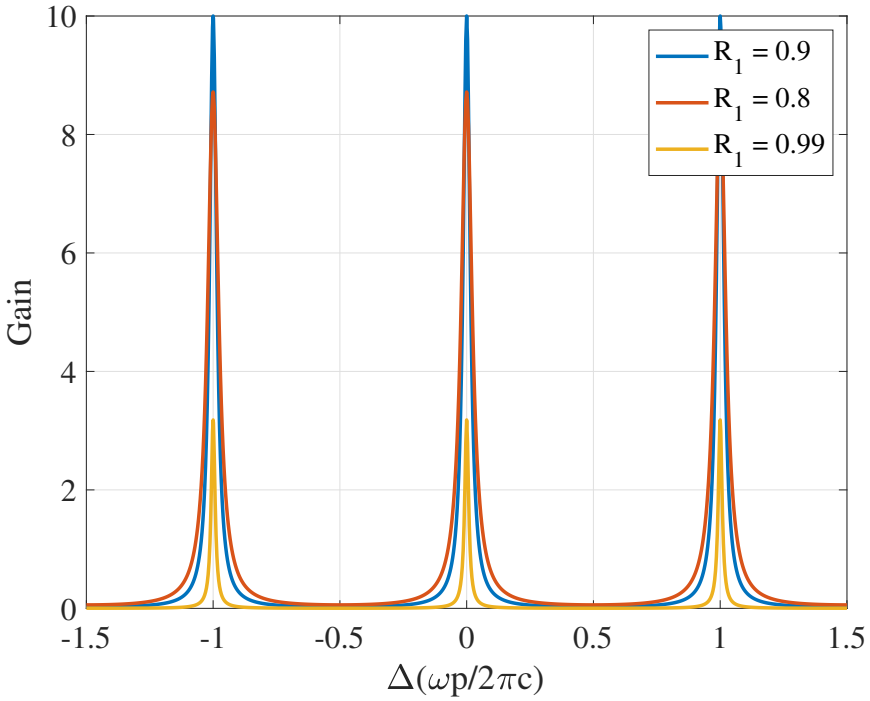

Figure 2-3: Cavity circulating power compared to input power for a number of different FP cavities assuming $R_2 = 0.9$.

[6]Note that this can be expressed in terms of half width at half maximum (HWHM) or full width at half maximum (FWHM); in this work we will try to be explicit.

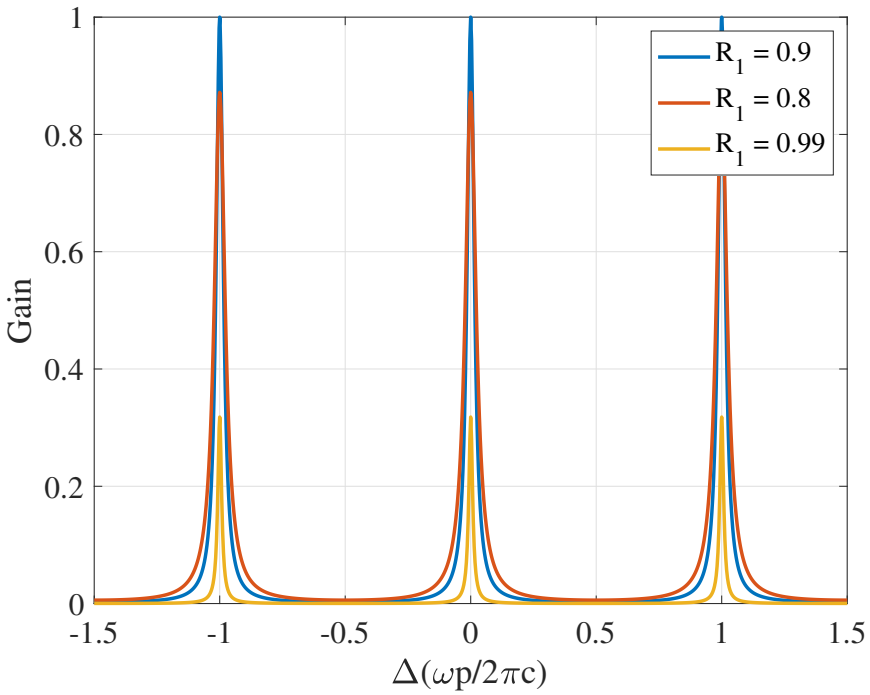

Figure 2-4: Transmission output power compared to input power for a number of different FP cavities assuming $R_2 = 0.9$. As expected, this plot nearly matches the circulating power plotted in figure 2-3, and it is only in the case when $R_1 = R_2$ that the cavity transmission is 1.

Figures 2-3 and 2-4 show the circulating and transmission gain of a passive FP cavity as a function of length (or frequency). Figures 2-5 and 2-6 show the magnitude and phase of the reflected beam from a FP cavity. Note the large phase shift in the reflected beam near resonance. This suggests a solution to the problem of maintaining the resonance condition: extracting this phase shift provides a control signal to keep the cavity on resonance.

2.2.1 Dynamics

Equations (2.9) assume a steady system, which occurs for perturbations at frequencies $\ll \Delta\nu$. A proper treatment of the dynamics of FP resonators [37] gives a storage time

$$\tau \equiv \frac{-1}{\nu_{\mathrm{FSR}} \ln(r_1 r_2)} \approx \frac{1}{\nu_{\mathrm{FSR}}(1 - r_1 r_2)}, \tag{2.14}$$

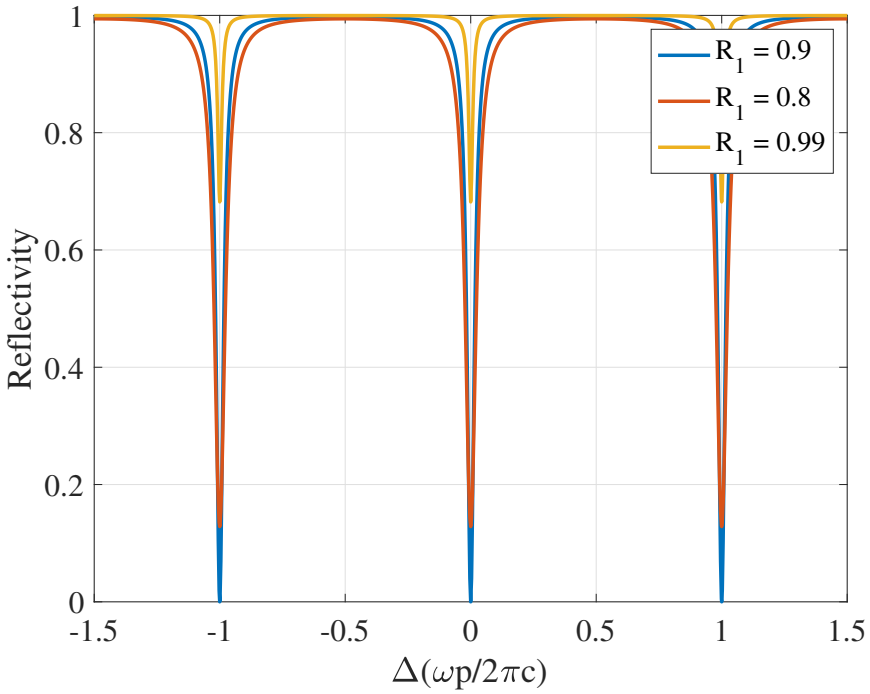

Figure 2-5: FP cavity reflected power compared to input power for a number of different optical cavities assuming $R_2 = 0.9$. Note in the case $R_1 = R_2$ the reflected power is zero.

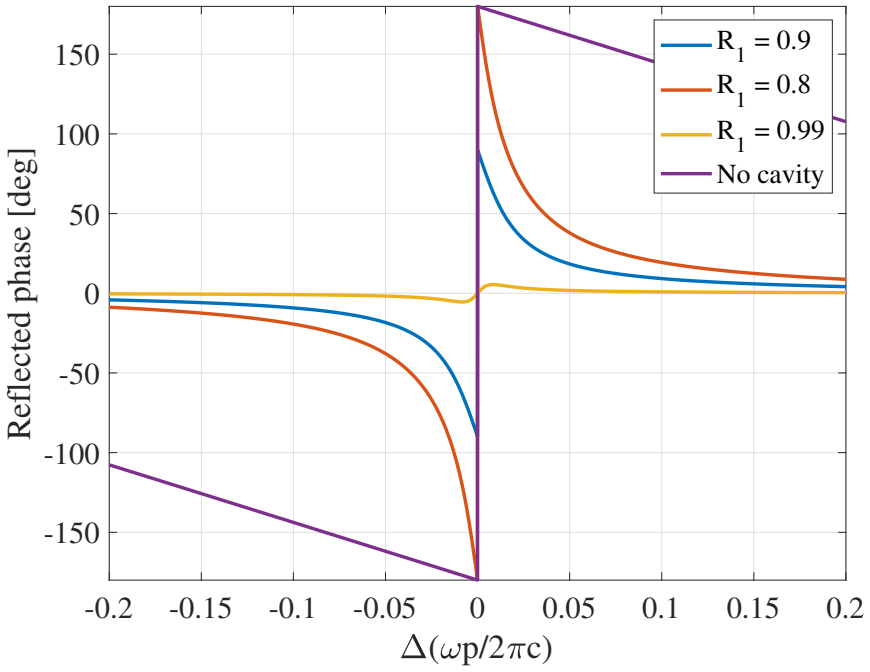

Figure 2-6: Cavity reflected phase for a number of different FP cavities assuming $R_2 = 0.9$. Also included is a trace for the phase shift for a single end mirror and no cavity. As expected, the cavity introduces a significantly larger phase shift than with just a single mirror.

for $r_1 r_2 \approx 1$. Additional complications appear for disturbances near the scale of a free spectral range (FSR), but for this work we will work in the limit where the above approximation is sufficient.

As we show in section A.2, any amplitude or (small) phase modulation at $\Omega \ll c/p$ can be written as a sum of a carrier at ω and sidebands at $\omega \pm \Omega$, where the relative phase of the sidebands determines the type of modulation. Thus, to determine the cavity response, we can consider how the carrier and sidebands appear in the cavity and write a superposition of the resulting beams (we will see an example of this in section 2.3.3). If the carrier is on resonance, the only terms that will oscillate at Ω will be proportional to

$$g_{\text{circ}}(\Omega) g^*_{\text{circ}}(0) \approx \frac{t_1^2}{1 - r_1 r_2} \frac{1}{1 - r_1 r_2 - i\Omega p/c} \tag{2.15}$$

This corresponds to a pole at angular frequency

$$\Omega_{\text{p}} = \frac{c}{p}(1 - r_1 r_2) \approx \frac{\nu_{\text{FSR}} \pi}{\mathcal{F}} = \gamma. \tag{2.16}$$

Hence, the cavity acts as a low-pass filter at frequency $\Delta\nu/2$.

2.2.2 Over- and Under-Coupled Cavities

As we can see in figure 2-6, the behavior of the reflected phase shift depends strongly on the chosen mirror parameters. This is hinted at in the expression for the reflected field in equation (2.9c), which is due to the interference of the promptly reflected field with the field leaking from the cavity. In the case of $r_1 > r_2$, when the promptly reflected beam dominates, the second term is smaller in magnitude and there is no sign flip. Similarly, when $r_2 > r_1$, most of the light reflected from the cavity must reflect off the end mirror and leak out of the cavity, and there is a large phase shift near resonance. The full version of equation (2.9c) simply includes the contribution of losses, but the intuition is the same.

When the second term is smaller than the first, the promptly reflected field dom-

inates this signal, and the phase shift is small; in this case, the cavity is said to be *under-coupled*. However, because the second term has a minus sign, when this term dominates, on resonance the reflected field undergoes a phase flip. In this case, the cavity is said to be *over-coupled*. In the edge case where these terms are equal, the cavity is *critically coupled* or *impedance matched*. In this case, the transmitted power is maximized, and as such the aLIGO mode cleaners are designed to be nearly critically coupled. Generally higher circulating powers and large phase shifts can be obtained for over-coupled cavities, so for this reason the LIGO arm cavities are over-coupled.

2.2.3 Transverse Modes

In the expression for the propagation phase, equation (2.10), we have assumed the incident light can be written as a plane wave. A more realistic description would be to write the beam as a superposition of transverse TEM_{mn} modes [2], with TEM_{00} as the fundamental Gaussian beam. These higher-order modes have an additional *Gouy phase* accumulation as compared to the fundamental mode

$$\psi_{mn} = (1 + m + n)\psi_{00}, \tag{2.17}$$

where ψ_{00} is the roundtrip Gouy phase shift of the fundamental TEM_{00} (Gaussian) mode. This means that the phase term in equation (2.10) should be modified to include this Gouy phase in addition to the usual propagation term. Because this phase depends on the transverse mode, in practice this means that the individual modes will resonate at different cavity lengths or laser frequencies. An example of this effect can be seen in figure 5-8 on page 136, in which the length of a cavity was scanned over a full FSR to look for the resonant conditions for these higher-order transverse modes[7].

[7]The TEM_{10} and TEM_{01} modes do not overlap in this figure because this scan was taken using a ring cavity, where the mirrors are not retroreflectors. In this configuration, there is an additional phase shift between vertical and horizontal transverse modes [38] and thus leads to different resonant frequencies.

2.2.4 Approximations

To simplify some of the above equations for the case of most optical cavities we will consider, we will assume we have a lossless FP cavity with high end-mirror reflectivity ($r_2 \approx 1$) held near resonance ($\phi \equiv 2\omega L/c \approx 2\pi n$). Using equation (A.3), we can write the change in phase on resonance as

$$\begin{aligned}\frac{\partial r_{\mathrm{cav},\varphi}}{\partial \phi} &= \Im\left\{\frac{r'_{\mathrm{cav}}}{r_{\mathrm{cav}}}\right\} = \Re\left\{\frac{-t_1^2 e^{i\phi}}{(e^{i\phi}-r_1)(r_1^2 - r_1 e^{i\phi} + t_1^2)}\right\} \\ &= -\frac{(1-r_1^2)}{(1-r_1)^2} = -\frac{1+r_1}{1-r_1},\end{aligned} \tag{2.18}$$

where we have used the relation $r_1^2 + t_1^2 = 1$. Further, the power gain on resonance is given by

$$|g_{\mathrm{circ}}|^2 = \frac{t_1^2}{(1-|g_{rt}|)^2} = \frac{1+r_1}{1-r_1}. \tag{2.19}$$

Thus, in this case, the reflected beam undergoes a phase shift that is $(1+r_1)/(1-r_1)$ times as large as for a single mirror[8], with a commensurate increase in circulating power. This is the advantage of using optical cavities. However, as we have already seen, this gain factor only occurs when we are near resonance.

If we make the further assumption that the input transmission is small ($t_1 \ll 1$), we can write the above in terms of the finesse $\mathcal{F}$ (2.13). We have

$$\mathcal{F} = \frac{\pi\sqrt{r_1}}{1-r_1} = \frac{\pi\sqrt{1-(1-r_1)}}{1-r_1} \approx \frac{\pi}{2}\frac{(1+r_1)}{1-r_1}. \tag{2.20}$$

Putting this all together, this means the light undergoes on average $2\mathcal{F}/\pi$ round trips before leaving the cavity.

We can also write the gain factors in a more convenient form. The circulating gain for the field (equation (2.9a)) is

$$g_{\mathrm{circ}} \approx \frac{it_1}{1 - r_1 \exp[-i\phi]} \approx \frac{it_1}{1 - r_1(1-i\phi)}. \tag{2.21}$$

[8]The extra 180° phase shift as seen when taking $r_1 \to 0$ is due to our phase convention for the transmission through a component.

Because the losses are small, we can also make the assumption that

$$r_1 = \sqrt{1 - t_1^2} \approx 1 - \frac{t_1^2}{2}. \tag{2.22}$$

Then the circulating gain can be written as

$$g_{\text{circ}} \approx \frac{it_1}{t_1^2/2 + i\phi - \cancel{i\phi t_1^2/2}} = \frac{2}{t_1} \frac{1}{1 + 2i\phi/t_1^2}, \tag{2.23}$$

where we have removed terms above order ϕ and t_1^2. The intracavity power gain is then

$$\frac{P_{\text{circ}}}{P_{\text{in}}} = |g_{\text{circ}}|^2 = \frac{4}{t_1^2} \frac{1}{1 + 4\phi^2/t_1^4} = \frac{4}{T_1} \frac{1}{1 + \delta^2}. \tag{2.24}$$

In this form, it's clear that the cavity response near resonance is Lorentzian with detuning $\delta \equiv 2\phi/t_1^2$ and linewidth (HWHM) $\gamma = t_1^2 c/4L$ (in units of angular frequency).

As we've seen, the benefits of using optical cavities only occurs in a narrow region near the resonance. The following section will deal with the generation of an error signal to ensure we remain near resonance.

2.3 Optical Cavity Sensing

As we have seen, the optical resonator is most sensitive near resonance. In addition to the main arm cavities, which are used to detect small changes in the phase of light down the arms that would indicate the passage of gravitational waves, a number of other optical cavities are used in aLIGO to keep the interferometer working. The main laser is frequency stabilized ("locked") to a stable reference cavity. Input and output mode cleaner cavities clean the spatial profile of the laser beam by locking to the TEM_{00} transverse mode (see discussion in section 2.2.3). A power-recycling cavity amplifies the input power and filters power fluctuations. A signal-recycling cavity increases the interferometer bandwidth. An optical parametric oscillator (OPO) uses the power amplification of a cavity to enhance nonlinear optical effects and generate squeezed light [17]. Even when a cavity is not being used directly for length sensing,

we must still generate error signals to ensure the cavity stays on resonance. We then have two related problems:

- How can one actually measure the phase shift near resonance?
- Given that the resonance features are significantly smaller than an optical wavelength, how does one keep the cavity near resonance?

The second item is straightforward to solve once we solve the first item and requires only that we keep $g_{rt} \approx 1$. Looking at the resonance condition (equation (2.11)), we see that we can achieve this by setting up a feedback loop [39] to adjust the cavity length via mirror actuators like piezoelectric stacks, magnetic coils, or radiation pressure [40]; change the laser frequency[9]; or add an additional phase shift inside the cavity.

In general, an ideal error signal is approximately linear in phase shift/cavity length/laser frequency near the resonance point. In the following sections, we will discuss some of the ways to generate a usable length error signal from optical resonators, including a novel technique that can be used for arbitrarily birefringent cavities. In all these cases, the signal scales as the input power, and thus, as we saw with the Michelson example, will generally have better shot noise sensitivity for higher input powers.

2.3.1 Transmission (Side-of-Fringe) Locking

The side of the cavity transmission peak shown in figure 2-4 is perhaps the simplest such signal that meets the requirements. Suppose we are on the left side of one transmission peak. Then an increase in power would indicate an increase in ϕ, and vice versa for a decrease in power. Using the output of a transmission photodetector, a feedback system can be used to ensure the transmitted power stays constant and the cavity stays "locked" nearly on resonance [41].

[9] It should be noted that lasers use optical resonators to select a lasing mode, and thus to change the laser frequency often these same properties of the lasing cavity must be changed.

There are a few issues that make this configuration non-ideal. First, any input power fluctuations will appear as length/frequency noise. Second, on resonance the circulating power is quadratic in cavity length and does not provide an error signal; thus, this method requires that the cavity be locked slightly off-resonance. This can change the response of the cavity, and in the case of weak mechanical oscillators, it can introduce optical spring [42] effects and can complicate the control problem. Additionally, because it is not a null measurement, it will have high levels of shot noise, even with a relatively stable input beam. For these reasons, this technique is generally not used for the most sensitive measurements, though it is simple to set up.

Note that an analogous technique can be used to lock to the side of a reflection dip as in figure 2-5, though this technique has the same issues as the transmission peak.

2.3.2 Dither Locking

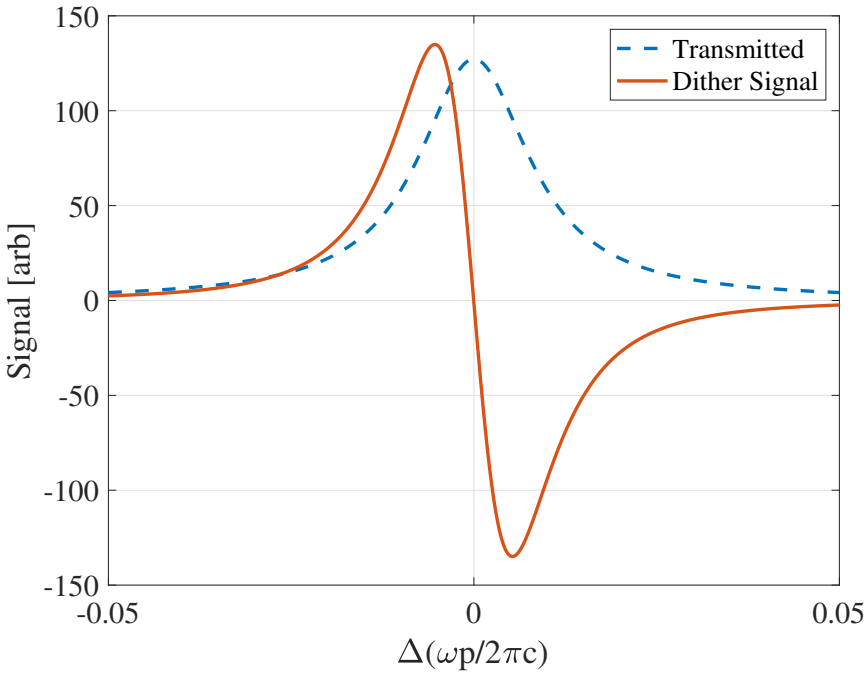

Figure 2-7: Dither locking signal

Another technique to acquire a lock signal are those that can sample the transmitted power as a function of laser frequency. This technique can be applied more

generally to find a local extrema and can be used, for example, to lock to an atomic resonance. In LIGO this technique is used to lock the OMC and to minimize the coupling of test mass angular motion to linear motion. This is also used in the experiment described in chapter 3 to find the optimal alignment.

Consider a system with independent variable x and output signal $T(x)$. As a concrete example, x could be either the laser frequency[10] or length of a cavity and $T(x)$ could be the transmitted power. For small variations about some x_0, we can Taylor expand the output signal as

$$T(x) = T(x_0) + T'(x_0)\Delta x + \frac{1}{2}T''(x_0)\Delta x^2 + \mathcal{O}\left(\Delta x^3\right), \tag{2.25}$$

where $\Delta x = x - x_0$ and $T'(x_0) \equiv \frac{\partial T}{\partial x}\big|_{x=x_0}$ (and similarly for $T^{(n)}(x_0)$). If we dither at angular frequency Ω_1 about x_0 with amplitude β[11], we can write this as

$$T(x_0 + \beta\cos\Omega_1 t) = T(x_0) + T'(x_0)\beta\cos\Omega_1 t + \frac{1}{2}T''(x_0)\beta^2\cos^2\Omega_1 t + \mathcal{O}(\beta^3). \tag{2.26}$$

When this signal is mixed with (multiplied by) a local oscillator at frequency Ω_2 and

[10] Generally this is done by modulating the phase, but since the phase and frequency are related by $\theta = \omega t$, at a single phase dither frequency Ω, this can be considered a frequency modulation:

$$\omega(t) = \frac{d\theta(t)}{dt} = \frac{d}{dt}(\beta\cos\Omega t) = -\beta\Omega\sin\Omega t$$

[11] The reader may wonder how "small" β must be given the Taylor approximation. For the case of locking a cavity, this assumption certainly holds for modulations smaller than the cavity linewidth.

some relative phase ϕ, we obtain

$$\begin{aligned}
&T(x_0+\beta\cos\Omega_1 t)A\cos(\Omega_2 t+\phi)= \\
&\quad A\cos(\Omega_2 t+\phi)T(x_0)+A\beta T'(x_0)\cos(\Omega_2 t+\phi)\cos\Omega_1 t \\
&\quad +\frac{1}{2}A\cos(\Omega_2 t+\phi)T''(x_0)\beta^2\cos^2\Omega t+\mathcal{O}(\beta^3) \\
&=A\cos(\Omega_2 t+\phi)T(x_0) \\
&\quad +\frac{A\beta}{2}T'(x_0)\left\{\cos\left[(\Omega_2-\Omega_1)t+\phi\right]+\cos\left[(\Omega_2+\Omega_1)t+\phi\right]\right\} \\
&\quad +\frac{A\beta^2}{4}T''(x_0) \\
&\quad \left\{\cos(\Omega_2 t+\phi)+\frac{1}{2}\cos\left[(\Omega_2-2\Omega_1)t+\phi\right]+\frac{1}{2}\cos\left[(\Omega_2+2\Omega_1)t+\phi\right]\right\} \\
&\quad +\mathcal{O}(\beta^3),
\end{aligned} \tag{2.27}$$

where we have used trigonometric identities to rewrite the second and third terms. When we use the same local oscillator to generate the dither signal ($\Omega_2=\Omega_1\equiv\Omega$) and then low-pass the signal to remove oscillations at Ω and higher frequencies, we have:

$$\text{error}=\frac{A\beta}{2}T'(x_0)\cos\phi. \tag{2.28}$$

By setting $\cos\phi=1$ to maximize this signal, the control signal is then proportional to the derivative of the output signal[12]. Figure 2-7 provides an example of this error signal for the over-coupled cavity presented in the previous section.

Thus, this method can be used to find extremal points of $T(x)$. In fact, if T is a function of multiple variables, each variable can be (de)modulated at a different frequency as long as those frequencies are separated sufficiently so that the beat frequency can be removed via the low-pass filter. For example, by dithering the angle of cavity optics, with each degree of freedom at a different frequency, the cavity alignment that maximizes transmission can be found with relative ease.

This technique has a few shortcomings. The dither frequency Ω must be low

[12]This same idea is used for lock-in amplifiers, where, with enough averaging, an injected signal can be extracted from far beneath the level of incoherent noise at the same frequency. Because the signal also depends on the amplitude of the local oscillator, increasing the drive of the local oscillator (LO) can also enhance the signal-to-noise ratio

enough for the signal $T(x)$ to respond approximately linearly with the drive. Practically this requires $\Omega \ll 2\pi\Delta\nu$. The low-pass cut-off frequency must also be sufficiently small to filter out signals at Ω, which limits the bandwidth of any control system. Additionally, because an external drive is required for this scheme, the effective range of the actuation for control is reduced. As a result, this technique can only be used for systems that do not require large gains and/or high-frequency actuation.

2.3.3 Introduction to Phase-Sensitive Detection: Pound-Drever-Hall

An improved locking scheme would still allow locking on resonance but without the bandwidth limitations of the dither scheme. The phase shift on reflection (figure 2-6) hints at a solution that does not require sampling the cavity power profile; if this phase shift can be converted to an optical power fluctuation, it can be read out with a photodetector. As suggested by the title of this work, the answer to this quandary is interferometry. By interfering the optical field that interacts with the cavity with a second optical field that is insensitive to the cavity resonance condition, this phase shift can be converted to an intensity signal. The canonical example of this technique is the Pound-Drever-Hall (PDH) technique [43, 44], which uses off-resonant frequency sidebands as the phase reference.

A full derivation and analysis of this technique can be found in [45], but here we review the salient points. Figure 2-8 shows the experimental setup for a PDH setup. An electro-optic modulator (EOM) modulates the phase of the laser beam at frequency Ω and modulation depth β, leading to a field incident on the cavity

$$\begin{aligned} E_{\text{in}} = E_0 e^{i(\omega t + \beta \sin \Omega t)} &\approx E_0 \left[J_0(\beta) + 2iJ_1(\beta) \sin \Omega t \right] e^{i\omega t} \\ &= E_0 \left[J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t} \right], \end{aligned} \quad (2.29)$$

where we have used the Jacobi-Anger expansion to write the exponential in terms of Bessel functions of the first kind and have assumed $|\beta| \ll 1$. In this form, it is clear

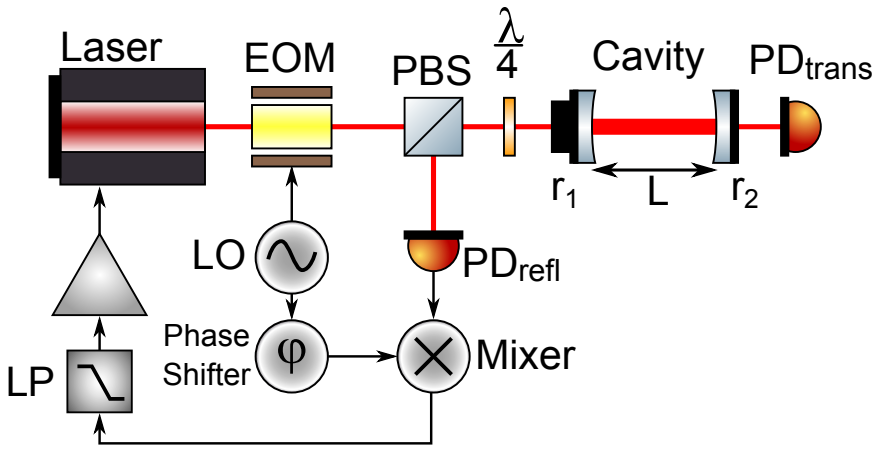

Figure 2-8: Setup for Pound-Drever-Hall (PDH) locking scheme. A local oscillator (LO) drives an electro-optic modulator (EOM) to modulate the incident laser beam phase and produce frequency sidebands. The light reflected from the cavity is measured and demodulated using a mixer and the same LO used for the phase modulation. A phase shifter and low-pass filter (LP) are used to generate the error signal.

that we have three beams incident on the cavity: a carrier at ω and two sidebands at $\omega \pm \Omega$. These sidebands will serve as the phase reference for the carrier, which will be on resonance for the cavity.

Since photodetectors measure power $P \propto |E|^2$, we can rewrite this in terms of the power of each component relative to the original beam power P_0:

$$P_c = J_0^2(\beta) P_0 \tag{2.30}$$

$$P_s = J_1^2(\beta) P_0 \tag{2.31}$$

$$P_c + 2P_s \approx P_0, \tag{2.32}$$

where $P_{c(s)}$ is the power in the carrier (sideband), and the last expression assumes the modulation depth is small so very little power is injected into higher-order sidebands. When this is reflected off the cavity, we have

$$\begin{aligned} E_{\text{refl}} = E_0[&r_{\text{cav}}(\omega) J_0(\beta) e^{i\omega t} \\ &+ r_{\text{cav}}(\omega+\Omega) J_1(\beta) e^{i(\omega+\Omega)t} - r_{\text{cav}}(\omega-\Omega) J_1(\beta) e^{i(\omega-\Omega)t}]. \end{aligned} \tag{2.33}$$

When this is put on a photodetector, the $|E|^2$ property will introduce interference

terms between each of these frequency components:

$$\begin{aligned} P_{\text{refl}} = P_c|r_{\text{cav}}(\omega)|^2 + P_s\{|r_{\text{cav}}(\omega+\Omega)|^2 + |r_{\text{cav}}(\omega-\Omega)|^2\} \\ +2\sqrt{P_c P_s}\{\Re[\theta(\omega,\Omega)]\cos\Omega t + \Im[\theta(\omega,\Omega)]\sin\Omega t\} \\ +(2\Omega \text{ terms}), \end{aligned} \tag{2.34}$$

where

$$\theta(\omega,\Omega) \equiv r_{\text{cav}}(\omega) r^*_{\text{cav}}(\omega+\Omega) - r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega-\Omega), \tag{2.35}$$

$\Re$ and $\Im$ refer to the real and imaginary parts, respectively, and * signifies complex conjugation. So we are left with a signal with DC components, components oscillating at Ω due to the beating of the sidebands with the carrier, and 2Ω terms due to the sidebands interfering with each other. As in the case of the dither lock in section 2.3.2, this signal will be demodulated at Ω so we will only worry about the terms oscillating at Ω.

Note that

$$\cos(\Omega t+\phi)(\sin\Omega t + \cos\Omega t) = \frac{1}{2}\left[\cos\phi - \sin\phi + \cos(2\Omega+\phi) + \sin(2\Omega+\phi)\right]. \tag{2.36}$$

Using this relation, when we demodulate the reflected signal (2.34) and low-pass to remove terms oscillating at Ω and above, we have

$$\text{error} = A\sqrt{P_c P_s}\left\{\Re[\theta(\omega,\Omega)]\cos\phi - \Im[\theta(\omega,\Omega)]\sin\phi\right\}. \tag{2.37}$$

Thus, by adjusting ϕ, we can extract the components of the signal oscillating in phase (real or "I") or in the quadrature phase (imaginary or "Q"), 90° out of phase with the LO[13].

For some intuition, we will consider two separate limiting cases. The first is when the modulation frequency is small ($\Omega \ll 2\pi\Delta\nu$). In this case, we can expand r_{cav} for

[13]It should be noted that often one can extract both quadratures with RF electronics and then "rotate" to the desired ratio once at DC frequencies.

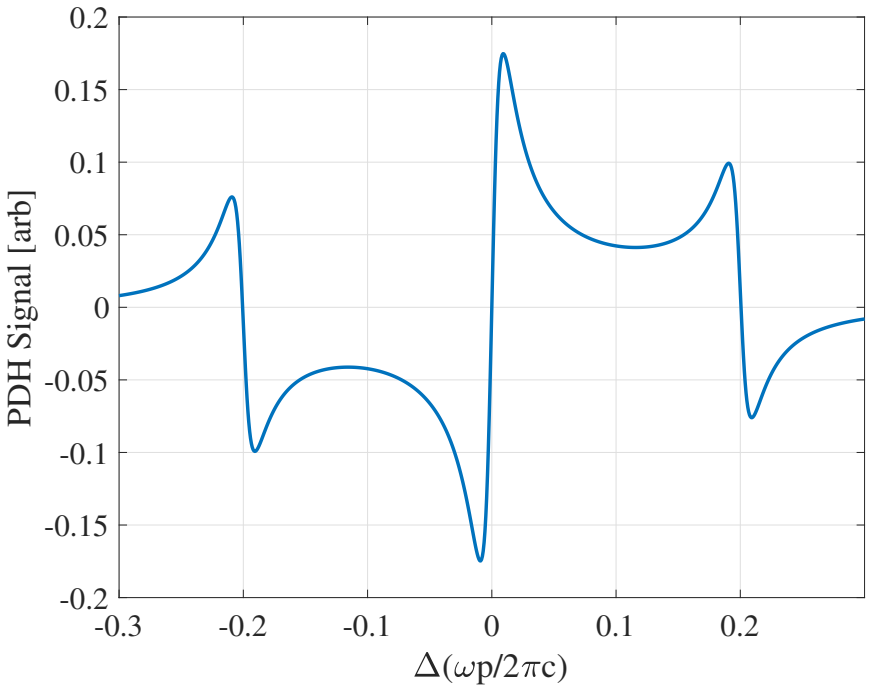

Figure 2-9: Pound-Drever-Hall (PDH) locking signal. The linear regions with opposite slopes from the central carrier are when the sidebands each go through the cavity resonance, in which case the original carrier acts as the phase reference.

small Ω around ω_0:

$$\begin{aligned}\theta(\omega_0,\Omega) &\approx r_{\text{cav}}(\omega_0)\left[r^*_{\text{cav}}(\omega_0)+\Omega r'^*_{\text{cav}}(\omega_0)\right]-r^*_{\text{cav}}(\omega_0)\left[r_{\text{cav}}(\omega_0)-\Omega r'_{\text{cav}}(\omega_0)\right]\\ &= \Omega\left[r_{\text{cav}}(\omega_0)r'^*_{\text{cav}}(\omega_0)+r^*_{\text{cav}}(\omega_0)r'_{\text{cav}}(\omega_0)\right]\\ &= \Omega\left.\frac{\partial|r_{\text{cav}}|^2}{\partial\omega}\right|_{\omega=\omega_0}=\Omega R'_{\text{cav}}(\omega_0),\end{aligned} \tag{2.38}$$

where $R_{\text{cav}}(\omega)$ is the cavity power reflectivity and $r'_{\text{cav}}(\omega_0)\equiv\frac{\partial r_{\text{cav}}}{\partial\omega}|_{\omega=\omega_0}$. This expression is purely real, so we are left with the $\cos\Omega t$ term in equation (2.34). As expected for small modulation frequencies, where the cavity has time to respond, we are essentially dithering the laser frequency and seeing how the cavity power reflection changes. This gives the same result as the error signal using the dither lock equation (2.28), where the function being sampled is $P_{\text{refl}}=P_0R_{\text{cav}}$ and we use equation (A.7) to approximate $\sqrt{P_cP_s}\approx P_0\beta/2$.

Usually PDH is used with high-frequency modulation ($\Delta\nu\ll\Omega/2\pi\ll c/2p$). In

this regime, we can assume that the sidebands are sufficiently far from resonance when the carrier is on resonance, so we can assume $r_{\rm cav}(\omega \pm \Omega) \approx r_1$. In this situation, for a given change in the resonance condition, the sideband reflection remains unchanged while the carrier goes through a large phase shift, which leads to a large change in power that we can observe. We then have

$$\theta(\omega, \Omega) \approx r_{\rm cav}(\omega) r_1 - r^*_{\rm cav}(\omega) r_1 = 2 r_1 \Im[r_{\rm cav}(\omega)]. \tag{2.39}$$

Because this term is imaginary, the $\sin \Omega t$ term in equation (2.34) will be most important, and we can extract the large phase shift of $r_{\rm cav}$ near resonance. Figure 2-9 shows this signal.

Note that this scheme requires a low-pass filter to cut out signal at the modulation frequency, which is significantly larger than the cavity linewidth. However, because the sidebands are always present, this technique can sense length fluctuations at much higher frequencies than the cavity pole and allows for higher bandwidths than a simple dither. Additionally, as with most phase-sensitive techniques, this method is first-order insensitive to input power fluctuations since $\Im\{r_{\rm cav}\} = 0$ exactly on resonance. Finally, the sidebands are pushed out to high frequencies, where shot noise is usually the only dominant source of noise. As a result of these characteristics, this technique can provide a sensitive readout of the resonance condition and allow for high control bandwidths.

While the math for these two cases is similar, they can be thought of as qualitatively different. In the case of low modulation frequency, the cavity (power) reflectivity is sampled to find the resonance condition; this is exactly what a dither lock is[14]. For high modulation frequencies, the frequency sidebands are used as a phase reference for the carrier beam, which obtains a large phase shift near resonance.

[14]While this technique samples the power, the cavity output power is itself determined by the relative phase of the input beam to the light already inside the cavity. Thus, even this dither technique is a form of interferometry!

2.3.4 Polarization Spectroscopy

We provide an additional example of a phase-sensitive locking method with a novel generalization of Hänsch and Couillaud (HC) polarization spectroscopy [46], which uses different polarizations as the phase references. This technique, as well as an experimental verification, was published in [47]. The HC method uses a polarizer in the resonator to produce a polarization-dependent resonance condition. Variations of this technique have been demonstrated in free space cavities by introducing a birefringent crystal into the cavity [48], using a non-planar ring cavity [49], using a triangular cavity [50], or using the birefringence of dielectric mirrors [51]. A special case of this technique has been demonstrated in fiber for strain sensing applications [52]. We generalize these locking techniques and describe a polarization spectroscopy locking method that introduces no additional losses, is simple to implement, and requires no knowledge of the cavity birefringence. These properties make this method particularly attractive for systems with inherent or unknown birefringence, such as whispering-gallery-mode (WGM) resonators [53, 54], cavities formed with crystal-coated mirrors [55, 56], and optical-fiber-based applications.

The resonance properties [57] and the polarization effects of fiber rings have been investigated previously [58, 59, 60]. Small stresses and imperfections in the core of the fiber produce polarization-dependent phase shifts that are rarely known *a priori*. This inherent unknown birefringence make this locking technique well suited to use with fiber ring resonators. In addition to the ease of alignment and mode-matching that comes with using single-mode optical fibers, fiber rings allow for long cavities and narrow linewidths in compact packages. Traditional methods such as PDH can be used to lock to fiber rings [61], but there are techniques unique to fiber optics, such as using Rayleigh scattering to provide optical feedback [62]. In contrast, this locking scheme requires no modulation or demodulation, which allows for simple electronics and a high-bandwidth error signal. This method can be used to produce compact prestabilized lasers for use in fiber-based systems that require stable coherent sources, such as telecommunication systems [63], lidar [64], fiber gyroscopes [65], and other

fiber-based sensors, such as strain sensors [52].

Polarization in optical resonators

A birefringent cavity will impart a different phase delay to each polarization component and may alter the polarization of the input light. Additionally, cavity losses may be polarization dependent, which can create additional interesting polarization effects. These polarization effects can be more easily understood by representing the electric field of fully polarized light as a two-component complex vector known as a Jones vector [38]. Conventionally, this vector is written using linear horizontal and vertical polarizations as the basis, so a polarized plane light wave at time t a distance z along the propagation axis will have the form

$$\vec{E} = \begin{pmatrix} E_{\mathrm{H}} \\ E_{\mathrm{V}} \end{pmatrix} = \begin{pmatrix} A_{\mathrm{H}} \\ A_{\mathrm{V}} e^{i\phi} \end{pmatrix} e^{i(\omega t - kz)}, \tag{2.40}$$

where A_{H} and A_{V} are the amplitudes of each polarization component, ϕ is the phase difference between each component, and ω and k are the angular frequency and wavenumber of the light, respectively. The last exponential is usually dropped since global phases have no effect on the polarization. In this framework, cavity losses, polarization transformations, and phase changes can be described by a 2×2 Jones matrix that acts on the polarization vector.

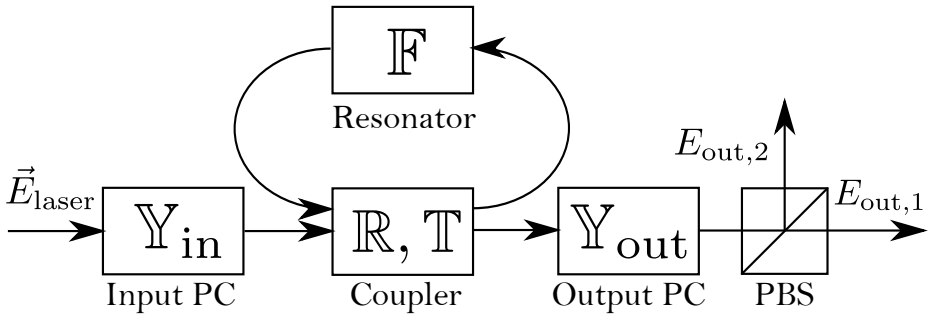

Figure 2-10: Block diagram of an optical experiment compatible with our locking method. Polarization controllers (PCs) on the input and output of the resonator are given by the Jones matrices $\mathbb{Y}_{\mathrm{in}}$ and $\mathbb{Y}_{\mathrm{out}}$. Matrices $\mathbb{R}$ and $\mathbb{T}$ represent reflection and transmission of the coupler, respectively, and $\mathbb{F}$ represents propagation through the resonator. A polarizing beam splitter (PBS) follows at the end.

In what follows, we will consider the polarization effects of a fiber ring resonator with a single input coupler, but this method is completely generalizable to any birefringent cavity. The amplitude transmittance and reflectance matrices of the input coupler are labeled $\mathbb{T}$ and $\mathbb{R}$, respectively. Let $\mathbb{F}$ be the matrix that represents the polarization transformation as the light propagates in the cavity (see Fig. 2-10). Included in this matrix is information about the overall (frequency-dependent) phase picked up and any losses incurred as light propagates from the input coupler through the cavity and back. The total cavity amplitude reflectivity matrix $\mathbb{R}_{\text{cav}}$, defined by $\vec{E}_{\text{refl}} = \mathbb{R}_{\text{cav}}\vec{E}_{\text{in}}$, is given by

$$\mathbb{R}_{\text{cav}} = \mathbb{R} - \mathbb{T}\mathbb{F}\left(\mathbb{I} - \mathbb{R}\mathbb{F}\right)^{-1}\mathbb{T}, \tag{2.41}$$

where $\mathbb{I}$ is the identity matrix. Since $\mathbb{R}\mathbb{F}$ represents one round trip of the resonator, the two eigenvectors of this matrix are the eigenpolarizations of the resonator. For positive real eigenvalues of $\mathbb{R}\mathbb{F}$, light will resonate.

We assume that the eigenvectors of $\mathbb{R}\mathbb{F}$ are also eigenvectors of $\mathbb{T}$ and $\mathbb{R}$, and therefore of $\mathbb{R}_{\text{cav}}$, for all frequencies ω. This assumption holds for input couplers with polarization-independent properties. For birefringent input couplers, this assumption is valid when the eigenpolarizations of the coupler match those of the cavity. For fibers, the polarization effects of the couplers are often much smaller than the effects due to birefringence in the fiber itself [66], so this is a reasonable assumption for fiber rings.

In this case, the eigenvalues of $\mathbb{R}_{\text{cav}}$ are given by, as in equation (2.9c),

$$r^j_{\text{cav}}(\omega) = r_j - \frac{t_j^2 f_j(\omega)}{1 - r_j f_j(\omega)}, \tag{2.42}$$

where t_j, r_j, and f_j are the eigenvalues of eigenpolarization j under $\mathbb{T}$, $\mathbb{R}$, and $\mathbb{F}$, respectively. In general, $f_j = \alpha_j e^{-ik_j p}$, where $k_j = 2\pi n_j \nu/c$ is the wavenumber for each eigenpolarization, n_j is the effective index of refraction for each eigenpolarization, ν is the frequency of light, p is the distance the light travels in one round trip of the cavity, and α_j is a complex term that accounts for losses in the cavity and additional

phase shifts not due to propagation.

In general, the wavenumbers for different eigenpolarizations will not be equal, so there will be a difference in the phase accumulated by each eigenpolarization after one round trip, resulting in two longitudinal modes that are on resonance for different ω. As one eigenpolarization moves through a resonance, it undergoes a large phase shift on reflection, while the other eigenpolarization does not. In this manner, the latter can be used as a phase reference for the former. This is easiest to implement when the resonances are shifted sufficiently far in frequency. Individual resonance peaks are well separated for round trip phase differences between eigenpolarizations $\Delta\theta$ that satisfy

$$\pm\Delta\theta \bmod 2\pi > \frac{2\pi\Delta\nu}{\mathrm{FSR}}, \tag{2.43}$$

where $\Delta\nu$ is the full width at half maximum of the cavity resonance and FSR is the free spectral range. This condition can be satisfied even in cavities with small birefringence so long as the cavity has a large finesse ($\propto \frac{\mathrm{FSR}}{\Delta\nu}$).

The error signal

To use the birefringence of the resonator to produce an error signal, light traveling to the resonator must first be put into the correct input polarization, which can be done using a polarization controller (PC). There are many ways to implement such a device, but a common arrangement consists of a quarter-wave plate, a half-wave plate, and a second quarter-wave plate, each of which can be rotated independently. These polarization controllers can map an arbitrary input polarization state to any other polarization state [67]. After reflecting from the resonator, light goes through another PC before propagating to a polarizing beam splitter (PBS) (figure 2-10). A photodetector (PD) at each output of the PBS measures the power. With the correct polarization control settings before and after the cavity, the difference in power at the two output ports of the PBS produces an error signal for the resonance condition of the cavity. Expressed in the Jones matrix formalism the error signal Δ is proportional

to

$$\Delta \propto |E_{\text{out},2}|^2 - |E_{\text{out},1}|^2, \tag{2.44}$$

where

$$\begin{aligned} E_{\text{out},1} &= \begin{pmatrix} 1 \\ 0 \end{pmatrix} \cdot \mathbb{Y}_{\text{out}} \mathbb{R}_{\text{cav}} \mathbb{Y}_{\text{in}} \vec{E}_{\text{laser}} \\ E_{\text{out},2} &= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \cdot \mathbb{Y}_{\text{out}} \mathbb{R}_{\text{cav}} \mathbb{Y}_{\text{in}} \vec{E}_{\text{laser}} \end{aligned} \tag{2.45}$$

are the respective projections of the final electric field amplitude onto the PBS polarization basis states. The proportionality constant will depend on the PD gain. In the above equations, $\mathbb{Y}_{\text{in (out)}}$ is the Jones matrix of the input (output) PC and $\mathbb{R}_{\text{cav}}$ is the cavity amplitude reflectivity matrix given by Eq. (2.41).

We further assume the cavity eigenpolarizations are orthogonal. This will be true for cavities with polarization-independent losses, in which $\mathbb{R}$, $\mathbb{T}$, and $\mathbb{F}$ are all unitary matrices multiplied by a constant loss term [59]. The eigenpolarizations will also be orthogonal for cavities in which one polarization mode is completely extinguished in the cavity, as in the original HC method [46]. In practice, this approximation holds true for most single-mode fiber rings [59].

Let $\vec{E}_a$ and $\vec{E}_b$ be the normalized eigenpolarizations of $\mathbb{RF}$. The input PCs are set so that the input polarization is an equal superposition of these eigenpolarizations:

$$\vec{E}_{\text{in}} = \mathbb{Y}_{\text{in}} \vec{E}_{\text{laser}} = \frac{E_0}{\sqrt{2}} \left(\vec{E}_a + e^{i\gamma} \vec{E}_b \right), \tag{2.46}$$

where E_0 is the amplitude of the electric field and γ is the phase difference between eigenpolarization components[15]. The reflected light will be in the polarization state

$$\begin{aligned} \vec{E}_{\text{refl}} = \mathbb{R}_{\text{cav}} \vec{E}_{\text{in}} &= \frac{E_0}{\sqrt{2}} \left(\mathbb{R}_{\text{cav}} \vec{E}_a + e^{i\gamma} \mathbb{R}_{\text{cav}} \vec{E}_b \right) \\ &= \frac{E_0}{\sqrt{2}} \left(r^a_{\text{cav}}(\omega) \vec{E}_a + e^{i\gamma} r^b_{\text{cav}}(\omega) \vec{E}_b \right). \end{aligned} \tag{2.47}$$

[15] Note this is not the same as the cavity linewidth γ

Because the polarization controllers can map an arbitrary input state to any given output polarization, there exists an arrangement of the output polarization controller that maps one eigenpolarization to an equal superposition of the PBS polarizations:

$$\vec{E}_a \mapsto \mathbb{Y}_{\text{out}}\vec{E}_a = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\ e^{i\delta}\end{pmatrix} \tag{2.48}$$

in the PBS basis for some δ[16]. The polarization controller produces a lossless, and therefore unitary, transformation on the Jones vector, so the other eigenpolarization will be mapped to a final polarization orthogonal to this with some relative phase shift ϕ:

$$\vec{E}_b \mapsto \mathbb{Y}_{\text{out}}\vec{E}_b = \frac{e^{i\phi}}{\sqrt{2}}\begin{pmatrix}-1\\ e^{i\delta}\end{pmatrix}. \tag{2.49}$$

The final polarization state in the PBS basis is then

$$\vec{E}_{\text{out}} = \mathbb{Y}_{\text{out}}\vec{E}_{\text{refl}} = \frac{E_0}{2}\left(r^a_{\text{cav}}(\omega)\begin{pmatrix}1\\ e^{i\delta}\end{pmatrix} + e^{i(\gamma+\phi)}r^b_{\text{cav}}(\omega)\begin{pmatrix}-1\\ e^{i\delta}\end{pmatrix}\right) \tag{2.50}$$

and the sum and difference of the photodiode signals are proportional to

$$|E_{\text{out},1}|^2 + |E_{\text{out},2}|^2 = \frac{E_0^2}{2}\left(|r^a_{\text{cav}}(\omega)|^2 + |r^b_{\text{cav}}(\omega)|^2\right) \tag{2.51}$$

and

$$|E_{\text{out},2}|^2 - |E_{\text{out},1}|^2 = E_0^2\,\Re\left\{(r^a_{\text{cav}}(\omega))^* r^b_{\text{cav}}(\omega)e^{i(\gamma+\phi)}\right\}, \tag{2.52}$$

respectively. Equation (2.52) forms the error signal.

The similarity to the reflection coefficients important for PDH in equation (2.35) should be noted, and the idea is the same: for widely separated resonances, near the resonance of one eigenpolarization the imaginary part of the cavity reflectivity undergoes a sharp change, while the reflectivity for the other eigenpolarization changes very little. By adjusting the additional phase term in equation (2.52), the steep imaginary

[16]Note that this is not the same δ as the cavity detuning in section 2.2.4

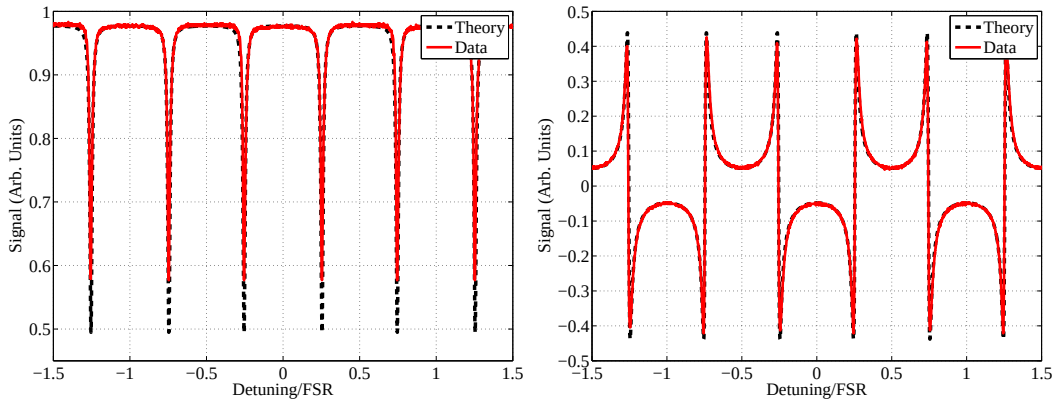

Figure 2-11: Observed and modeled (a) sum and (b) difference of the two output ports of the PBS for a 20 m fiber ring resonator setup. The difference signal provides the error signal.

part of the individual reflection coefficients can be extracted. For widely separated resonances, this occurs for $\gamma + \phi \approx \pm\pi/2$. The output polarization controllers vary ϕ to produce an ideal error signal; therefore, it is not necessary to know or control γ. This means that the input polarization state can be chosen without regard to the relative phases of the eigenpolarizations, and that only the sum signal is needed to set the input PCs.

This is a general method for producing an error signal for a birefringent resonator with two longitudinal modes; for a chosen input polarization satisfying equation (2.46), one can always find a physically realizable Jones matrix $\mathbb{Y}_{\text{out}}$ that produces an error signal. Figure 2-11 shows the theoretical and measured cavity signals from a 20 m fiber ring resonator; the full details of this setup can be found in [47].

Fundamentally, this method is limited only by shot noise. As we will see in section 5.1, in fiber systems stimulated Brillouin scattering [68] puts a limit on the input power, which prevents the reduction of shot noise by turning up the optical power. However, in practice this method will be limited by acoustic/vibrational noise and thermorefractive effects in the fiber.

2.3.5 Recap of Locking Techniques

We have reviewed two phase-sensitive methods for locking optical resonators, however there are a number of other phase-sensitive techniques. Higher-order modes can be used as reference beam [69, 70], with the caveat that quadrant photodetectors (QPDs) or bullseye PDs must be used, as otherwise there is no interference between orthogonal transverse modes. The reference beam could also be supplied from an external source, as in homodyne detection [34]. The differential arm degree of freedom of the aLIGO detectors uses a form of homodyne detection called DC readout [23]. Because these techniques are usually locked at a null point, they are also insensitive to input intensity fluctuations to first order. Finally, angular displacements can change the relative phase of odd higher-order modes relative to the fundamental, so similar techniques can be used to sense the angular degrees of freedom [71, 72, 73],

These phase-sensitive techniques provide for some of the most sensitive measurements of optical phase arising from variations in laser frequency or mirror position. In general, the signal generated from these techniques will scale as the geometric mean of carrier and reference power, and thus the shot-noise sensitivity will improve with higher input powers. In all cases, the E^2 property of the photodetector will give a measured signal with a term similar to $(r^a_{\rm cav})^* r^b_{\rm cav}$ for carrier a and reference b (or vice versa) that allows for extraction of the steep phase shift near resonance.

In the following section, we will provide a formal introduction to quantum noise and see that shot noise is not the only relevant quantum noise term in the context of optical components that are free to move.

2.4 Quantum Optomechanics and the Standard Quantum Limit

The LIGO detectors, with suspended test masses, are an example of optomechanical systems, where light couples to the motion of a mechanical oscillator via radiation pressure (figure 2-12). While such systems are useful for detecting gravitational waves,

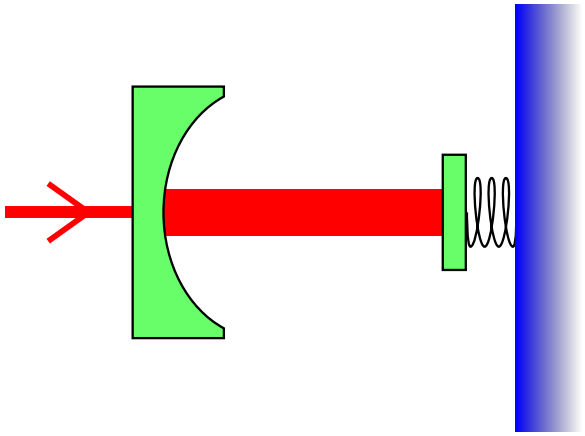

Figure 2-12: An optomechanical setup.

they are interesting to study in their own right. Due to the low entropy of laser light, these systems can be used to cool mechanical oscillators to the motional ground state, which allows study of the quantum mechanical properties of macroscopic objects. Conversion of the quantum properties of light from optical frequencies to lower mechanical frequencies may allow for quantum storage mediums. See a recent review in [74].

In contrast to many other optomechanical experiments, in LIGO and the other experiments discussed here, we are working at frequencies far above the mechanical resonance and so we will use the free mass approximation for these oscillators. Additionally, due to the low mechanical frequencies, we can treat the mechanical oscillator as a classical oscillator. However, due to the sensitivity of these instruments, the quantum mechanical fluctuations of the light fields used to probe the mirror position are not negligible so we will use a semiclassical treatment.

We can write the electric field in terms of the raising operator a_ω for a given mode [34]:

$$E(t)^{(+)} = \int_0^\infty \sqrt{\frac{2\pi\hbar\omega}{\mathcal{A}c}} a_\omega e^{-i\omega t} \frac{d\omega}{2\pi}, \tag{2.53}$$

where $\mathcal{A}$ is the effective mode area, $E(t) = E^{(+)}(t) + E^{(-)}(t)$ and $E^{(-)} = \left(E^{(+)}\right)^*$. We can write this using a two-photon formalism developed by Caves and Schumaker [75, 76]. Let's define the annihilation operator for sidebands about some carrier frequency

ω_0 as

$$a_+ \equiv a_{\omega_0+\Omega} \qquad\qquad a_- \equiv a_{\omega_0-\Omega}. \tag{2.54}$$

This allows us to write the total electric field in terms of these sidebands:

$$E^{(+)} = \sqrt{\frac{2\pi\hbar\omega_0}{\mathcal{A}c}} e^{-i\omega_0} \int_0^\infty \left(a_+ e^{-i\Omega t} + a_- e^{i\Omega t}\right) \frac{d\Omega}{2\pi}, \tag{2.55}$$

where we are only considering frequencies $\Omega \ll \omega_0$ and have expanded the integral limit to ∞ for notational convenience.

Let's consider the operators representing amplitude and phase fluctuations, respectively, at Ω, given by

$$a_1 \equiv \frac{a_+ + a_-^\dagger}{\sqrt{2}} \qquad\qquad a_2 \equiv \frac{a_+ - a_-^\dagger}{\sqrt{2}i}. \tag{2.56}$$

Note that these each depend on Ω but we have suppressed this for notational convenience. The commutation relations for these operators is given by

$$\left[a_1, a_2^\dagger\right] = -\left[a_2, a_1^\dagger\right] = i2\pi\delta(\Omega - \Omega'), \tag{2.57}$$

with all others zero. Equation (2.57) implies an uncertainty relation between these two variables.

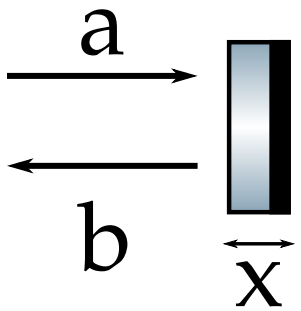

Figure 2-13: Field operators for light incident on and reflected from a perfect reflector free mass. Because we assume a perfectly reflecting mirror, we can ignore the fields on the other side of the mirror.

Consider a free mass perfect reflector with mass m and incident power P and

wavevector $k = \omega/c$. We can write the output field b in terms of the input field a. We must account for the fact that the light impinging on the mirror will introduce a fluctuating radiation pressure force. The fluctuating power (in spectral density units) is given by

$$\delta P = 2\sqrt{P}\sqrt{\hbar\omega}a_1. \tag{2.58}$$

This produces a fluctuating radiation pressure force which turns into a changing displacement (back action)

$$\delta F = \frac{2\delta P}{c} \qquad\qquad \delta x_{\mathrm{BA}} = \frac{-\delta F}{m\Omega^2}. \tag{2.59}$$

This will induce a phase shift on the reflected field, so the reflected phase quadrature operator can be written as

$$b_2 = 2k\sqrt{\frac{P}{\hbar\omega}}(x + \delta x_{\mathrm{BA}}) + a_2, \tag{2.60}$$

where x represents the motion of the mirror due to other forces. In matrix form, we have

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -\mathcal{K} & 1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} + \begin{pmatrix} 0 \\ 2k\sqrt{\frac{P}{\hbar\omega}} \end{pmatrix} x, \tag{2.61}$$

where we have defined

$$\mathcal{K}(\Omega) = \frac{8kP}{m\Omega^2 c}. \tag{2.62}$$

Thus, the amplitude fluctuations are converted to phase fluctuations by the motion of the mirror; this effect is known as quantum radiation pressure noise (QRPN). First shown rigorously to affect the sensitivity of GW interferometers by Caves [77], as of this writing this noise has not yet been directly measured in any GW detectors. However, it is expected to limit the design sensitivity of aLIGO (figure 1-4).

We can write this noise in length units. Neglecting the contribution from other

mirror motion, the measured noise spectrum[17] is given by

$$\begin{aligned} G_{xx}(\Omega) = \Delta x^2(\Omega) &= \frac{\hbar\omega}{4k^2P}\left\langle|\Delta b_2|^2\right\rangle \\ &= \frac{\hbar\omega}{4k^2P}\left(\left\langle|\Delta a_2|^2\right\rangle + |\mathcal{K}|^2\left\langle|\Delta a_1|^2\right\rangle - 2\Re\left\{\left\langle\mathcal{K}a_1a_2^\dagger\right\rangle_{\text{sym}}\right\}\right), \end{aligned} \tag{2.63}$$

where we have defined $\langle RS\rangle_{\text{sym}}$ for operators R, S as

$$\langle RS\rangle_{\text{sym}} \equiv \langle(RS)_{\text{sym}}\rangle = \left\langle\frac{1}{2}(RS+SR)\right\rangle, \tag{2.64}$$

and

$$|\Delta R|^2 \equiv (\Delta R\Delta R^\dagger)_{\text{sym}}. \tag{2.65}$$

In the case of uncorrelated input quadratures, which is true for coherent states of light[18], we have

$$\left\langle|\Delta a_2|^2\right\rangle = \left\langle|\Delta a_1|^2\right\rangle = \frac{1}{2} \tag{2.66}$$

$$\left\langle\mathcal{K}a_1a_2^\dagger\right\rangle_{\text{sym}} = 0, \tag{2.67}$$

so

$$G_{xx}^{\text{quantum}}(\Omega) = \frac{G_{xx}^{\text{SQL}}(\Omega)}{2}\left[\frac{1}{\mathcal{K}(\Omega)} + \mathcal{K}(\Omega)\right] \geq G_{xx}^{\text{SQL}}(\Omega), \tag{2.68}$$

where we have defined the free-mass standard quantum limit (SQL) [78]

$$G_{xx}^{\text{SQL}}(\Omega) = \frac{2\hbar}{m\Omega^2}. \tag{2.69}$$

This represents the lowest noise level attainable with coherent states of light (i.e., a_1 and a_2 are uncorrelated). The first term in equation (2.68) is the shot noise and, as expected, scales as $P^{-1/2}$. The second term is QRPN and scales as $P^{1/2}$. It should

[17] We are sweeping a few things under the rug here. Note that a_i is not Hermitian. Because we measure these quadratures with interferometry, we ultimately measure the (Hermitian) symmetric version of these operators, as defined in equation (2.64); see [75] for a more complete discussion.

[18] The cross-term in equation (2.63) suggests that with a proper correlation between the input phase and amplitude quadratures, we can dip below the minimum noise from coherent states. This is the principle behind the injection of squeezed light.

be noted that equality is only attained for $\mathcal{K} = 1$ (i.e., the shot noise and QRPN contributions are equal), which is only true for a single Ω for a given circulating power.

Qualitatively, we can think of this noise term in the following way: the arrival of photons at the photodetector is a Poisson process. For a given observation time, an average of n photons arrive at the detector, with variance $(\Delta n)^2 = n$. However, our signal scales as the power detected, which scales as n. Thus, our shot-noise-limited sensitivity will scale as $1/\sqrt{n} \propto 1/\sqrt{P}$, as we have already seen.

However, this variance is also important when we try to measure the position of a free mass. Upon reflection from a mirror, each photon will provide an impulse of $2\hbar\omega/c$. This will produce a fluctuating force on the mirrors that will look like a GW. By the same argument as above, the variance of this force will scale as n, and so the noise due to this quantum radiation pressure noise (QRPN) is proportional to $\sqrt{P}$.

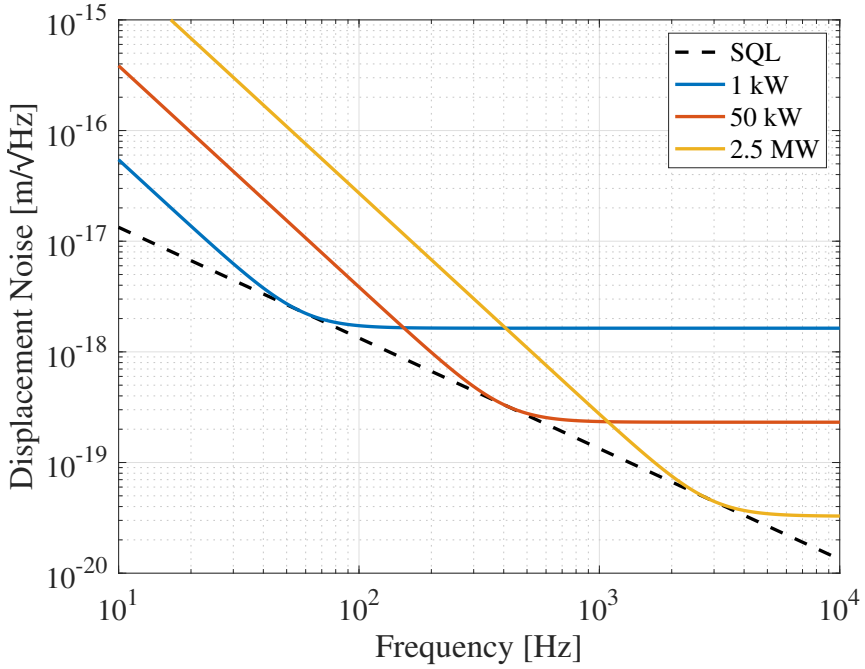

Figure 2-14: Standard quantum limit (SQL) and quantum-noise-limited sensitivity for a 0.3 g test mass optic for varying power impinging on the mirror. At high frequencies, the noise is limited by shot noise, which has a flat spectrum. At low frequencies, the quantum radiation pressure noise (QRPN) dominates and has a f^{-2} spectrum.

Figure 2-14 shows the quantum noise for different powers impinging on a free mass

compared to the SQL[19]. While we have only considered the case for a single free mass, these general results will hold for a LIGO-like interferometer [80, 79] with a $\mathcal{K}$ term that depends on the cavity parameters and uses the reduced mass of the test masses.

2.5 Conclusion

We've reviewed the properties of optical resonators and have shown how they can improve the sensitivity of interferometric detectors. We have also reviewed a few different techniques to generate an error signal that can be used to keep such cavities on resonance. We have shown that the shot-noise-limited sensitivity can be improved with increasing laser power. Due to the generalized Heisenberg uncertainty principle, this leads to an increase in mirror displacement, reducing the interferometer sensitivity. While this seems straightforward in principle to reach a regime dominated by this force, in practice there are a number of other noise sources that can limit the interferometer sensitivity where QRPN could dominate. In the remainder of this thesis, we will discuss a few experiments that work to bring this regime closer to reality.

[19]The observant reader may wonder why the shape of the quantum noise here does not match that for the aLIGO noise curve in figure 1-4. This is due to two factors: first, above the cavity pole the response of the interferometer to displacements goes as $\sim f^{-1}$. Thus, the flat shot noise turns into an upturned spectrum in units of length. Second, the presence of a signal recycling cavity changes the shape of the shot noise in the full interferometer [79], though we will not discuss this effect in this work.

Chapter 3

Radiation Pressure with a Gram-scale Oscillator

This chapter describes a FPMI with gram-scale end mirrors that has been used to study the radiation-pressure-dominated regime of a mechanical oscillator. Previous work [81] had determined that the ultimate sensitivity of this instrument was limited by the thermal noise of the end mirrors. This chapter describes work to mitigate this thermal noise and attempts to reach the QRPN-dominated regime.

3.1 Background

The optical layout is shown in figure 3-1. This setup has been used to probe the dynamics of a suspended optic system when dominated by radiation pressure effects: it has been used to demonstrate parametric instabilities and an optical spring [82], producing an all-optical stable optical spring [83], and optical dilution and cooling of macroscopic mirror to the mK regime [84]. Recently this system was used to measure optomechanically induced transparency (OMIT) [85] at sub-Hz linewidth [86, 87], the narrowest yet observed.

This setup also makes an excellent testbed to study the effects of QRPN in geometries similar to advanced GW detectors. The vast majority of optomechanics experiments use the mechanical resonance to enhance the effect of radiation pressure

fluctuations. Such experiments are only able to reach the QRPN-limited regime in a narrow frequency range around the mechanical resonance, as in the first observation of QRPN in a macroscopic object [88]. However, GW detectors use suspended optics with resonant frequencies around 1 Hz as the mechanical oscillator. In the case of these "soft" oscillators, where the frequencies of interest are far above the mechanical resonance frequency, the oscillator can be approximated as a free mass and broadband radiation pressure effects can be observed. Only recently was QRPN measured in the audio band [89], and it was demonstrated that such noise could be significantly reduced through the use of squeezed light [90]. While this noise source does not yet limit the aLIGO detectors, as the circulating power increases, such noise will eventually limit the low-frequency sensitivity.

The use of soft mechanical oscillators has a number of advantages in its own right. Classical radiation pressure forces can significantly alter the dynamics of these systems, which can provide optical dilution and drastically enhance the quality factor of these oscillators. By measuring off resonance, the effects of QRPN can be observed across a wide frequency band, particularly one that is relevant for GW detectors.

These setups also have significant challenges. The low frequencies of the mechanical oscillators means that cooling such systems to near the motional ground state is more challenging. Vibrational, acoustic, and other sources of noise are significantly higher in the audio frequency band, and as such complex vibration isolation and differential readout schemes are required.

QRPN is expected to limit the sensitivity of future GW detectors, so an understanding of the behavior of this noise source is essential. Additionally, once we reach the QRPN-dominated regime, we can induce correlations between amplitude and phase quadratures (section 2.4) and could make an optomechanical squeezer, which would allow for sub-SQL sensitivity.

In this chapter, we will review the experimental setup, prior work, and describe attempts to reach the quantum back-action-dominated regime. While we were ultimately unable to reach the sensitivity required to measure QRPN, we have improved the thermal noise of the end mirrors to obtain the best sensitivity yet. The lessons

from this experiment will be helpful for other audio-band quantum-limited room-temperature optomechanical experiments.

3.2 Experimental Setup

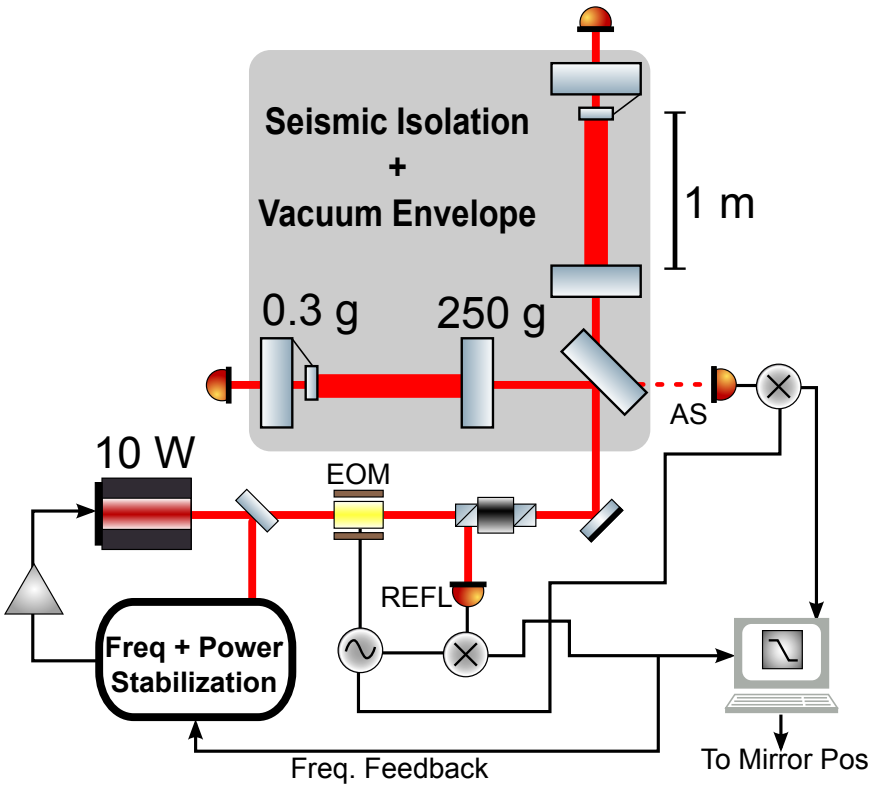

Figure 3-1: Layout of the gram-scale oscillator experiment. The experiment is a Fabry-Pérot Michelson interferometer (FPMI) with gram-scale end mirrors. Readout is via the Pound-Drever-Hall (PDH) scheme. Actuation is provided by an analog path that adjusts the laser frequency and a digital control system that actuates on the mirror length and angular degrees of freedom. EOM: electro-optic modulator, AS: anti-symmetric photodetector, REFL: reflected photodetector.

The optical layout is shown in figure 3-1. A description of the setup is found in [91] and [92]; we review the important characteristics and describe some updates to the setup since that work. Like LIGO, the experiment consists of a FPMI. However, in lieu of recycling cavities, to enhance radiation pressure effects, the arms are high finesse (8000 vs 400 for the observatories) with end mirrors of mass 0.3 g (vs 40 kg for observatories). The input mirrors are significantly heavier than the end mirrors

at 250 g, so the effects of radiation pressure on these optics can be ignored. The differential arm length (DARM) degree of freedom (DOF) is the figure of merit, and allows for the cancellation of common-mode noise from the laser.

An iLIGO master oscillator power amplifier (MOPA) [93, 94] provides the laser for the experiment. This system consists of a double-passed free-space Nd:YAG amplifier seeded by a non-planar ring oscillator (NPRO) operating at 1064 nm, which generates 10 W of output power. This beam is transmitted through an iLIGO pre-mode cleaner (PMC) [95] (not shown) to remove the higher-order transverse modes and passively suppress noise around the RF sidebands. The laser is locked to an ultrastable in-vacuum reference cavity to stabilize the frequency, and the relative intensity noise is stabilized to $10^{-7}/\sqrt{\text{Hz}}$.

The input mirrors and beamsplitter are iLIGO small optic suspensions [96]. These optics are suspended from a steel tower to reduce seismic coupling above the pendulum resonant frequency of 1 Hz. A set of shadow sensors and voice coils known as optical sensor and electromagnetic motors (OSEMs) and associated magnets glued to the mirror allow for local sensing and actuation of the optic in the length, pitch, yaw, and side degrees of freedom. Because the end mirrors have limited actuation authority, the main length degrees of freedom are controlled by input mirror and beamsplitter OSEMs, while the laser frequency provides an additional high-frequency actuation path for common arm length (CARM). An aLIGO computer system is responsible for data acquisition and digital controls. The entire interferometer is located inside a vacuum chamber and has both a passive isolation stack [97] and active external seismic isolation stage [22].

The length degrees of freedom are read out using a variation of the PDH locking scheme [35]. As in the case of single cavities, the interferometer is only sensitive when both arms are on resonance and the Michelson degree of freedom is locked on a dark fringe. Rather than wait for all degrees of freedom to simultaneously reach resonance, lock acquisition occurs in a controlled manner using different readout mechanisms than the final low-noise configuration. Lock acquisition is entirely automated and, in good seismic conditions, takes a few minutes to reach low-noise configuration. First,

each arm cavity is locked using a side-of-fringe lock in transmission. Using a slightly misaligned input beam (not shown in figure 3-1), the Michelson DOF is locked at a dark fringe. The path of this misaligned beam is adjusted so that it does not enter the arm cavities (and is thus insensitive to CARM/DARM) and is approximately 90° out of phase with the main beam. This latter property allows for DC locking the Michelson DOF using this misaligned beam while keeping the main beam at a dark fringe.

The arms are then brought to $\sim 90\%$ of the transmission fringe and the CARM and DARM DOFs are handed off to RF control, at which point the arms can be brought fully to resonance. Even at 30 mW, we observe an optical spring [42] at 100's of Hz. As the arms are moved to resonance, the dynamics of this spring change drastically, so a number of digital filters are engaged and disengaged during this process to compensate for the moving optical spring frequency.

Once on resonance, the Michelson DOF is handed off to RF control, the misaligned input beam is shuttered, and a number of other control loops are engaged to bring the interferometer to the optimal configuration. The injected power can then be increased to a few watts.

Auxiliary DOF Control

The local angular motion of the large optics is sensed and controlled with the OSEMs. Unfortunately, these cannot be used for the motion of the small end mirrors, so optical levers are used to sense the position, pitch, and yaw motion of the end mirrors, and control is fed back via the outer ring OSEMs for local damping. In full lock, the optical levers add noise, so the angular degrees are instead controlled by sending a band-passed signal from DARM and CARM at a narrow band around the measured mechanical frequencies.

Low-bandwidth alignment dithers (as described in section 2.3.2) on all optics are used to maximize the arm transmission and improve the dark port contrast. Dither loops constantly monitor the gain of control loops and adjust the digital gain to keep the system stable as the power changes. These same injected excitations are used to

remove any DC offsets in the length DOFs by sensing the common and differential transmission signals (CARM and DARM offsets) or the anti-symmetric (AS) DC power (for the MICH offset).

Calibration

The known laser wavelength (1064 nm) is used to calibrate the instrument. First the end mirrors are misaligned so that only the MICH DOF is locked using the side of fringe. The distance between fringes is given by the laser wavelength, so we can drive the ITM actuators and this knowledge to calibrate those actuators. As expected, the response in length has a f^{-2} dependence up to 2 kHz; above this frequency we ran out of actuation range. A full closed-loop DARM transfer function was then performed in full lock to convert between measured signal to a calibrated open-loop noise spectrum.

This actuator calibration method was tested using a number of different alignments, input powers, and photodetectors, and in each case the actuation strength varied by no more than 3% between measurements.

3.3 Gram-Scale Mirrors and Thermal Noise

Previous work with this setup [81] suggested the interferometer sensitivity was limited by thermal noise due to low-Q high-frequency mechanical modes of the end mirrors. These mechanical modes, while in the 10's of kHz range and far outside the measurement band, involved large stresses of the lossy epoxy attaching the mirror to the fiber supports. Modeling suggested that by using flat attachments, less epoxy would be required and those modes would be shifted to higher frequencies. Both effects would significantly reduce the thermal noise in the measurement band. As such, the end mirrors were rebuilt in an attempt to reduce the suspension thermal noise.

Figure 3-2 shows the two suspension systems. The new optics were cut from the original round cavity optics and should have similar optical properties. Testing of the optic prior to installation discovered irregularities in the coating and a position-dependent transmission. Measurements prior to installing the new mirrors suggested

that the irregularity was sufficiently off-center to not affect operation (assuming the beam could be kept near the center of the optic), and the curvature was measured to be in line with the original value. Measurements of the optic in a small Michelson interferometer to measure fringes indicated the optical quality should be sufficient to create a cavity of a few thousand.

Optic	Mode	Frequency [Hz]	Q Factor ($\times 10^3$)
ETMX	Pendulum	11.4116	?
	Yaw	89.7789	337
	Pitch	116.5445	391
ETMY	Pendulum	11.174	?
	Yaw	92.938	450
	Pitch	128.6634	439

Table 3.1: Measured ETM suspension modes. Excess noise around 10 Hz precluded a reliable measurement of the Q factor of the pendulum modes, however, finite-element modeling predicts a Q factor around 400 000.

3.3.1 Thermal Noise

The thermal noise is intimately tied to the damping of mechanical oscillators. When the oscillator is rung up, damping will convert mechanical motion to heat, which leaks into the environment. Conversely, at thermal equilibrium at finite temperatures, the mechanical system will constantly be exchanging energy with the environment, which will occasionally provide "kicks" to the system and produce mechanical noise. Thus, the damping term is related to the rate at which energy is exchanged between the mechanical oscillator and the environment. This relationship is formalized in the fluctuation-dissipation theorem [98]:

$$G_{FF}(\Omega) = 4k_B T R(\Omega), \tag{3.1}$$

where $G_{FF}(\Omega)$ is the power spectrum at frequency Ω, k_B is the Boltzmann constant, T is the absolute temperature, and $R(\Omega)$ is the resistance of the system, which is related to its damping. In this form, the relation to the standard Johnson-Nyquist

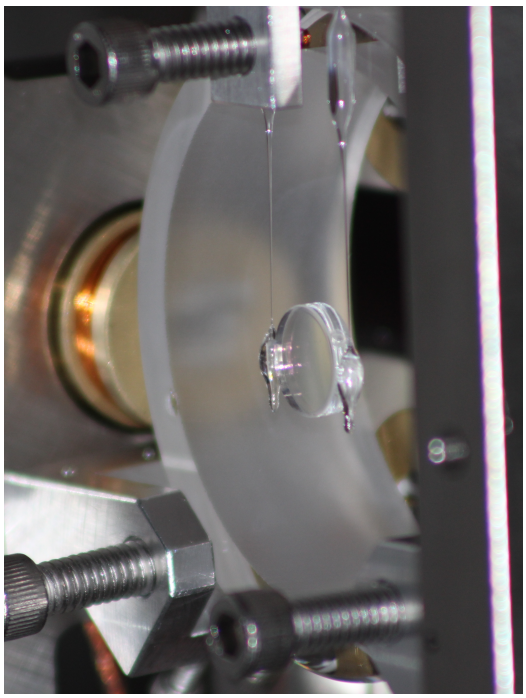

(a)

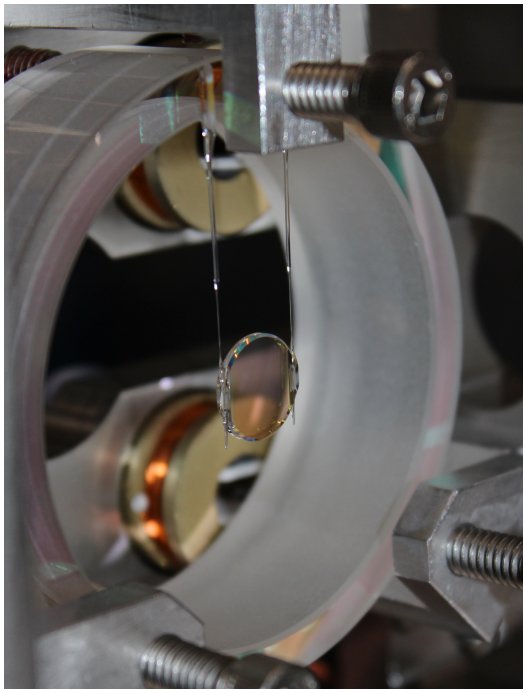

(b)

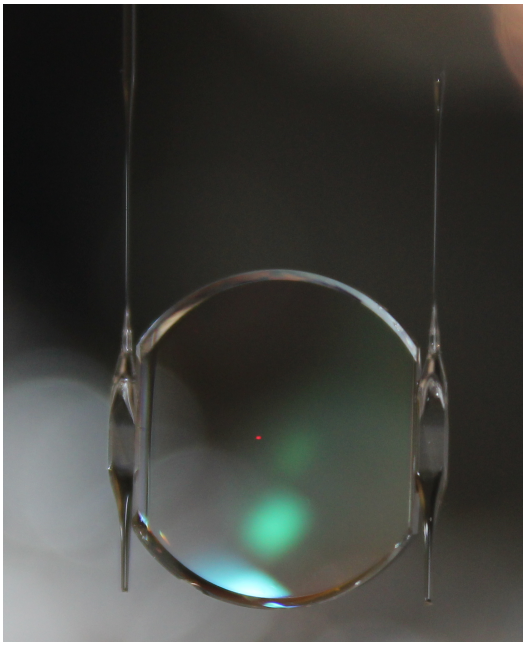

(c)

Figure 3-2: Comparison of the (a) original end mirrors and (b) new end mirrors with new support structure and lighter mass. (c) Closeup of new suspensions. The new mirrors are approximately 10 mm in diameter. The outer ring is suspended from the tower by steel wire.

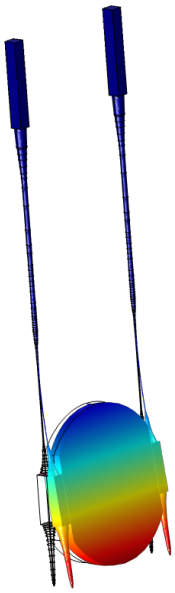

(a) Fundamental

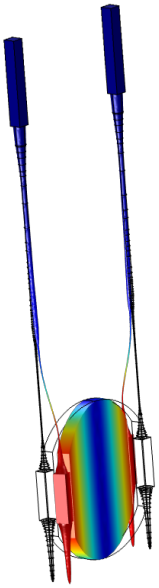

(b) Yaw

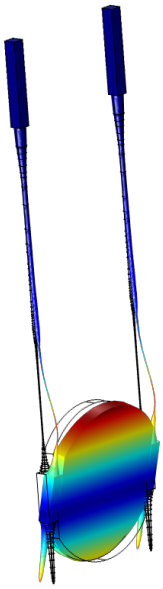

(c) Pitch

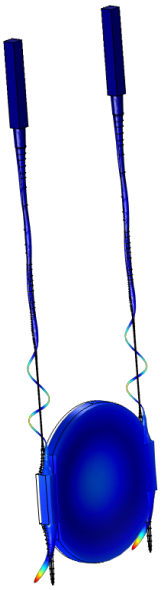

(d) Drumhead

Figure 3-3: Modeled gram-scale mirror mechanical modes most relevant for thermal noise. The measured properties of these modes is given in table 3.1. The drumhead mode was not measured directly but is expected to be around 55 kHz with Q factor around 30 000.

noise in resistors is clear [99, 100]. To extend the concept of "resistance" to mechanical systems, we can write a mechanical version of Ohm's law which relates force to a "flow":

$$F(\Omega) = Z(\Omega)v(\Omega), \tag{3.2}$$

where $v(\Omega)$ is the velocity of the mechanical oscillator, and $Z(\Omega)$ is the mechanical impedance. As in the electrical case, the resistance is defined as the real part of the impedance $R(\Omega) \equiv \Re(Z(\Omega))$. In the frequency domain, we have $v(\Omega) = i\Omega x(\Omega)$, so we can define a mechanical susceptibility $\chi(\Omega)$ that relates force and displacement:

$$F(\Omega) = Z(\Omega)(i\Omega x(\Omega)) = \frac{1}{\chi(\Omega)} x(\Omega). \tag{3.3}$$

We can then write the thermal noise power spectrum as

$$G_{xx}(\Omega) = |\chi(\Omega)|^2 \, G_{FF}(\Omega) \tag{3.4}$$

$$= \frac{4k_B T R(\Omega)}{\Omega^2 \, |Z(\Omega)|^2} \tag{3.5}$$

$$= \frac{4k_B T}{\Omega^2} \Re\left(\frac{1}{Z(\Omega)}\right). \tag{3.6}$$

In the case of mechanical oscillators with natural frequency Ω_m and mass m, we can write the damping term as an additional imaginary component to the spring constant called the *loss angle* $\phi(\Omega)$:

$$K_m(\Omega) = m\Omega_m^2[1 - i\phi(\Omega)]. \tag{3.7}$$

The loss angle is related to the mechanical resistance by

$$\phi(\Omega) = \frac{\Omega}{m\Omega_m^2} R(\Omega). \tag{3.8}$$

In the case of viscous damping ($F \propto v$), the loss angle is proportional to frequency: $\phi(\Omega) = \Gamma_m \Omega / \Omega_m^2$, where Γ_m is the mechanical damping rate. In the case of structural damping, which seems to describe loss due to internal friction of a material [101],

the loss angle is constant: $\phi(\Omega) = \Gamma_m/\Omega_m = 1/Q_m$. However, the loss angle need not have such a simple relation to frequency, and in the case of the aLIGO optical coatings, the loss angle was found to have a $f^{-0.1}$ slope [102].

Near the mechanical resonance frequency, these loss mechanisms look very similar. But far off-resonance, where this measurement is made, the thermal noise will differ. As has been previously measured in this and similar systems, structural damping seems to be the dominant loss effect in these systems [81].

The thermal noise is estimated with a finite-element model of the suspension (section 3.3.2) using the normal mode decomposition developed by Gillespie and Raab [103]. The derivation is provided in section A.3. This noise depends on the overlap between the optical beam and the mechanical motion of the mirror; figure 3-4 shows how this noise depends on beam position. This was estimated using the first 50 mechanical modes, reaching up to 80 kHz.

It should be noted that we obtained similar results using the "direct method" of Levin to calculate thermal noise [104]. This technique was found to be more computationally expensive, and the normal-mode picture is more intuitive for the ultimate limitation of beam position. Other simulations using this method have proved to be sufficiently accurate for these types of systems [105].

3.3.2 Finite-Element Model

Previous modeling [81, 91] suggested that this contribution to the thermal noise could be mitigated by using flat supports for the suspension attachments instead of the original "turkey leg" supports. This comes about for two primary reasons. First, the flat attachment is significantly stiffer, pushing the resonance frequency of relevant modes much higher and reducing the thermal noise contribution to the QRPN measurement band. Second, the quality factor of these modes is dominated by losses in the epoxy used to these flat attachments require less epoxy.

A new model was developed in COMSOL Multiphysics® using the as-built dimensions of the silica fibers to determine the exact contribution with the new suspensions. The measured profile of the fibers was used in the new model (minimum fiber diame-

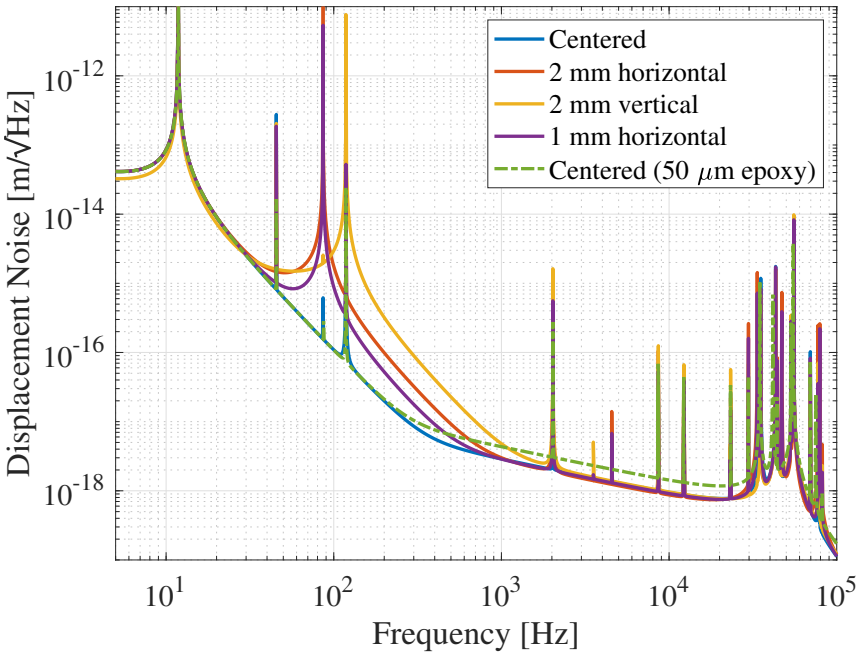

Figure 3-4: Estimated thermal noise contribution from a single end mirror based on beam position, either located at the node of the pitch and yaw modes ("centered") or offset. The dashed curve shows the estimated thermal noise when assuming 50 µm of epoxy instead of 10 µm. Depending on the beam position, thermal noise from the yaw and pitch modes can be observed. These modes are at 86 and 118 Hz in this model, respectively, and are shown in figure 3-3.

ter 100 µm) with minor adjustments to obtain similar frequencies and quality factors for the position, pitch, and yaw modes.

There is uncertainty in the thermal noise level at high frequencies, as this number will depend on the amount of epoxy applied, the exact loss angle, and the exact frequencies of the higher-order modes. That last variable depends on the smaller features of the mechanical model, which we do not have the resolution to accurately model with certainty. However, this level is consistent with the upper limit established with a measured displacement spectrum (section 3.5).

3.4 Experiment Upgrades

As previously mentioned, both end mirrors were replaced with new mirrors using flat suspension supports. Additionally, these optics were lighter than the original mirrors

(0.3 g vs 1 g) in the hopes of observing radiation pressure effects at lower input powers.

In addition to the lighter end mirrors with improved thermal noise, a few others improvements were implemented. Due to high-frequency optical springs introduced when moving to resonance, the DARM control loop requires unity-gain frequency (UGF) around 1 kHz. The digital control system has a finite delay that introduces phase loss and can make control loops unstable. To improve the reliability of this system, the digital control system was upgraded to an aLIGO digital system, with custom modifications to allow the system to run across multiple cores. In this system, we obtained intrinsic delays of 33 µs, which leads to $\sim 12^\circ$ phase delay at 1 kHz. This is sufficient to allow us to produce a stable control loop, as a proof of principle that this system could be used if additional bandwidth was required in other LIGO systems.

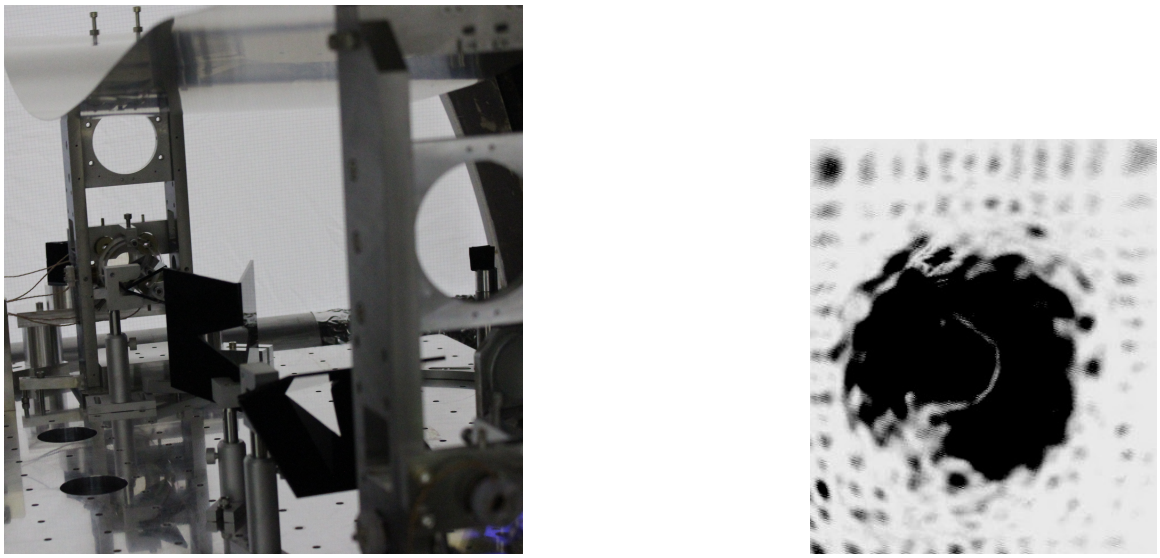

(a) (b)

Figure 3-5: (a) Looking down an arm cavity, showing the baffles installed to catch high-angle scattered light. (b) Inverted image of AS port beam at 400 mW input power. A slight mode mismatch between the arms can be seen, but additionally a high-order mode can also be seen. This is believed to be generated by scattering on ETMY, which then resonates in the Y arm.

To avoid scattered light [106, 107, 108] that could reflect off the vacuum chamber or optical table, arm cavity "baffles" were installed along each arm (figure 3-5a). Two types of baffles were constructed: a triangular baffle that would catch wide-angle scattering from the mirror surface, and a wedge that could be placed in the middle

of the arm to catch light scattered more narrowly. Unfortunately, the former baffle limited the ability to determine the beam position on the end mirrors, and as such, a scaled-down version of the optic baffle was installed instead. There was no observable difference in performance after the installation of these baffles, likely due to low-angle scattering back into the cavity axis being the dominant noise mechanism.

Concerned about the effect of scattering, the cavity length was adjusted by a few centimeters to change the transverse mode overlap condition. Unfortunately, the scattering was so severe that even much higher-order modes would be excited. Figure 3-5b shows one such mode at the anti-symmetric port.

While this experiment did not have an output mode cleaner as in aLIGO [109], we attempted to simulate this effect by installing a coiled single-mode fiber at the output port to remove higher-order modes reaching the photodetector. This had a detrimental effect on the interferometer noise at low frequencies, most likely as a result of angular motion changing the coupling to the fiber or phase noise due to length fluctuations of the fiber.

3.5 Results

Figure 3-6 compares the thermal noise of the old and new suspensions. As previously reported [81], the old suspension spectrum is limited by thermal noise due to high-frequency bending modes of the optic supports and the lossy epoxy attaching the fibers to the mirror. The curve for the new suspensions is taken as the minimum value for many different measurements at different optical powers. Included is the thermal noise model for the new suspensions with the beam "centered" at the node of the pitch and yaw modes, eliminating the coupling of the thermal noise of these modes to the spectrum.

We also investigated the effect of beam position on the end on the noise spectrum. In this setup, our probe of the mirror position is also the source of radiation pressure effects. By driving a spectral line in the laser intensity stabilization servo, we could drive the pitch and yaw modes of the end mirrors and observe a similar line in the

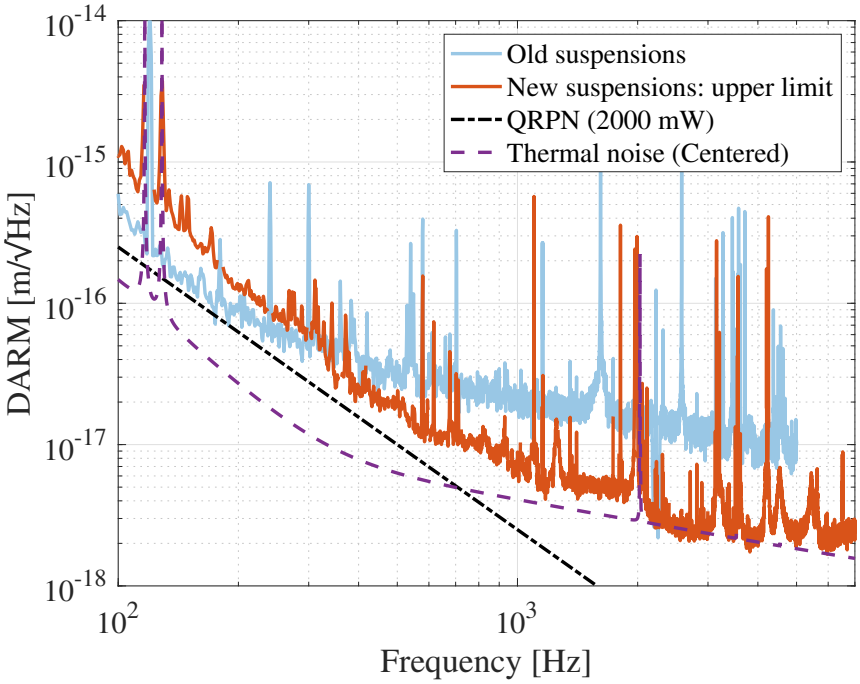

Figure 3-6: Comparison of the thermal noise of the old suspensions (measured) with an upper limit of the thermal noise for the new suspensions. Also included is the modeled thermal noise of the new suspensions with a "centered" beam (definition in text) and the goal sensitivity to measure QRPN. The sensitivity of the interferometer using the new suspensions has been limited by a number of factors, so this plot is an composite of multiple measurements, where the high-frequency portion is obtained at high input powers, and the low-frequency portion is from low input powers.

readout. By changing the spot position on the mirrors, this line could be minimized[1], which in the process would minimize the coupling of the thermal noise of this mode to the readout. Figure 3-7 shows the results of using this technique to minimize the coupling to the pitch mode of one end mirror. Before moving the beam position, the spectrum is modeled well by thermal noise due to the pitch mode, and even with increased power, this still overwhelms the QRPN contribution. Moving the beam vertically reduces the noise at low frequencies, however, there still is a mystery noise at low frequencies. More problematic, attempts to increase the operating power proved challenging. Figure 3-8 shows the effect on the spectrum of increased laser power. The low-frequency noise significantly increases in power for unknown reasons.

[1] As the circulating power is increased, the static radiation pressure force introduces slight alignment shifts, which change the coupling to the angular modes. As such, to ensure the angular error signal keeps the same sign, in practice we cannot operate exactly at the angular mode node.

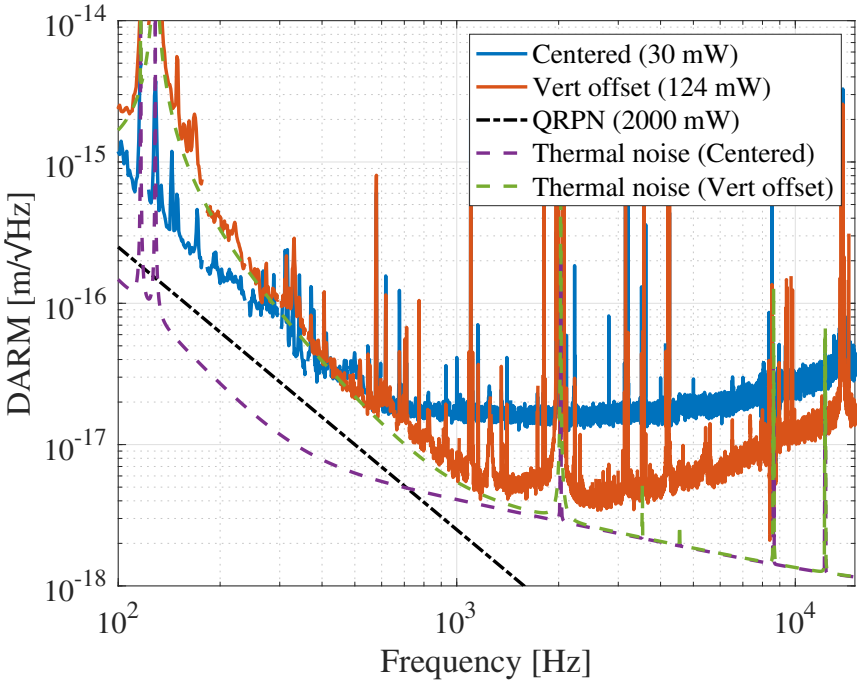

Figure 3-7: Displacement sensitivity as a function of beam position on the end mirror. While there are smaller optical losses for a beam offset from the mechanical mode nodes, there is excess low-frequency noise believed to be due to the pitch mode thermal noise. By moving the beam toward the node of this mode this noise is reduced, but we are unable to reach high circulating powers due to high losses.

Not only did this degrade instrument sensitivity, the larger forces required to keep the instrument in the low-noise configuration tended to saturate actuators and unlock the interferometer.

There are a few hypotheses on the origin of this noise. The static radiation pressure force may introduce additional alignment offsets that will change the noise coupling; indeed, implementing an alignment feedforward system to move the end mirrors as the power was increased proved to improve the instrument reliability. The increase in motion may lead to increased nonlinear effects, some which have been observed in the LIGO detectors as well [32]; unfortunately the origin of these effects is still unknown. While the relative power fluctuations should stay constant, in absolute terms classical radiation pressure effects will increase with power. The relative intensity noise is measured with an out-of-loop sensor to be as low as $1 \times 10^{-7}/\sqrt{\text{Hz}}$ in the measurement band. However, when the prompt reflection is measured from the input mirrors, this is

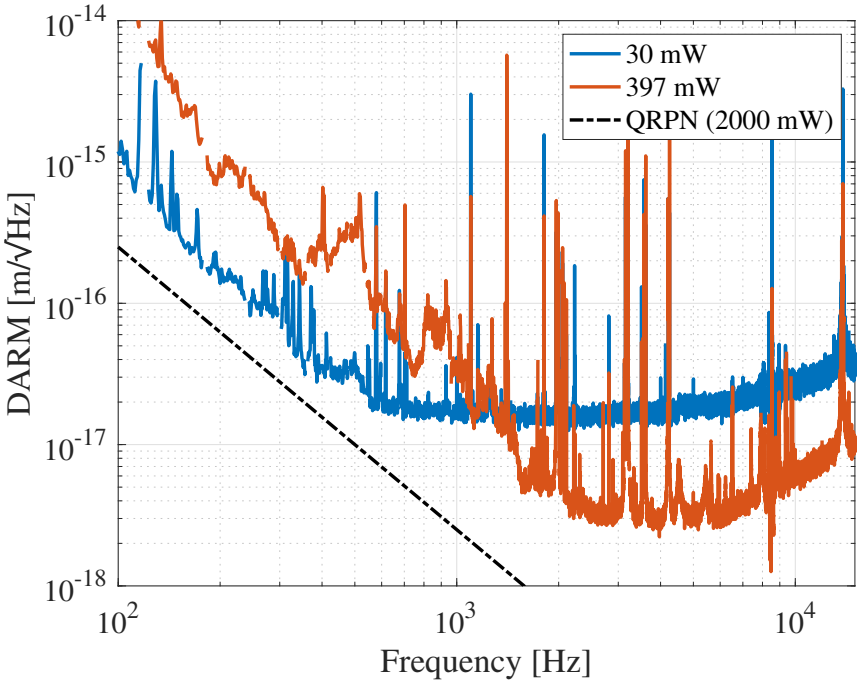

Figure 3-8: Sensitivity decrease for higher input powers. While the shot-noise-limited region at high frequencies is significantly reduced at higher input powers, the low-frequency noise becomes significantly higher. This is thought to be due to scattering into higher-order transverse modes of the cavity (see discussion in text), and possibly radiation pressure noise due to increased absolute laser intensity noise.

found to be a factor of 3 higher, suggesting additional intensity noise coupling after the intensity stabilization servo (ISS). The reasons for this were not understood; attempts to move the ISS detectors closer to the injection point proved to be ineffective.

Finally, the most likely suspect is scattered light. Unfortunately, the process used to chop off the ends of the end mirrors seems to have damaged the optical coating. As a result, scattering losses on the surface of the end mirrors was strongly position dependent. The positions that had minimal scattering would couple in thermal noise, while at the ideal operating position, we observed a decrease in the transmitted light of 30%. It is believed that light was scattered into other transverse modes that would periodically resonate in the arm cavity, with a strongly nonlinear classical radiation pressure force. Figure 3-5b shows an image of the AS port at high powers, showing a high order mode that was resonating in one arm. The high order of the mode suggests that scattering is particularly bad, and that changing the length of the cavity would

only make another transverse mode resonate instead.

Unfortunately, continued technical difficulties due to aging hardware precluded a full exploration and quantification of these effects.

3.6 Conclusions

We have reviewed the most recent results of a FPMI with gram-scale mechanical oscillators. New mirror suspensions with improved thermal noise were constructed and installed, and a new finite-element model was developed. Above 400 Hz the new suspensions have improved thermal noise, as expected, up to a factor of 5 at high frequencies. Below this, the noise spectrum is dominated by thermal noise of the optic pitch and yaw modes. In principle, this excess thermal noise can be removed by positioning the beam at a node of these modes. However, non-uniform optical losses and scattering points have precluded a favorable beam position that also allows buildup of intracavity power. Additional low-frequency noise at increasing powers appears, making it difficult to reach a QRPN-dominated region. Aging infrastructure and other technical issues have prevented the removal of the remaining technical noise and injecting higher powers. Nevertheless, we were able to validate the original thermal noise model and show that this term could be significantly reduced. These measurements have been a stark reminder that the thermal noise from mechanical modes far outside the measurement band must be considered. This important concept has been used in the successful design of other mechanical oscillators for QRPN-dominated experiments [105, 110].

Chapter 4

Test Mass Charging and Electric Field Coupling

While the aLIGO design sensitivity is limited at most frequencies by quantum noise, in practice there are other limiting noise sources. Figure 1-5 shows the current noise budget for LLO, showing the other known and unknown sources of technical noise. Both detectors, as well as many other radiation-pressure dominated experiments [111], are limited by excess noise at low frequencies. Because the effect of stray electric fields is another potential force on the optics, it can very easily mask the effects of radiation pressure, so understanding this noise term is important for reaching QRPN-limited measurements. This section will describe work carried out at LLO to understand and mitigate the effects of test mass charge and stray electric fields near the test masses. This broad investigation will investigate how actuation forces are modified by charge on the test mass, describe techniques to assess charge distribution on the test mass, model noise coupling in a charged test mass interferometer model, and describe other investigations on stray electric fields.

4.1 Test Mass Actuation

Excess charge can modify the length actuation of the test mass. To understand these effects, here we review the aLIGO quadruple pendulum suspension ("quad") and test

mass electrostatic actuation. Additional details can be found in [18].

The quad consists of a four-stage pendulum suspended from the aLIGO internal seismic isolation (ISI) platform. At the bottom of the suspension is the 40 kg fused silica test mass, which is attached to the penultimate mass (PUM) via a set of fused silica fibers. A secondary "reaction chain" hangs behind the main chain and contains actuators for the lower stages. The fact that it is suspended allows for actuation without additional seismic coupling. These actuators are used to maintain the arm cavity resonant condition by adjusting the test mass length and angular degrees of freedom. All stages except the test mass are actuated via electromagnetic coils and magnets affixed to the main chain masses. The upper stages have greater actuation range than lower stages to counteract ground motion. As a result, the upper stages have more actuation noise, but this is passively filtered by the suspension pendulum before reaching the test mass. A metal cage surrounds the suspension and provides protection in the event of a fiber breakage. Additionally, a ring heater (RH) and metal shield is mounted on the cage and is used to compensate thermal distortion of the test mass due to optical absorption [30]. Though not used for test mass actuation, the proximity of the ring heater and shield to the test mass requires us to consider how it affects the actuation. The test mass and its surroundings are shown in figure 4-1.

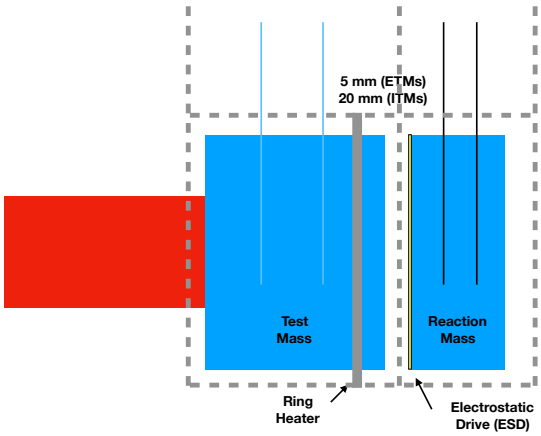

Figure 4-1: Simplified side view of an end test mass/reaction mass and surroundings, including the arm cavity laser beam, the surrounding cage, and ring heater shield. The transmitted laser is omitted. Recall that the test mass is the last stage of the aLIGO quad suspension (figure 1-3).

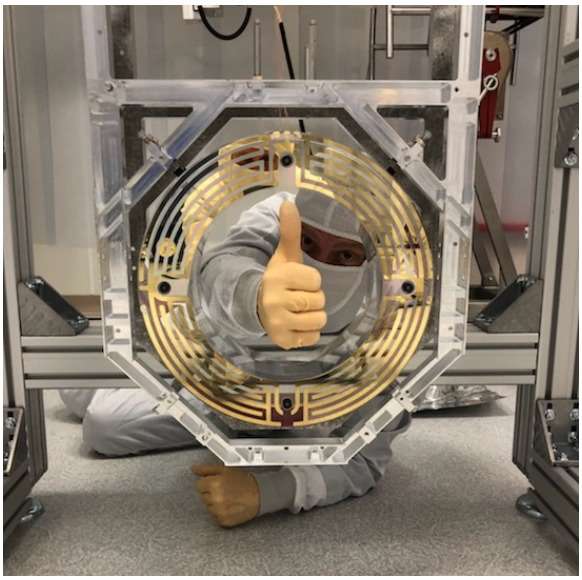

Figure 4-2: The installed end reaction mass (ERM) showing the installed electrostatic drive (ESD) pattern. Prior to O3 the original end reaction masses (ERMs) were replaced with annular ERMs, in which the center section was removed to reduce the effects of squeezed film damping. Photo credit: Gary Traylor and Danny Sellers

Due to concerns about excess thermal noise and "crackling" due to magnetic domain flipping (Barkhausen noise) [112, 113], no magnets are attached to the test mass and it is instead controlled via electrostatic actuation. An overview of the design and initial performance of the electrostatic drive (ESD) can be found in [114]. In the following section we review this scheme.

4.1.1 A Simple Model

The ESD consists of a pattern of interdigitated gold electrodes on the end reaction mass (ERM), the bottom stage of the reaction chain (figure 4-2). One electrode (the "bias") is held at a static, high voltage. The other set of traces (the "signal") has a switchable connection to a high-voltage driver or a low-noise, low-voltage driver. The former is used for lock acquisition, while the latter is used for low-noise operation. Figure 4-3 shows the principle of operation: the high voltage on the bias trace polarizes the test mass and produces a static inhomogeneous electric field, which attracts the test mass to the reaction mass. Changes in the voltage of the signal trace will change

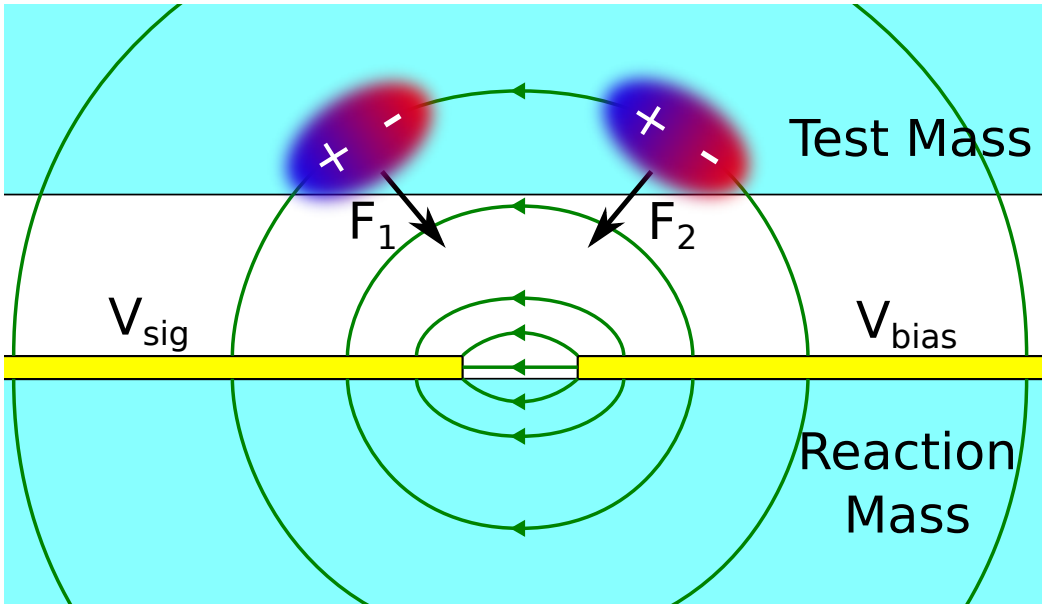

Figure 4-3: Cartoon illustrating the ESD principle of operation. A static high voltage is applied to the bias trace to produce an electric field (green) that polarizes the test mass. Because this field is inhomogeneous, the test mass experiences a net attractive force toward the ESD. The voltage of the signal trace is adjusted to change this force. Note that here it is assumed that $V_{\rm bias} > V_{\rm sig}$, but this need not be the case.

the gradient of this electric field, producing a force on the test mass according to

$$F = \alpha \Delta V^2 = \alpha (V_b - V_s)^2 = \alpha (V_b^2 - 2V_b V_s + V_s^2), \tag{4.1}$$

where ΔV is the voltage difference between ESD traces and α is a factor that depends on the geometry of the generated electric field. Note that in the limit of small changes in signal voltage, the V_s^2 term can be neglected, and the only fluctuating term is linear in both the signal and bias voltages[1]. Note that the bias trace is a single trace, while the signal trace is separated into four separate quadrants that can be driven individually. In normal operation all signal traces are driven together, so we will consider them together as a single trace[2].

[1]Note that this squared term cannot be neglected during lock acquisition when the full ESD range is required, however, the digital control system can calculate and compensate for this term. We will concern ourselves only with the noise properties in the low-noise state, where the squared term has been measured to be negligible and will be ignored from here.

[2]The segmented nature of the ESDs is used to damp body modes of the test masses to eliminate PIs [115], but we will not discuss this here.

4.2 Prior Work and Motivation

A number of observations suggested that charge might be affecting the performance of gravitational-wave detectors, including the ESD actuation. Three effects in particular demanded a more complex actuation model:

1. **Weaker than expected actuation:** compared to simulations [114], the ETM ESD actuation strength was approximately 30% lower than expected [116].

2. **Actuation force not zero at zero bias:** as seen in equation (4.1), for a purely dipolar force, if there is no bias voltage and no signal voltage offset, for small signals there should be no force on the test mass. However, the ESDs still produced a small linear force even for zero bias.

3. **Actuation strength changing in time:** over the time scale of weeks, the actuation strength was observed to be changing by a few percent [117]. An example of this change at LLO during O2 is shown in figure 4-4.

Effect 1 was determined to be due to the suspension cage surrounding the optic and the ring heater shield (see figure 4-1), which is held at the same potential as the chamber. When considering the electric field lines going to the cage from the ESD, this explains the weaker actuation [118]. Note that the ITMs didn't have as large a discrepancy due to the ESD being farther away from both the optic and cage.

Effects 2 and 3 suggest the presence of a time-varying charge distribution on the test mass. Instead of merely coupling to the polarized test mass, the ESD additionally couples to the separated charges. Because of the insulating properties of the test mass, the charge distribution is relatively stable but is believed to be pushed around due to the near-constant application of the bias voltage. Tabletop experiments using ESD-like electrodes and fused silica samples have observed this charge separation on long time scales [119], as well as discontinuous charge jumps in silica samples [120]. This charge separation is proposed to be due to ionic imperfections in the fused silica [121, 122] or surface charges due to adsorbed layers of water molecules [123, 124] that are then separated by the applied bias voltage.

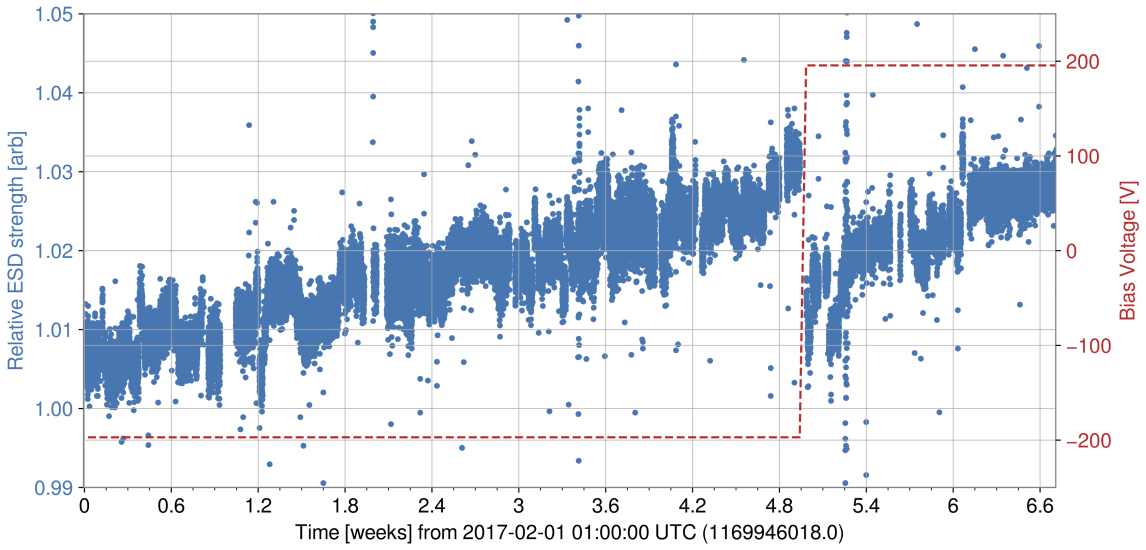

Figure 4-4: Track of the ESD actuator strength at LLO during O2 and average ESD bias voltage applied. The discontinuity in the actuator strength around week 5 is due to a change in the applied bias voltage. In general, for a constant bias voltage, the ESD strength has been observed to increase in magnitude.

In full-size gravitational-wave detectors, various anomalous observations suggested the effect of stray electric fields was not fully understood, and that charging of the test masses could be contributing to unexplained excess noise at low frequencies:

- After a power loss at the GEO 600 interferometer, an uncontrolled test mass accidentally contacted an ESD that remained at a high voltage [125]. Subsequently the interferometer would not lock, and the ESD actuation strength was abnormal. Only after venting the chamber and exposing the test mass to UV light was the actuation returned to normal.

- In July 2017, an earthquake in Montana likely caused a test mass at LIGO Hanford Observatory (LHO) to contact one of the earthquake stops that keeps the test mass safe during large ground motion. While there was no damage, there was excess broadband low-frequency noise. Additionally, there was significant coupling between the suspension cage motion and the test mass motion, suggesting image charge [126] attraction in the cage. Much of this excess noise and all of the cage motion coupling was eliminated by opening the chamber and discharging the test mass.

- At LLO, high ESD bias voltages produce excess broadband displacement noise (figure 4-5).

Additional theoretical work suggests that charge effects could reduce the sensitivity of the detectors, though none of these effects has been unambiguously determined to be a limiting noise term. Stray electric fields due to surface charge moving on the test mass is one possible source of low-frequency noise with an expected f^{-1} force amplitude noise dependence [127]. A charged test mass will produce image charges [126] in the cage, which would couple cage motion to the test mass [128]. Large uncertainties on the actual surface charge have complicated estimates of these noise terms, though this charge fluctuation has been measured in tabletop experiments [129]. A summary of all the known effects of charge on detector operation is shown in figure 4-6.

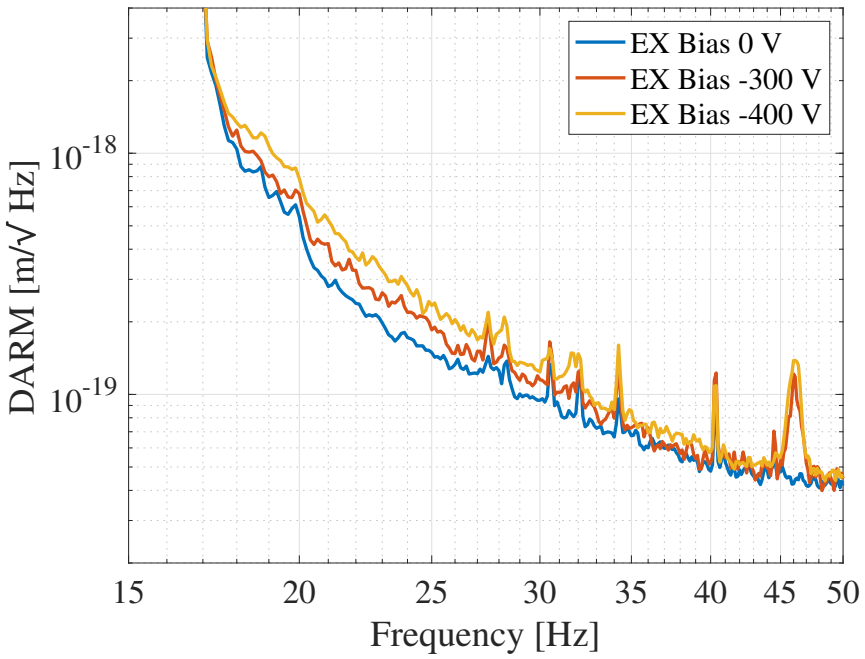

Figure 4-5: Broadband increase in the noise floor when a bias voltage is applied to the ETMX ESD. Additional narrow-band features also begin to appear at high bias voltages, which are likely due to electronic pickup via other mechanisms.

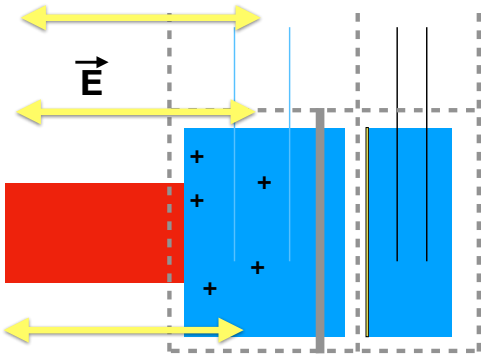

(a) Large-scale varying electric fields in combination with net charge on the test mass.

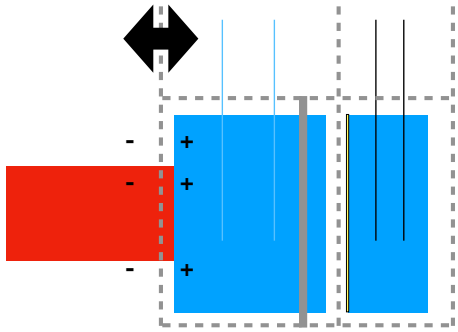

(b) Image charges in the optic cage can couple cage motion to test mass motion.

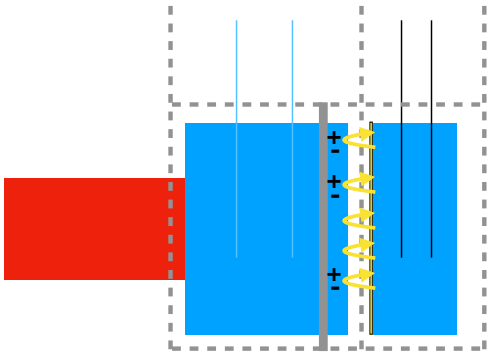

(c) Local charge separation near the ESD traces can change the actuation strength.

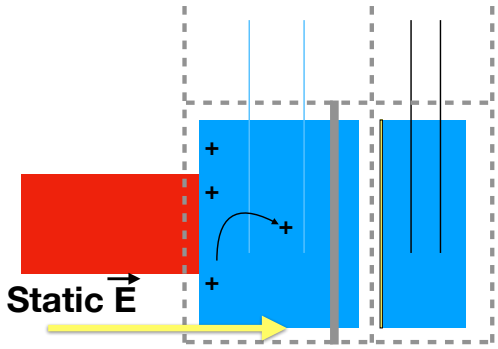

(d) Charge "hopping," in tandem with static electric fields, can couple in additional noise (including transients).

Figure 4-6: A summary of the different ways charge can change the operation of the interferometer and/or couple in excess noise. While the charges are on the test mass in (c) and (d), charge separation or hopping on the reaction mass near the ESD can also generate noise or modify the actuation.

4.2.1 Complete Actuator Model

When we consider all the above effects on the ESD actuation, we can write a more complete model for the actuation force[3] [130, 131]:

$$F = \alpha(V_b - V_s)^2 + 4\gamma\left(\frac{V_b + V_s}{2} - V_{\text{ref}}\right)^2 + 2\beta\left(\frac{V_b + V_s}{2} - V_{\text{ref}}\right) + \beta_2(V_b - V_s) + \delta, \quad (4.2)$$

where $V_{b(s)}$ is the voltage on the bias (signal) trace and V_{ref} is the voltage on the suspension cage and ring heater shield. Note that generally $V_{\text{ref}} = 0$, but as we will see in section 4.5, this voltage is non-stationary and can couple in noise. The coefficients represent:

- α: the dipole force as in equation (4.1). This term depends on the geometry of the ESDs and test mass and dielectric properties of the test mass (e.g., distance between traces, distance between ESD and test mass, etc.).
- γ: the effect of the cage and ring heater shield on the electric field gradient generated by the ESD. Depends on the ESD geometry relative to the cage.
- β: describes the effect of charge between the ESD and the cage, which will produce a force linear in the electric field between the ESD and cage.
- β_2: represents the effect of charge separation near the ESD traces, either on the test mass or ERM.
- δ: describes all other forces that do not depend on ESD or cage voltages. We will ignore this term for now.

Note that this model says nothing about the *net* charge on the optic, as the ESD can only probe the charge arrangement near the electrodes. Discussion of net test mass charge, including the effects on sensitivity, can be found in section 4.6.

[3] The seemingly arbitrary inclusion of numerical factors in front of each coefficient is to ensure this equation follows previous conventions for values of the actuation coefficients while retaining the physical intuition behind each term.

Some caveats that complicate the analysis: First, the local distribution of charge on the optic is almost certainly not uniform and hard to measure directly[4]. Second, small differences in the relative position of the reaction chain and the test mass can provide different measurements of the actuator strength [132]. This distance predicts short-term variations in actuator strength quite well[5] [133].

Note that for the ITMs, the ESD trace is bonded to the compensation plate (CP), which is analogous to the ERMs. However, because the CPs are farther away from the ITMs than the ERMs are from the ETMs (20 mm vs. 5 mm) and the actuator strength scales approximately as d^{-2}, in what follows we will consider only the ETMs. Additionally, at LLO, the ITM ESDs are not required for lock acquisition or low-noise operation, so they were shorted at the vacuum feedthrough prior to O3 and were inoperable (see further discussion in section 4.5).

4.3 Charge Origin

The fused silica test mass is an excellent insulator, with bulk electrical resistivity $>10^{16}\,\Omega\,\mathrm{m}$. Additionally, the thin silica fibers connecting the test mass to the penultimate stage makes the transfer of charge from upper stages unlikely. Thus, any *net* charge accumulation must come via direct contact with the test mass. Note that even if the optic is not charged, local charge localization can occur due to the persistent high-voltage bias; we will consider this "charging" as well.

A few proposed sources of optic charging include:

- **Triboelectric charging:** Friction between two surfaces can transfer charge. In the case of the test mass, the largest transfer of charge comes from the removal of First Contact Polymer, a polymer used to clean the tests masses prior to exiting and evacuating the vacuum chambers. This is mostly removed via a

[4]It's not clear whether the charge separation is most important on the test mass or the ERM. Fortunately, this can all be wrapped up in (4.2) as they relate to the instrument sensitivity, so we will not concern ourselves further.

[5]Lest the reader think this effect could explain entirely the observed actuation strength changing over time, know that this effect would change the α and γ terms of the actuation strength. This is not what is observed.

discharge system consisting of nitrogen gas flowing through a thin nozzle held at high potential, applied just before chamber exit [134]. An electrometer is placed near both the front and back of the test mass to confirm the optic has been discharged. It is assumed that the test mass at LHO was charged due to an earthquake stop rubbing against the test mass during the Montana earthquake (section 4.2).

- **Dust:** Residual particulates in the chamber may be able to move charge onto the optic. However, after pumping down this seems like an unlikely source. First, the ultra-clean environment inside the chamber is unlikely to support significant amounts of particulate matter. This is supported by the fact that the test masses remain in the chambers for years without a significant optical quality degradation once at low pressure, which might occur due to dust hitting or accumulating on the front surface. Second, because the test masses are off the chamber floor, in the absence of any air currents, most of the dust will settle to the chamber floor and has no way to reach the test mass. Finally, the optic itself is surrounded by a cage and other metal structures, all of which are grounded and would thus attract dust by image charges; it is unlikely that a dust particle would even reach the optic as opposed to other surfaces (unless the optic already has a large net charge, which is not observed).

- **UV photons:** Ultraviolet photons generated from ion pumps near the test masses [135] can ionize the test mass (or worse, damage the optical coating). Before O3, a set of baffles was installed on each ion pump to keep these photons from reaching the test mass; it remains to be seen if they are effective at keeping the test mass charge low.

- **Ionic imperfections in substrate:** Under certain conditions, charge carriers have been observed to appear in fused silica samples [136], likely due to defects. Low-mobility charge carriers could be responsible for the long-term test mass polarization.

- **Water monolayer:** A monolayer of water will exist on most surfaces inside the chamber [137]. Charge carriers in this monolayer can transport charge around the surface of the optic. The rate of polarization would then be expected to depend on the cleanliness or humidity, which has been observed in tabletop experiments [138].

- **Cosmic rays:** The effect of air showers from cosmic rays on the detector sensitivity has been investigated previously and was found to be negligible [139]. However, these showers can in principle deposit charge on the optic. This might be observed as a net charge on the optic after remaining in the chamber for a long period; tabletop experiments have shown that charge remains on the isolated fused silica indefinitely [140]. Like the UV photons above, long-term measurements of the optic *net* charge are required to rule this out.

It should be noted that every time the vacuum system is pumped down any large electric fields will be discharged by virtue of a decreasing dielectric breakdown voltage of air at lower pressure. This effect is discussed further in section 4.6.4.

4.4 Measuring Charge Distribution

While the model laid out in section 4.2.1 has been used successfully to explain the observed actuator strength and subsequent changes, this work is the first to comprehensively investigate the effect of optic charge and stray electric fields on interferometer sensitivity.

For simplicity, let's set the driver common voltage as our reference voltage, i.e., define the potential of the ESD driver common as $V = 0$. We will separate the voltages in equation (4.2) into offsets and fluctuating components:

$$V_i = V_{i,0} + \delta V_i.$$

With this simplification, we can write the net force on the optic in terms of the

fluctuating terms

$$\begin{aligned}\delta F =&\delta V_s \left[2(\gamma-\alpha)V_{b,0} + 2(\alpha+\gamma)V_{s,0} + \beta - \beta_2\right] \\ &+\delta V_b \left[2(\alpha+\gamma)V_{b,0} + 2(\gamma-\alpha)V_{s,0} + \beta + \beta_2\right] \\ &+\delta V_{\text{ref}} \left[-4\gamma(V_{b,0}+V_{s,0}) - 2\beta\right],\end{aligned} \tag{4.3}$$

where we have assumed no DC offset on the cage and have dropped constant terms and those of $\mathcal{O}(\delta V_i^2)$. In the nominal operating configuration, the signal trace also has no offset ($V_{s,0} = 0$), but we include it for full generality[6]. Driving the bias and signal traces in phase is equivalent to driving the cage with opposite sign.

Setting $V_{s,0} = 0$, we can rewrite equation (4.3) as

$$\begin{aligned}\delta F =&\delta V_s \left[2(\gamma-\alpha)(V_{b,0} - V_{\text{eff}})\right] \\ &+\delta V_b \left[2(\alpha+\gamma)(V_{b,0} - V_{\text{eff},2})\right] \\ &+\delta V_{\text{ref}} \left[-4\gamma(V_{b,0} - V_{\text{eff,gnd}})\right],\end{aligned} \tag{4.4}$$

where we have defined the following "effective voltages":

$$V_{\text{eff}} \equiv \frac{-(\beta-\beta_2)}{2(\gamma-\alpha)} \tag{4.5}$$

$$V_{\text{eff},2} \equiv \frac{-(\beta+\beta_2)}{2(\alpha+\gamma)} \tag{4.6}$$

$$V_{\text{eff,gnd}} \equiv \frac{-\beta}{2\gamma}. \tag{4.7}$$

This form makes clear that the strength of each excitation is linear in the bias voltage with some offset due to non-uniform charge arrangement. The slope of the bias voltage dependence is determined by geometry and the dielectric properties of the optic and is independent of the (immobile) charge on the optic.

By driving the signal and bias traces at a number of bias offsets, the slope and effective voltages of the first two terms can be determined. Figure 4-7 shows the results

[6]While we are not able to put a large voltage offset on the signal traces in low-noise operation, we have checked that we obtain similar values for α and γ when using the signal trace as a bias and driving the bias path instead. Note that, as expected, the signal and bias traces have analogous relations to α and γ; the only difference is related to the charge distribution nearest the ESDs (β_2).

of these measurements. While the ESD actuator strength is continuously monitored and has traditionally been observed to increase in time (figure 4-4), by occasionally measuring each actuator coefficient we can better understand the changing charge distribution on the optic and keep updated models of actuator and stray field noise.

Prior "charge" measurements were performed by driving individual quadrants of the ESD at full range and high voltage and measuring the angular response using optical levers outside the vacuum chamber. While these measurements do not require a locked interferometer, they are generally noisier than the in-lock measurements. They can measure an "effective voltage" that roughly tracks the changing actuator strength [117], but these values could not be used to make an estimate for the sensitivity impact of various electric field noises.

These length measurements have been scripted and can now be reliably run regularly, taking approximately 5 minutes per optic to measure all signal quadrants at once, and 15 minutes to measure all quadrants individually. Unlike the optical lever measurements, they provide a much more relevant measurement of the effect of charge on the ESD actuators. Additional long-term measurements are required to draw additional conclusions about charge motion on the test masses.

4.4.1 Calibration and Measurement Uncertainty

A note on the calibration for this measurement is required, especially considering the statistically significant spread between the measurements that depend on calibration. The uncertainties in figure 4-7 are the one-sigma uncertainties at each drive configuration as determined by the magnitude squared coherence between DARM and the drive; the uncertainties in figures 4-8 and 4-9 are determined based on the goodness of the linear fits of these measurements. They do not include calibration uncertainty. Certain values do not depend on calibration, like zero crossings (effective voltages) or actuator coefficients scaled by α, and should provide a more accurate view of the variation in time.

The actuation force is calibrated using the full interferometer DARM calibration [142, 143] and the mechanical response of the test mass, using the free mass

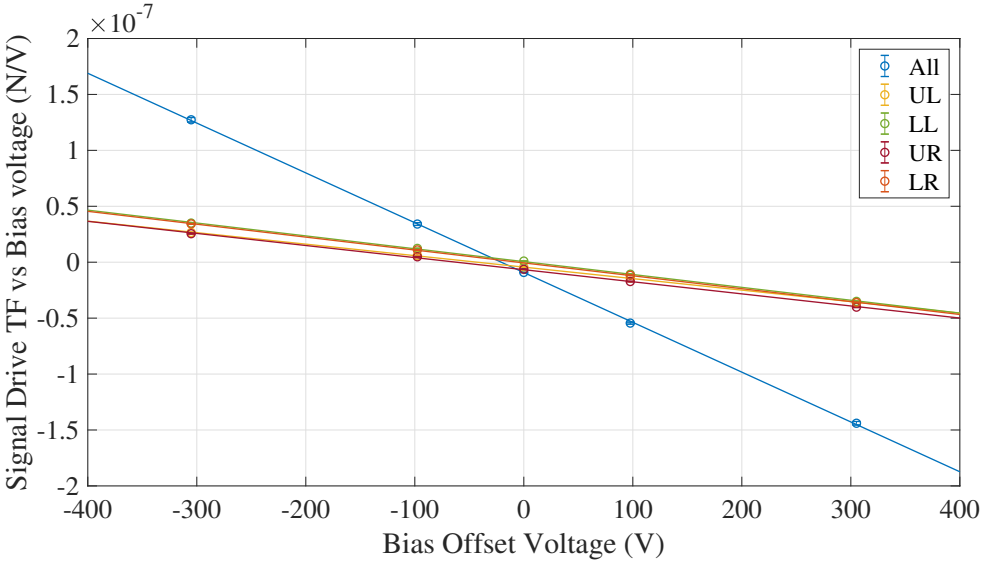

(a) ESD signal trace actuation strength as a function of bias trace voltage. This measurement was also performed with the individual signal trace quadrants (shallow lines) to demonstrate some variation between quadrants.

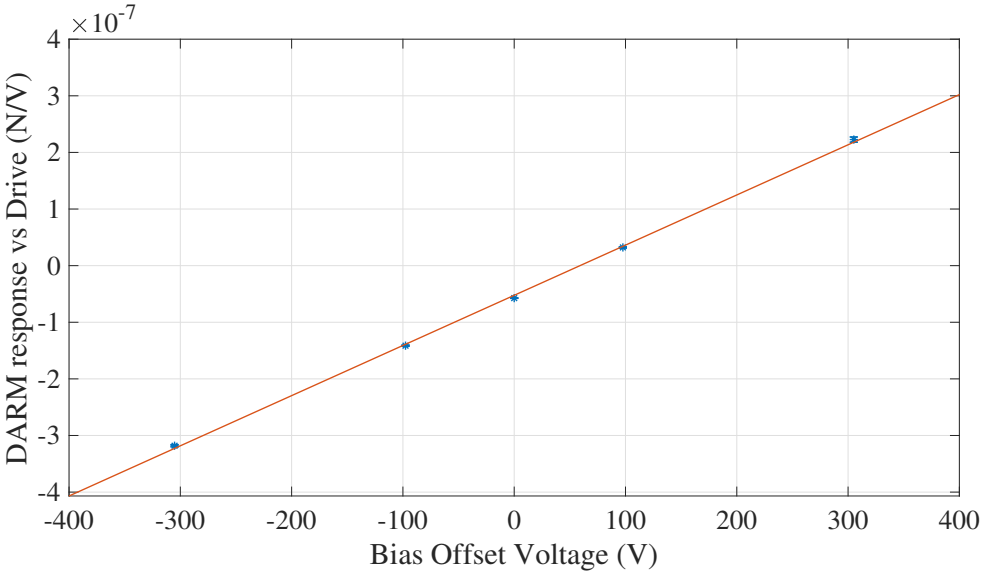

(b) Bias trace actuator strength as a function of DC bias voltage.

Figure 4-7: Example results of a charge measurement showing the linear dependence of the (a) signal and (b) bias actuation vs. DC bias voltage for ETMX. Note that in normal operation the bias trace is held at constant voltage, but it can be used as a diagnostic.

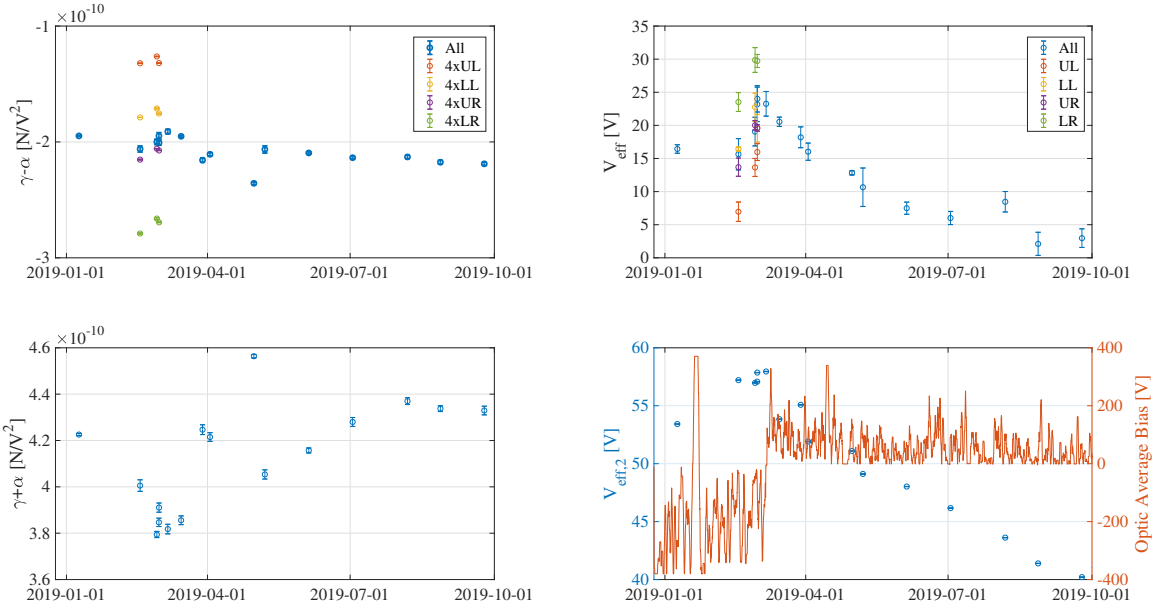

(a) "Raw" measured ESD actuator terms over time. Where relevant, the corresponding terms for the individual quadrants of the signal trace are also included.

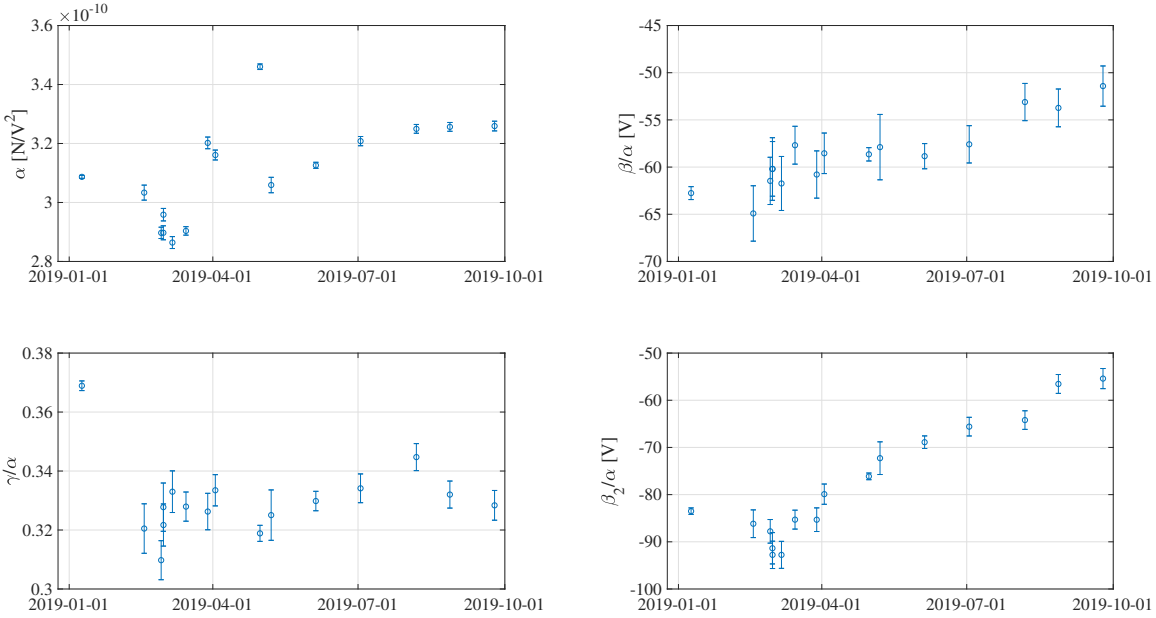

(b) Derived ESD actuator coefficients from the measurements in (a). To avoid uncertainties due to calibration, these coefficients are scaled by α.

Figure 4-8: Trends of the ESD actuator coefficients for ETMX at LLO, including a moving average (24-hour window) of the applied bias voltage. Note that at LLO, the ETMX ESD is only used for lock acquisition before DARM control is transferred to ETMY for low-noise operation and the ETMX bias is set to zero. As a result, the bias voltage is nonzero on average.

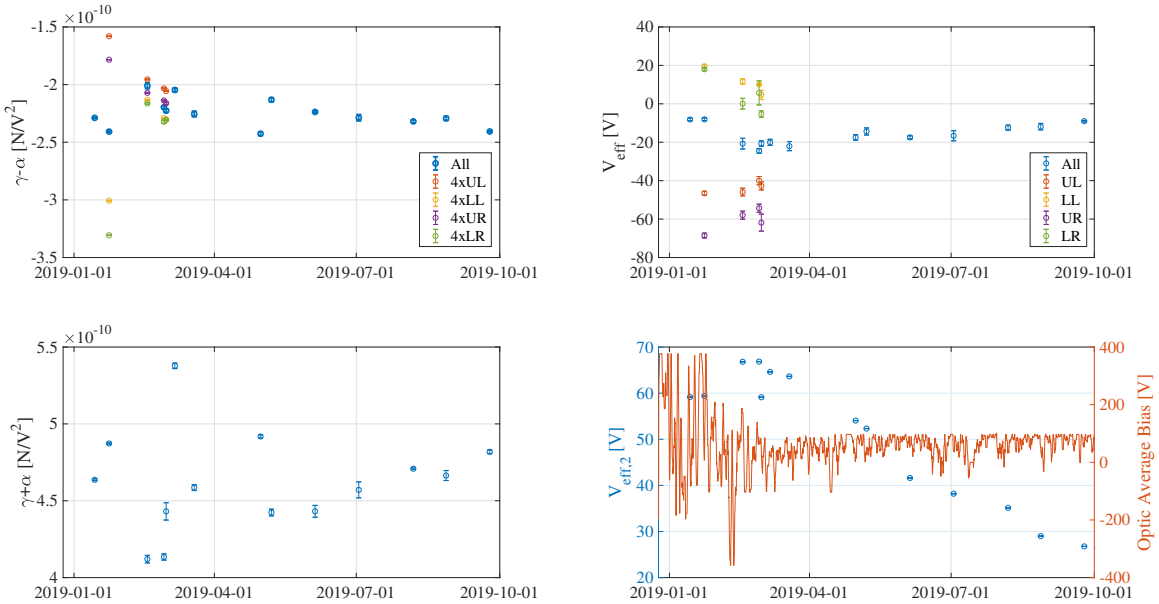

(a) "Raw" measured ESD actuator terms over time. Where relevant, the corresponding terms for the individual quadrants of the signal trace are also included.

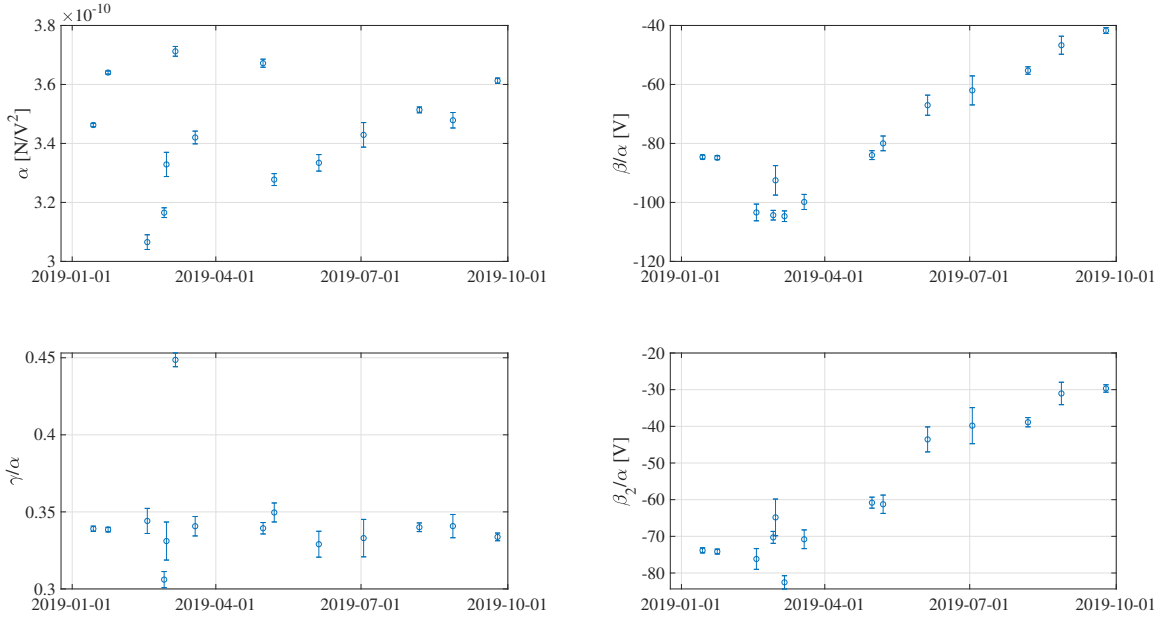

(b) Derived ESD actuator coefficients from the measurements in (a). To avoid uncertainties due to calibration, these coefficients are scaled by α.

Figure 4-9: Trends of the ESD actuator coefficients for ETMY at LLO, including a moving average (24-hour window) of the applied bias voltage. Note that at LLO, the ETMY ESD is only used for low-noise DARM control. As such, when the interferometer is down, the bias voltage can be set opposite to the value used for low-noise operation to maintain an average bias voltage close to zero. The sudden change in spread of the quadrant actuation strength was due to a change in beam spot position on the optic in February [141].

approximation[7]. The details of this system can be found in [143], but we review the basic features here.

This system uses the known radiation pressure continually applied to a test mass using an auxiliary laser system (the "photon calibrator" [144]) to convert the optical readout to meters or strain. Additionally, the control signals used to keep the interferometer locked must also be compensated to produce an accurate out-of-loop strain signal. This is done by driving each stage of the suspension at a different frequency and comparing the strength of these excitations in DARM with the photon calibrator (Pcal) line to estimate the strength of each actuator stage. By fitting these results to a model of the suspension actuators, the full DARM control loop can be estimated to generate a calibrated signal. Before O3, an online system to measure these actuator coefficients was installed at the sites, allowing for real-time calibration updates at each observatory[8]. At LLO, the standard operating condition uses the ETMY ESD for DARM actuation. However, to perform actuator measurements on ETMY, DARM control must be transferred to the ETMX ESD. Because these actuator strengths are not matched exactly, it takes some time for the calibration to settle when handing off DARM control. In the interest of using as little observing time as possible, these scripts have been optimized to continue once the calibration is sufficiently settled but still changing, which can affect the final measurement and introduces some excess spread in the data.

Other effects can also introduce excess variation in these measurements. The strength of the ESD actuation is strongly dependent on the relative distance between the main suspension chain and the reaction chain [132, 133], which can change from lock to lock. Large transients during measurements can confuse measurements, though a minimum coherence between drive and DARM is required to use a measurement. Finally, between observing runs, small tweaks are made to the calibration system that can change the results of these measurements.

[7] At 50 Hz the difference between the free-mass approximation of the mechanical response $\hat{x} = -\hat{F}/m\omega^2$ differs from the susceptibility calculated using the full quad model by less than 1%.

[8] Note that a similar system had already existed for observing data but was not easily accessed from the control room. Additionally, this new system was not used for gravitational wave signals, but merely as a convenience for commissioning.

Nevertheless, this variation between measurements is small enough that a long-term trend is clearly visible. Additionally, because the actuator strength is constantly being measured, the fact that it is slowly changing in time is not a major impediment to instrument operation. As long as the control loops remain stable, the calibration system will compensate for the changing actuation strength, and when the effective bias voltage reaches a certain level, the running bias voltage sign can be flipped.

4.5 Interferometer Sensitivity

Now that these actuation coefficients are measured and tracked over time, this allows us to predict the various noise contributions that limit sensitivity. To make an accurate prediction, we must understand the following sources of noise:

- **ESD driver signal noise:** The ESD driver is an analog gain/filter stage that takes as input the signal from a digital-to-analog converter (DAC) and sends the signal to the ESD inside the vacuum chamber. The driver can be switched remotely from a high-range, high-voltage configuration used for lock acquisition to a low-noise, low-voltage configuration for normal operation. This latter configuration has additional analog filters to eliminate excess DAC noise at higher frequencies. This noise can be estimated by measuring the voltage noise of each signal quadrant relative to the driver chassis. Note that each signal quadrant is driven by its own channel in the driver (and corresponding DAC channel) [9].

- **ESD driver bias path noise:** The bias path is similar to the signal path, however, even in the low-noise configuration it uses a high-voltage amplifier. Because this path is normally held at a constant offset, the driver has an analog second-order low-pass filter around 1 Hz to filter this excess noise. Prior to this work, the input signal was not otherwise filtered.

[9] Because these channels are independent, there is $\sqrt{4}$ greater noise coupled into DARM as compared to a single signal trace and driver.

- **Driver reference voltage fluctuations:** While the vacuum chamber and suspension cage are electrically connected to the ESD driver chassis, any voltage fluctuations between these surfaces will generate a fluctuating electric field. This noise is estimated by performing a differential measurement of the ESD driver chassis voltage and the chamber voltage. We explore the origin of this effect below.

The first two sources are relatively straightforward and are due to unavoidable noise from the driver electronics. However, it's not obvious why the driver reference voltage will fluctuate relative to the chamber, especially if they are electrically connected. Understanding this requires some knowledge of the general electronics layout and design at each observatory.

Generally, each custom circuit is designed with differential inputs or outputs with a common shield. To avoid ground loops, where possible the shield is only connected at one end of the cabling, in accordance with best practices [145]. This configuration has the advantage of reducing electronic pickup when traversing the long paths between electronics racks and the vacuum systems. Additionally, the differential signals avoid the issue of having the vacuum chamber, analog electronics, and digital systems referenced to separate common voltages.

A few issues come up in practice. First, for safety reasons the beam tube is grounded at a number of points along its length. Due to differences in the local ground potential, this generates currents through the beam tube as large as 6 A [146], which produce local potential differences even between points that are nominally "ground." Second, for various reasons, additional unintended ground connections exist in the system, sometimes inside the vacuum chamber. Going through such a complex system to find and remove these connections would be prohibitively time consuming. Thus, it's often best to understand how these connections affect the performance of the interferometer and fix the most pressing issues. Lastly, while the ESD signal and bias traces are driven from the same analog driver with filtered output (and thus have very low differential noise), the reference voltage can fluctuate relative to the cage surrounding the optic. This will produce an electric field between the ESD and

the cage that isn't filtered and couples to DARM via the γ and β parameters (see last term of equation (4.3)). Figure 4-10 demonstrates the mechanism for this noise source.

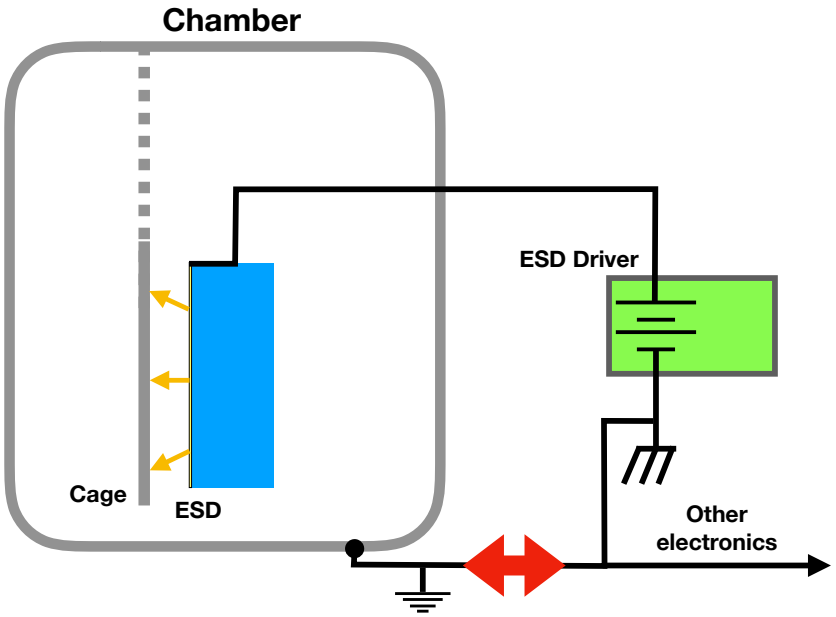

Figure 4-10: Schematic demonstrating how ground currents can produce a voltage differential between ESD driver common and the chamber, producing an electric field between ESD traces and the suspension cage. Further, analog filters in the driver, which will remove output voltage fluctuations relative to the driver reference voltage, will not eliminate this source of noise.

By measuring the voltage fluctuations between the chamber and driver, the contribution to DARM can be estimated using the measurement of the ESD actuation coefficients. Fortunately the ground noise fluctuations at the end stations, where a large bias voltage is applied during observing runs, were small ($<100\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$). However, we observed significant voltage differences between the corner station vacuum system and the ESD driver references, with DC voltage differences of 60 mV and fluctuating voltages of $5\,\mu\mathrm{V}/\sqrt{\mathrm{Hz}}$.

Because the chamber and ESD driver reference are nominally connected electrically, it is challenging and disruptive to accurately measure γ directly. However, as seen in figure 4-11, our estimate of γ is sufficient to use as an estimate of the noise contribution. These results validate this proposed mechanism as the cause of this

noise, and suggest this model of ESD actuation is sufficiently accurate.

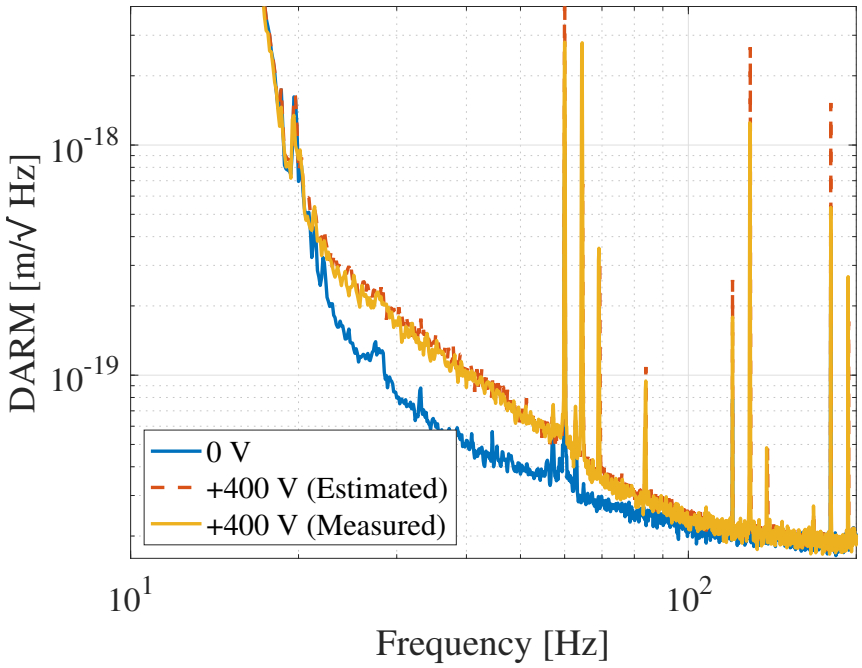

Figure 4-11: Test of the coupling of voltage fluctuations between vacuum chamber and ESD driver common to DARM using measured value of γ. The solid traces show the DARM spectrum without a bias voltage on the ITM ESDs (normal operating conditions) or a 400 V bias voltage on ITMX, while the dashed line is the predicted spectrum for the latter configuration. The prediction is calculated using measurements of the voltage fluctuations between vacuum chamber and ESD driver common and the measured γ and β actuation parameters.

The easiest way to mitigate this noise source is to minimize the resistance between the ESD driver reference and the chamber. Indeed, when the driver was directly connected to the chamber, this noise was reduced by a factor of 2 [147]. Removing electronics that share the ESD driver rack (and therefore reference voltage) can also eliminate this noise, and this likely explains the reduction of spectral lines in [148]. If the driver is referenced to the chamber and all other grounding points are removed, this noise should be eliminated. However, modifying the corner station electronics to do this would be prohibitively time consuming, and because the ITM ESDs are not required for full interferometer operation, these ESD traces were disconnected from the driver and were shorted to the chamber.

Similar predictions for the end test masses were significantly lower than the mea-

sured bias-dependent broadband noise (see figure 4-5 on page 89). Shorting the ESD driver reference to the chamber had no effect. Additionally, there was no coherence observed between chamber-driver voltage fluctuations and DARM when a large bias voltage was applied and the excess noise appeared, suggesting an altogether different noise source.

Measurements of the noise of the suspension actuator DACs (General Standards 18AO8) for various drive amplitudes [149] in most cases found a broadband noise floor around $1\,\mu V/\sqrt{Hz}$, which was sufficient for the required performance. There was no static offset voltage for these measurements. However, we found that the noise of the DACs increased as the offset voltage was increased, which produced a corresponding increase in the HV driver noise (figure 4-12). The origin of this noise is not fully understood, but it is likely intrinsic to the hardware and not due to the LIGO digital system [150]. Because it is approximately linear in offset, it is assumed to be amplified noise of the DAC reference voltage. The expected contribution of this noise to DARM (using the relevant actuator coefficients in (4.3)) matched the measured value. An additional analog low-pass filter between the DAC and high-voltage driver eliminated this bias-dependent noise and lowered this noise contribution far below current sensitivity (figure 4-24).

Finally, it should be noted that a similar dependence is observed for voltage noise of the ring heater:

$$\delta F_{\mathrm{RH}} = \delta V_{\mathrm{RH}} \left[A(V_{b,0} - V_{\mathrm{eff,RH}}) \right] \tag{4.8}$$

Due to the differential drive of the ring heater, this coupling is significantly smaller than that due to the other terms in equation (4.4) and such measurements cannot currently be non-invasively performed in the normal low-noise configuration. To make these measurements, the ring heater current driver must be disconnected, and all quadrants are driven together. An example of this measurement at ETMX is shown in figure 4-13, and a compilation of these measurements is provided in table 4.1. Investigating the time dependence of this coupling is one way to better understand the charge distribution on the optic. The broadband RH driver noise sits around

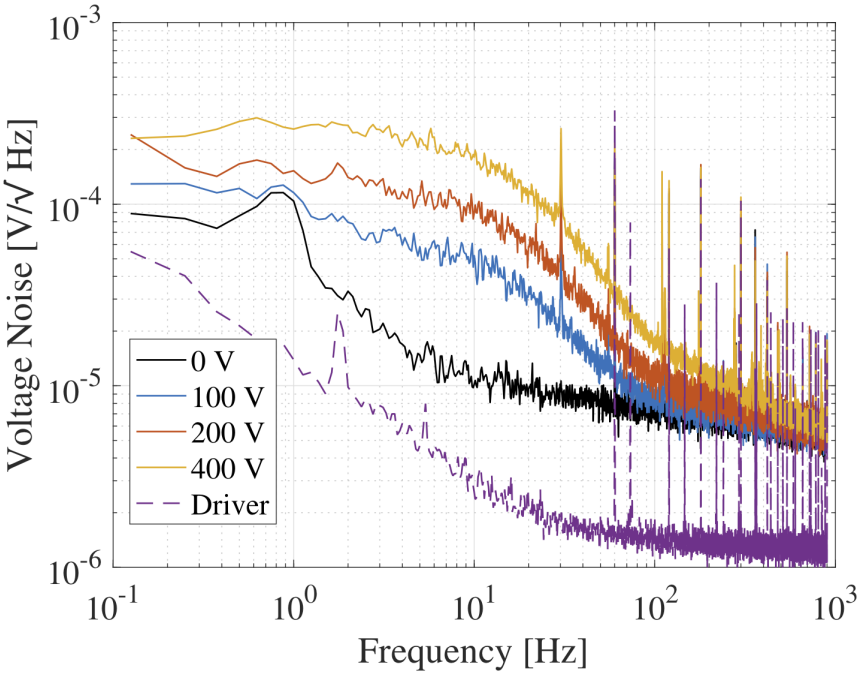

Figure 4-12: ESD high-voltage driver noise as a function of output voltage using General Standards 18-bit DACs. As the DC bias voltage is increased, the noise grows to more than an order of magnitude larger than with no offset. Also included is the driver noise with a low-noise input voltage source. An additional low-pass filter was installed at the driver input to eliminate this DAC noise.

$1\,\mu V/\sqrt{Hz}$ [151], and an estimate of this coupling to DARM is given in figure 4-24.

This model suggests a way to mitigate the effect of excess noise from one particular source. Setting the bias voltage to one of the effective voltages defined in equations (4.5) will eliminate the associated noise from coupling to DARM, and could be used to eliminate excess noise on one electrode. Fortunately, at LLO there was not excess broadband noise entering at the end stations, so this technique was only able to minimize lines in the spectrum [152] but was not run in this configuration for observation.

4.5.1 Electrostatic Drive Transients

While most of the focus has been on the stationary noise performance, because the LIGO detectors are looking for transient astrophysical events, we must also consider how charging events can produce transients in the detectors. Unfortunately, a lack

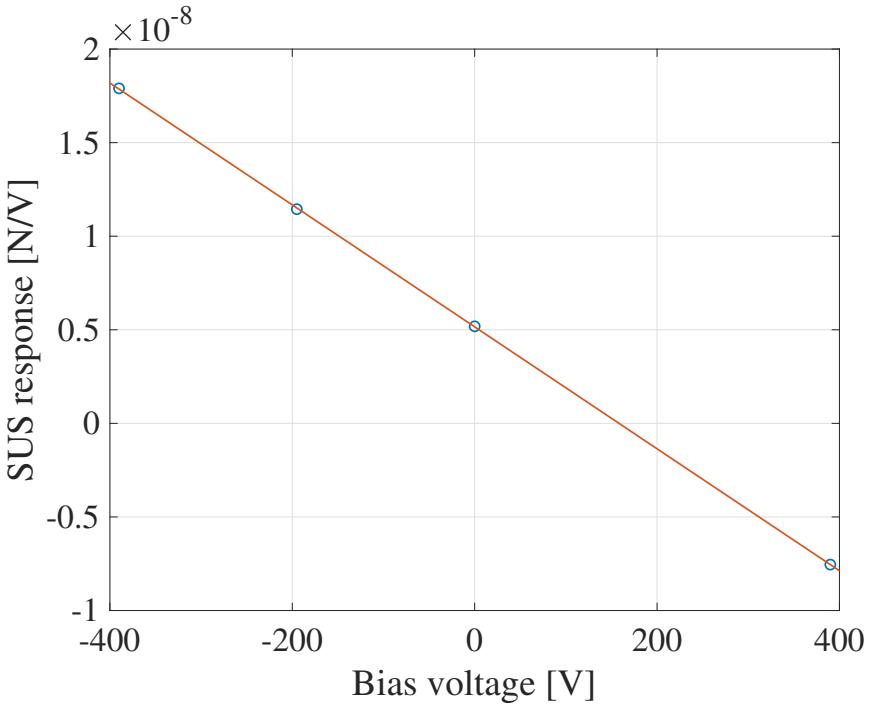

Figure 4-13: Actuation strength of ETMX ring heater with all quadrants driven together vs. bias voltage. Note that the strength of this coupling is about an order of magnitude smaller than from the ESDs, though the voltage noise on the RH drivers is higher than on the ESDs.

of monitor equipment makes direct measurements of charge flow on the test mass difficult. However, there exists some evidence that discharging events can produce transients in DARM. Small-scale experiments with a fused silica mass and mock ESD have observed charge transfers even under vacuum [153]. Figure 4-14a shows the effect of slowly ramping the ESD bias voltage from lower to upper voltage limit and the associated transients in a filtered DARM time series. These transients can be deterministically produced by ramping across the range of the ESD driver voltage and seem to persist for tens of seconds after the ramp has finished. In particular, repeatedly ramping between 0 V and one voltage limit does not produce glitching. This behavior strongly suggests some sort of dielectric breakdown and current flow.

Of note, the ESD monitor circuit is before the low-pass filter, and when these transients appear there is not a corresponding signal in the monitor. In fact, even when there is a glitch in the monitor channel, there are no associated transients in DARM (figure 4-14b). This suggests some sort of discharge occurring after the ESD

Optic	A [N/V^2]	$V_{\mathrm{eff,RH}}$ [V]
ITMY	-5.7×10^{-12}	3.4
ETMX	-3.3×10^{-11}	158
ETMY	4.24×10^{-11}	286

Table 4.1: Ring heater coupling to DARM using parameters from equation (4.8). Note that this measurement was not made for ITMX but it is assumed to be comparable to ITMY and small compared to the ETMs. Additionally, because the effective ring heater bias voltage is very small, the coupling from the ITMs is even smaller than for the ETMs.

driver electronics, either in the cable leading to the chamber, between the ESD traces, or charge hopping on the test mass. Following this discovery, direct measurements of the output of spare HV drivers were unable to observe this repeatable glitching. Additional work is required to understand the origin and whether they are directly related to transients seen during observing periods [154].

Running the Gravity Spy transient identification system [155] on these glitches showed that they closely matched the "tomte" class of glitches observed in prior observing runs (figure 4-15). Such glitches occurred 1.3 times per day of observing time during O2 and produced a large background in the high-mass region of the binary black hole (BBH) template bank. Using a narrow Kaiser window to isolate the glitch and then produce a spectrum of the transient, we find they have approximately a $1/f^3$ spectrum. This is consistent with a discharge-induced step function excitation ($1/f$) filtered through the inertial response of the test mass ($1/f^2$).

Since the bias voltage is held constant during normal operation, it is unlikely that the ESD is responsible for all the observed tomte glitches, however, this insight does provide an additional diagnostic tool. As an example, looking through past observing data for occurrences of this glitch morphology using Gravity Spy, an excess was found at LLO when a ring heater driver was found to be glitchy. The prevalence of this type of glitch in DARM without being seen on any auxiliary monitors would suggest faulty equipment near the test mass and may signal an impending equipment failure.

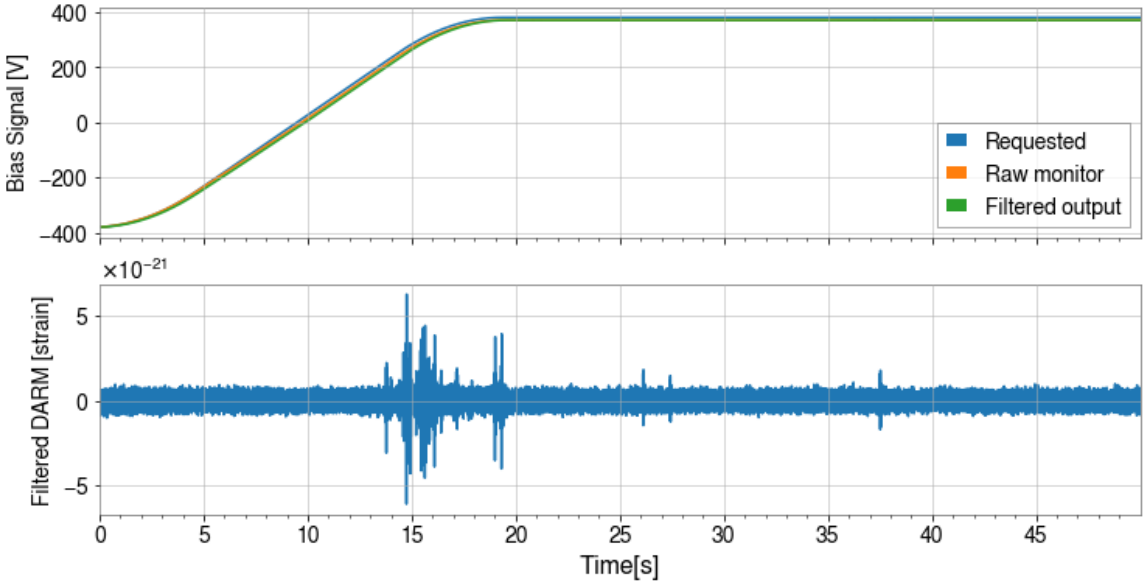

(a)

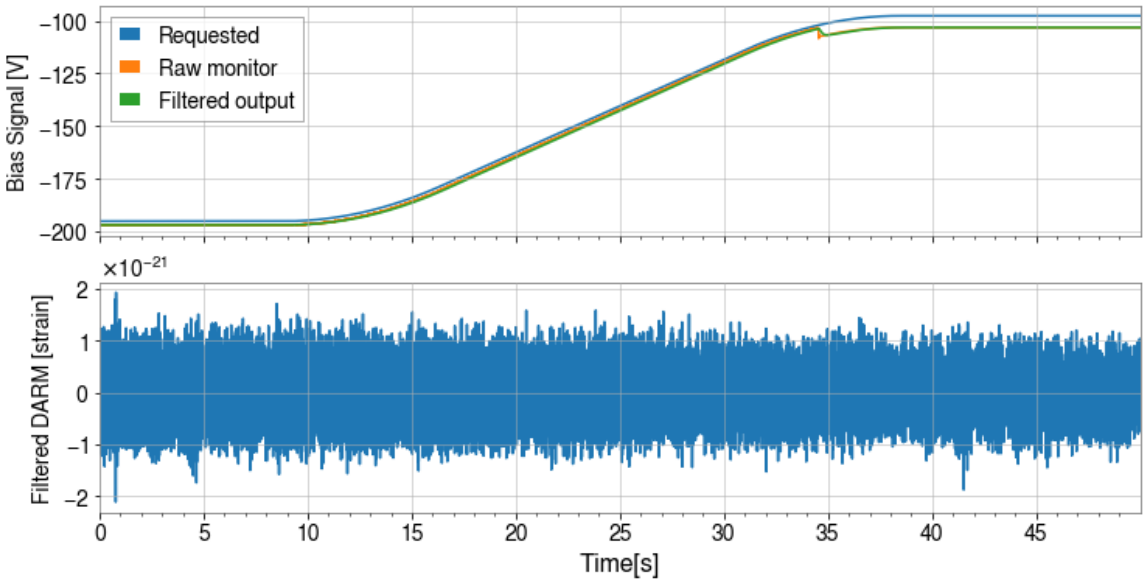

(b)

Figure 4-14: Induced "glitches" in DARM by ramping the ESD bias voltage. The upper plots show the requested output voltage, the output voltage as measured by the monitor circuit, and the estimated actual applied voltage to the ESD due to an analog second-order low-pass filter at 1.8 Hz. The lower plots show a DARM time series filtered by a bandpass filter with passband 30–250 Hz. Figure (a) shows these glitches appearing in DARM for a large bias voltage change for an otherwise smooth ramp, while figure (b) shows a bias ramp with a measured voltage step but no apparent glitch in DARM, suggesting discharge after the ESD driver electronics.

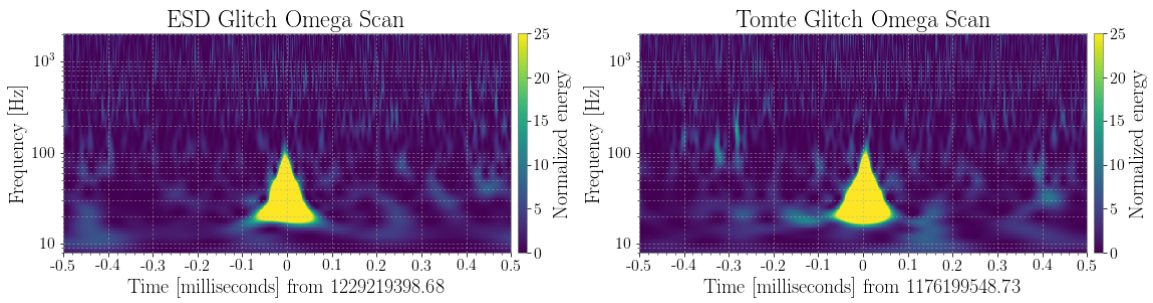

(a) (b)

Figure 4-15: Time-frequency representation based on the Q-transform ("Q-scan") [156, 157] of DARM during (a) a bias voltage ramp and (b) a "tomte" glitch seen during O2.

4.6 Electric Field Meter

Since the first LIGO observing runs, the detectors have suffered from unexplained low-frequency noise. Given the aforementioned experiences, a natural phenomenon to investigate was the effect of optic charging and stray electric fields inside the vacuum chambers as a source of this excess noise. Nominally the large steel vacuum chambers act as Faraday cages and isolate the optics from electric fields entering from outside the chamber. However, viewports and in-vacuum electronics can still allow electric fields to reach the test masses. To measure the magnitude of these fluctuations, an electric field meter (EFM) was installed at one end station at each site to measure the scale of these fluctuations. The electric field meter (EFM) consists of a cubic body held at the chamber potential and four "sensing plates" separated from the main cube by ceramic standoffs and a 1 TΩ resistor. Figures B-3 and B-2 show the mechanical drawings and sensing circuit of the EFM, respectively. The high impedance between plates allows for a differential measurement of the voltage on each plate relative to the body; because they are physically separated, this provides an estimate of the electric field around the EFM along each axis with common-mode cancellation. An additional set of "calibration plates" can be optionally attached to produce electric fields for calibration (section 4.6.2). The sensing plates can be shorted to the main body when the chamber is open to prevent accidental electrostatic discharge that may damage the sensitive electrometer components.

At LLO, the EFM was installed in the ETMY vacuum chamber (figure 4-16). It hangs from the ISI platform from above and sits at approximately the same height as the test mass but offset from the main beam. Viton straps hang from the ISI platform to damp its motion.

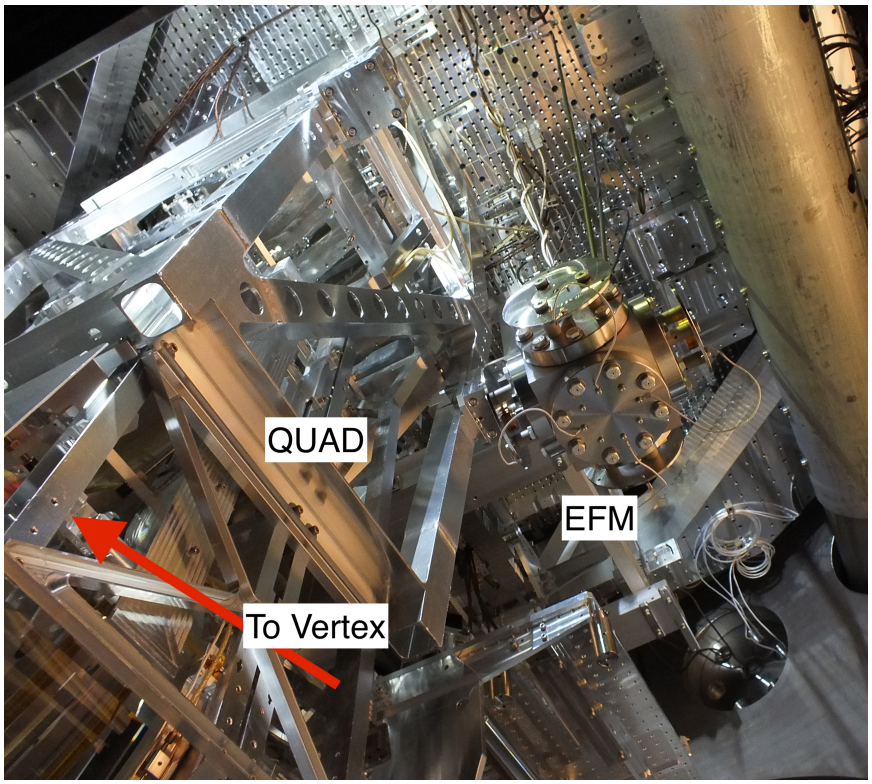

Figure 4-16: Electric field meter (EFM) in the ETMY vacuum chamber at LIGO Livingston Observatory (LLO) as seen from below. The EFM is mounted to the internal seismic isolation (ISI) platform by a steel rod with viton straps to damp its motion. The main cube is held at the same potential as the electronic racks, while the voltage on each plate is measured to estimate the electric field along the X and Y directions (as referenced by the interferometer). To the left is the cage surrounding the ETMY quad suspension. The arrow points down the Y arm toward the interferometer vertex.

4.6.1 Characterization and Sensitivity

The installation of new equipment into the vacuum chamber, especially near the arm laser beam, has the potential to become a new light scattering surface. If there is excess scattered light noise, the upper stage can be excited at the measured mechan-

Mode	Frequency [Hz]	Q Factor
Pendulum	0.597	50
Twist	4.715	207
Mode 3	5.872	178
Mode 4	3.994	?

Table 4.2: As-installed EFM mechanical modes. The signal for mode 4 was too small to measure the Q factor.

Axis	RC Time [s] / Sense-body Capacitance (inferred) [pF]
+X	17.8
-X	14.1
+Y	14.6
-Y	17.3

Table 4.3: EFM electrical properties. Because $R = 1$ TΩ, the RC time in seconds is equal to the capacitance between the sensing plate and EFM body in pF.

ical frequency of the EFM assembly to eliminate it as a source of scattered light. Additionally, these resonances can produce excess noise in the instrument readout. These modes and quality factors were measured by gently pushing the installed EFM cube in-air and fitting the ringdown of the various excited modes as seen in the EFM readout. The as-installed mechanical properties were measured and are listed in table 4.2.

Measurements of the RC ring down times were also performed and are presented in table 4.3. The EFM was excited by quickly moving a body part in the chamber far away from the EFM and then holding it until the ring down was complete.

At atmospheric pressure, the low-frequency spectrum is dominated by impulse (dis)charging events that disappear at low-pressure. Under vacuum, the noise curve is at the instrument noise floor, which is dominated at low frequencies by the input current noise of the input amplifier through the $10^{12}\,\Omega$ resistance (filtered by the RC low-pass filter). Indeed, when there is no ESD drive, neither axis is coherent with either DARM or each other, suggesting they are only seeing instrumental noise.

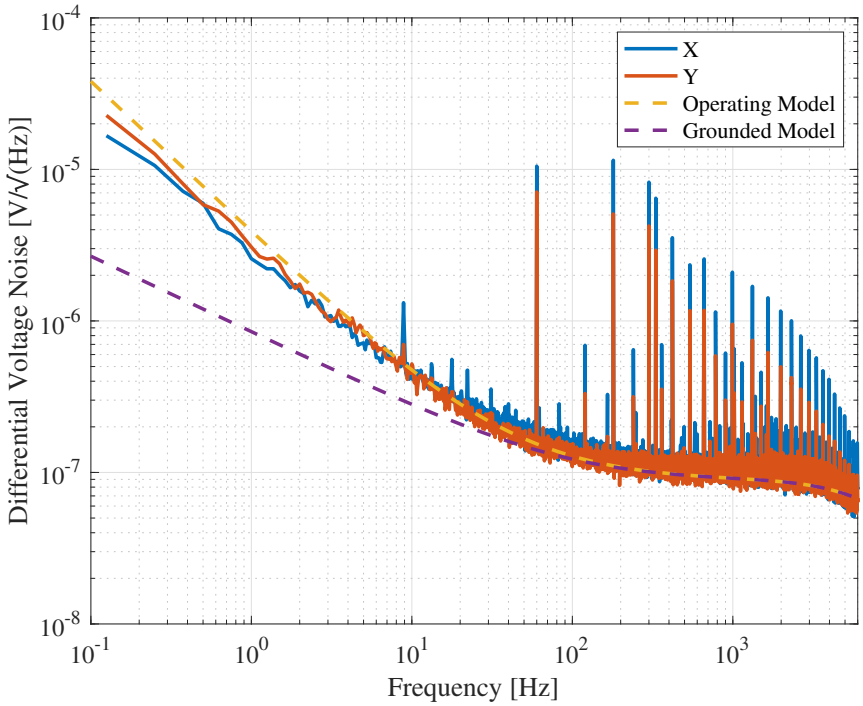

Figure 4-17: EFM noise spectrum under vacuum and LTSpice [158] models of the expected noise from the sensor electronics alone. Once under vacuum, the broadband noise floor matches the expected noise floor; spectral lines are mostly harmonics of the power mains. For reference, a noise model of the EFM with the sensing plates shorted to the body is also included.

4.6.2 Calibration and Readout

The (differential) output of each axis is sent out of the chamber and into the standard aLIGO digital acquisition system. A set of calibration plates can be attached to opposite sensing plates and driven together to measure the common-mode rejection ratio (CMRR) along each axis, while an internal digital potentiometer is set to maximize this ratio. For both axes, >26 dB CMRR was achieved.

Because it is a differential measurement between two plates, the electric field is then given by [159]

$$E = \frac{V_{\text{diff,plates}}}{2d} = \frac{V_{\text{diff,out}}}{2dG} \tag{4.9}$$

where $V_{\text{diff,plates (out)}}$ is the plate (output) differential voltage, d =1/2 inch is the sensing plate separation from the main cube, and $G = 15.8$ is the electronic gain. All data in this work are given in terms of differential voltage as measured at the plates or

measured electric field.

An additional digital channel to approximately whiten the EFM signal was deployed and could be used to easily measure charge transients.

4.6.3 Electric Field Injections and Coupling to DARM

The broadband magnitude-squared coherence [1] between DARM and the EFM signals is $< 10^{-2}$ above 10 Hz in normal operation (figure 4-18). Nevertheless, we can attempt to set upper limits on the noise due to stray electric fields by injecting electric fields, as described in [32]. A viewport cover was modified to include an additional isolated metal plate that could be driven relative to the chamber and inject electric fields into the chamber (figure 4-19).

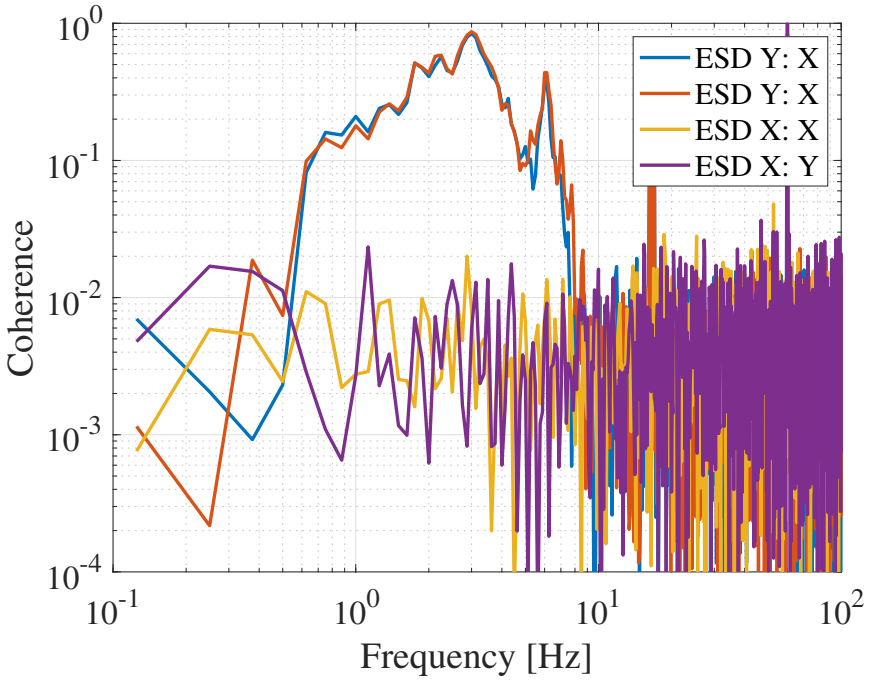

Figure 4-18: EFM signal magnitude squared coherence with DARM when actuating with the ETMX or ETMY ESD. The signal is only marginally coherent with DARM at low frequencies when the ETMY ESD is used (and limited by statistical noise elsewhere), suggesting that the only electric field that the EFM measures is coming from the ESD. Further, above 10 Hz there is no broadband coherence in either case.

Because the EFM is only installed in one chamber, we are unable to make precise estimates for this noise term. Additionally due to the location of the EFM relative to

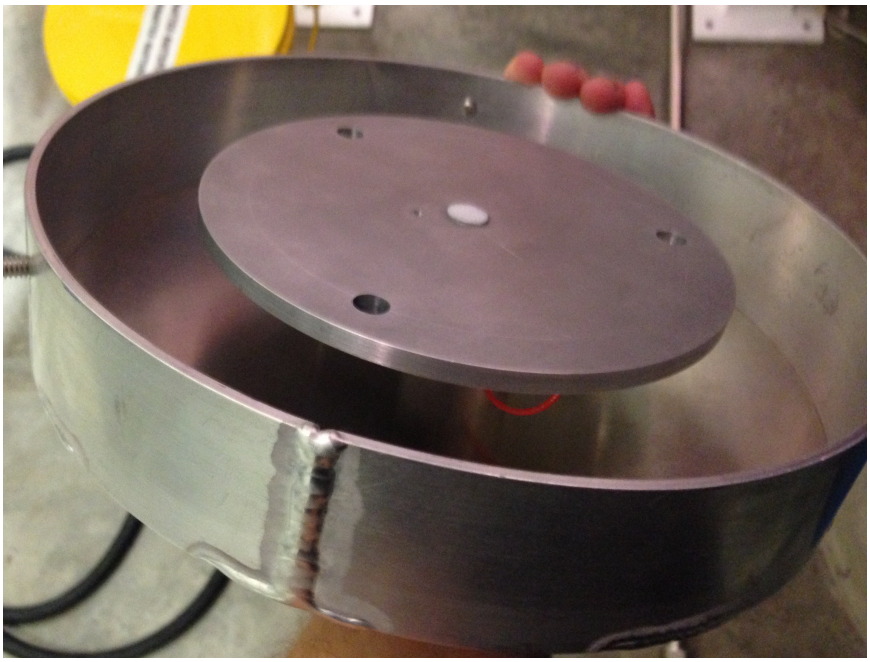

Figure 4-19: Viewport "capacitor" used to inject electric fields into the vacuum chamber. The center plate can be driven at high voltages relative to the outer portion (and the rest of the vacuum chamber).

the optic, for this estimate to be accurate, we must assume that the electric injected at the viewport is typical of those that can move the optic and is fairly uniform. The viewport is ∼20 cm in diameter and is a few meters away from the test mass, which is surrounded by metal on most sides, so this is not an unreasonable assumption. The viewport capacitor was placed on the viewport closest to collinear with the arm cavity to produce a field parallel to the arm and generate the largest coupling to DARM.

Fortunately, the EFM is a good monitor of the in-chamber electric field and the coupling to DARM is low. Broadband injections produced no signal in DARM, so the capacitor was excited at 50 Hz with 100 V peak-to-peak, far larger than any fluctuating voltage that would normally exist near an uncovered viewport. Figure 4-20 shows this excitation appearing in both EFM axes and orders of magnitude above the broadband EFM noise without appearing in DARM[10]. This demonstrates that the EFM is a good witness for large-scale electric fields, and fortunately, the coupling

[10]Smaller bandwidths (i.e., longer measurements) would produce an even larger signal in the EFM and might have revealed a signal in DARM, however, the measurement time was somewhat limited due to the presence of large transients in DARM ("glitches") that raised the spectrum far above the normal sensitivity.

of these fields to DARM is small.

Only when the bias voltage was maxed out (400 V) did a small signal appear in DARM at the excitation frequency. For a few reasons, this excitation is believed to be coupling to DARM via one of the local fluctuating voltages from equation (4.3) rather than from the bias voltage enhancing the coupling of large-scale electric fields. First, as described previously, the injected field should be fairly uniform near the test mass, and a bias voltage would not change the coupling of homogeneous fields. Second, for safety reasons the high-voltage amplifier used for the viewport injection is driven relative to ground, which can produce potential differences between the chamber and ESD driver reference. Finally, this excitation showed up in a number of unrelated monitors such as the ESD power supply voltage monitors, even when the HV amplifier was powered off. This suggests this injection was coupling to DARM via other mechanisms.

As a result, we believe the net charge on the optic is low and we can only set upper limits of the effect of large-scale in-chamber electric fields on DARM. Using the measured calibrated transfer function between DARM and the calibrated EFM signal, we can estimate the net charge of the optic. In the frequency domain, we have

$$\hat{x} = -\frac{\hat{F}}{m\omega^2} = -\frac{q\hat{E}}{m\omega^2}. \tag{4.10}$$

For a measured transfer function of $\hat{x}/\hat{E} \approx 1.8 \times 10^{-23}$ m/(V/m) at 50 Hz, we obtain an effective ETMY charge of 7×10^{-11} C. The coherence of this measurement was low (figure 4-20b) and the phase was not as expected ($-130°$), lending further evidence that such excitation likely entered via the ESD. As such, this should be considered an upper limit on the net charge of the optic. Note that this "effective" net charge would be different depending on the location of the charge, as the dielectric test mass would partially shield the front surface of the optic from these injected electric fields. However, the fact that both axes of the EFM see this injection at comparable levels provides confidence that the injected field is fairly uniform near the test mass.

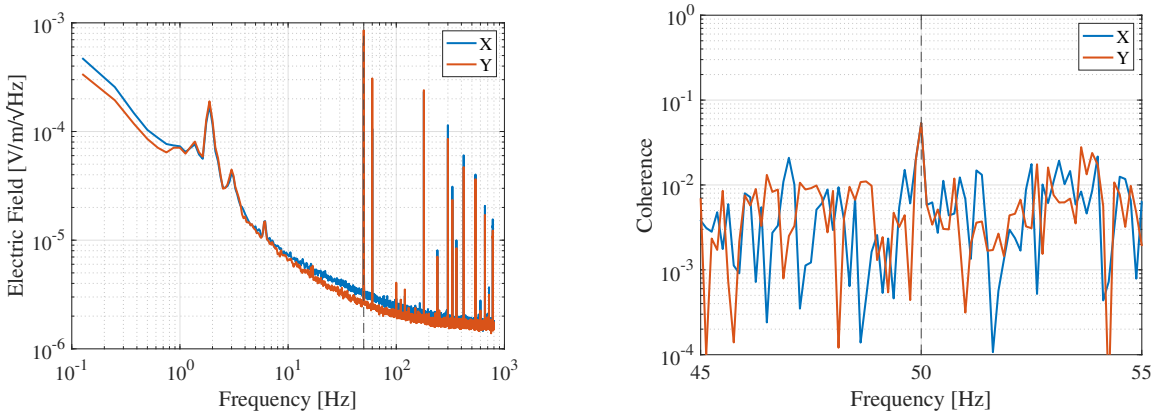

(a) (b)

Figure 4-20: (a) EFM spectrum while injecting a signal at 50 Hz into the chamber via the viewport capacitor, and (b) magnitude-squared coherence with DARM during excitation. This injection appears orders of magnitude above the EFM noise floor and still the signal is minimally coherent with DARM, suggesting low net charge on the test mass.

4.6.4 EFM Transients

While the EFM was not sensitive enough to observe the ambient in-chamber electric fields, it did provide some clues about the origin of charge on the test mass. Figure 4-21 shows a time series for the EFM after the chamber doors were replaced after a chamber incursion but before the chamber was pumped down. There are clear discontinuities in the raw time series, which suggest the accumulation of charge on the sensing plates. Unfortunately, due to the differential nature of the EFM, we are unable to say whether these are discharging or charging events, but because they seem to appear after a vacuum incursion they are hypothesized to be due to particulates sticking to the plates, having been stirred up and possibly triboelectrically charged. The known capacitance of the plates can be used to estimate the amount of charge transferred during each event $q = CV$, with $C \approx 15\,\text{pF} = 15\,\text{pC/V}$. With this relation, these individual transients transfer 100's to 1000's of e to the EFM plates, which is consistent with particulate matter charged by friction [160]. Of course, we cannot definitively state that this is the origin of these transients, as the EFM may instead be sensitive to discharge events further away from the sensor.

During the evacuation of the vacuum chamber, similar but larger transient events

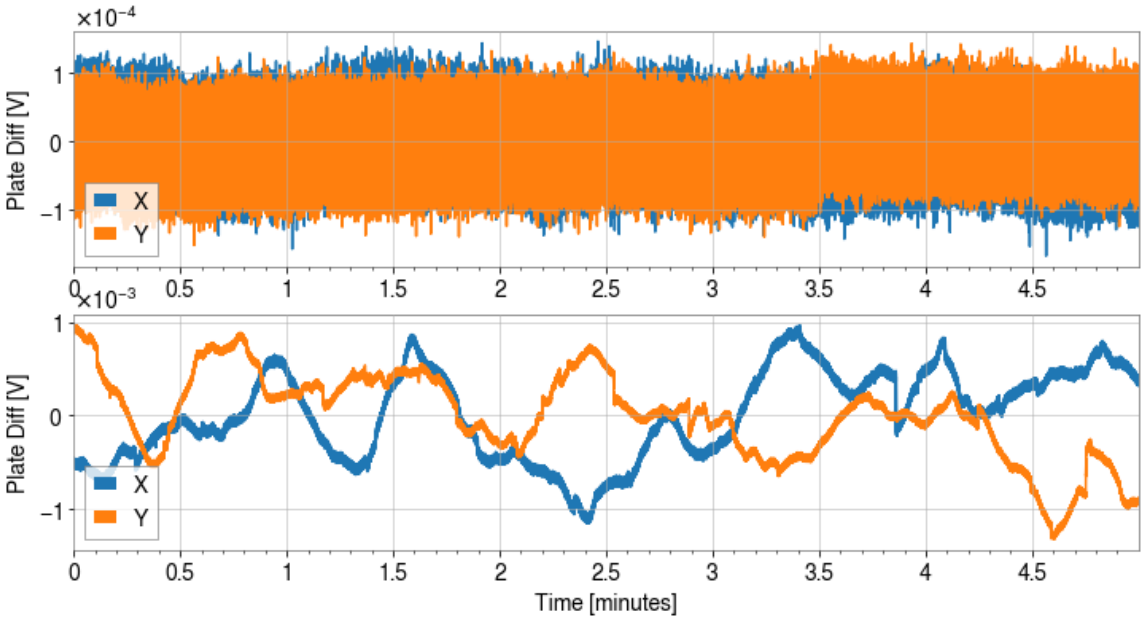

(a) Raw signals

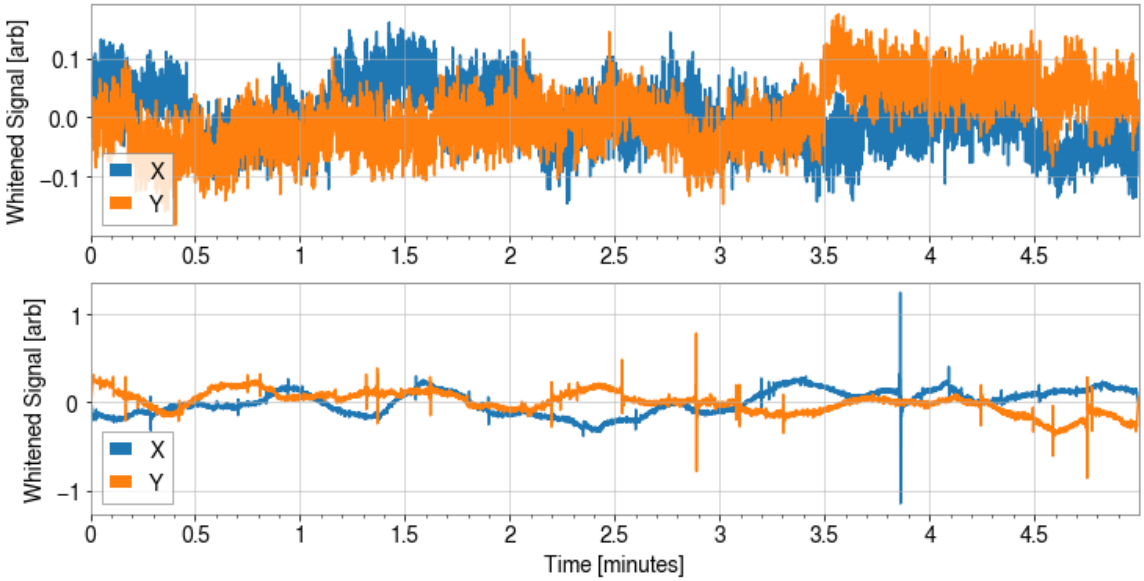

(b) Whitened signals

Figure 4-21: Ambient EFM signals in air. For comparison, the top sets of plots show the noise of the EFM with the sensing plates grounded to the body. There are some non-Gaussian transients in (a) the raw signal, but they become especially clear in (b) the whitened time series.

were observed in the EFM signals. Figure 4-22 shows these transients, which were often large enough to saturate the signal. The decay of the sensing RC circuit is clear after each impulse. They began to appear once the chamber reached approximately 0.5 atm and gradually tail off once the pressure reached a few torr. These transients were not observed when the chamber was vented to atmosphere. Additionally, when the chamber door was replaced but the chamber was not pumped down for a few days, these transients also became significantly smaller and less frequent.

Two possible origins were considered most likely, and they are explored in a test described in section 4.7. The first is the effect of particulates moving across the surface of the sensing plates and depositing charge via a triboelectric interaction. Because the amount of charge deposited increases with particle velocity [160], it makes sense that large transients are seen when the chamber is being evacuated and disappear when the pumping stops (figure 4-22). At atmospheric pressures, the gas density is sufficiently high that repeated collisions keep the dust charge low. At lower pressures, most of the dust will have settled to the chamber floor and there is insufficient gas pressure to bring it to the height of the EFM, which would explain the eventual disappearance. When venting to atmosphere, the gas flow is significantly lower at similar pressures than when pumping down. Finally, if the chamber is not pumped down until a few days after closing the door, the dust has time to settle to the chamber floor and is far less likely to get kicked up toward the EFM.

The second, which is believed to be the actual cause, is due to dielectric breakdown of the air in the chamber as the chamber pressure drops. Any net charge will produce an electric field, which, if sufficiently large, can ionize the air in the chamber. At the breakdown voltage, a stripped electron will accelerate due to this electric field and collide with another atom at high enough energy to ionize this atom, freeing another electron that can ionize additional atoms in a runaway process. This breakdown voltage varies according to Paschen's law. At high pressures, the mean free path of stripped electrons is small enough that they cannot accelerate to an ionizing speed. At low pressures, the gas density is so low that very few ionizing collisions occur. As the air pressure is reduced, residual charge on insulators near the EFM (most likely

the ceramic standoffs) induces dielectric breakdown of the air, which causes a small current to briefly flow, which is what is seen on the EFM. Since these insulators have already been discharged, these blips are not seen when venting to atmosphere. Waiting a few days before pumping down allows residual charges in the air to neutralize the charge on the insulators, reducing the incidence of these discharge events.

While these discharging events are not directly relevant to the interferometer sensitivity, the fact that they do appear suggests similar charging may occur on other insulators in the chamber, even when not directly manipulated during a chamber incursion. Specifically, it seems that even objects far away from chamber activity can become inadvertently charged during a vacuum incursion. To avoid having to pump out adsorbed water, the chamber is purged with dry air when open, which will further inhibit discharge at atmospheric pressure.

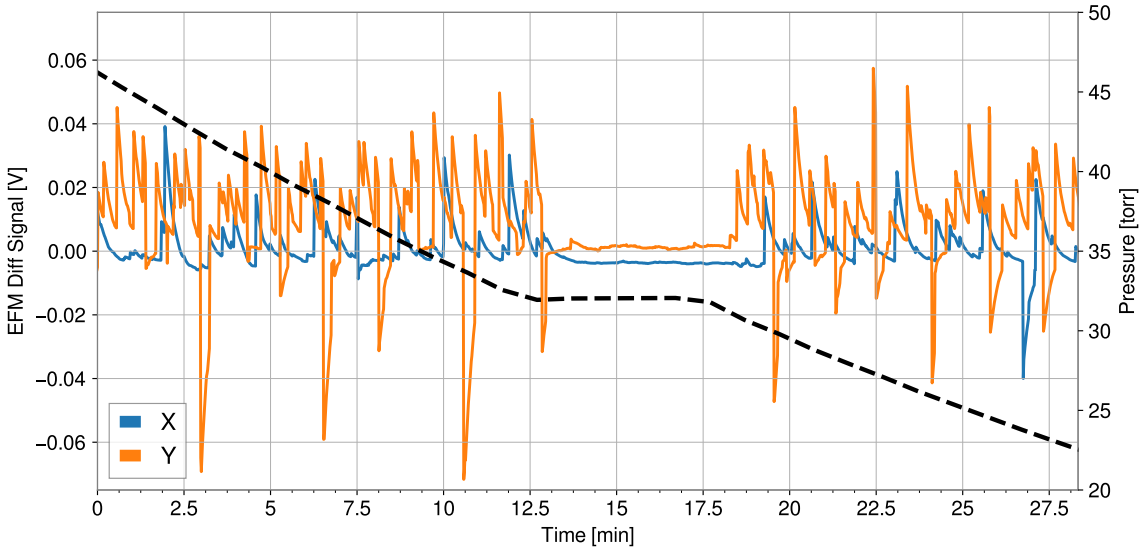

Figure 4-22: Glitches during the evacuation of the ETMY vacuum chamber. The roughing pump was briefly valved out to examine the effect on these transients.

4.7 Additional Charging Tests

These results all suggest that the discharging procedure described in section 4.3 is effective at reducing the net optic charge to values significantly below that which

might couple noise into DARM. For completeness, we review some additional tests that were performed.

- To confirm that net optic charge was low, the suspension cage was driven 50× above ambient motion and no increase in DARM was observed.

- Because the EFM had not seen any in-chamber electric fields in differential mode, we decided to run with a single plate grounded to see if the common-mode noise might be coupling to DARM. While various spectral lines appeared much larger in the EFM spectrum, there was no broadband increase in the spectrum, and no additional coherence with DARM.

- To understand the origin of the glitches during pump down, a test set up using a single EFM sensing plate and associated electronics was set up in a small vacuum chamber [161]. Figure 4-23 shows the set up. The following were explored as possible causes of the transients:

  – *Water evaporation*: It was hypothesized that some of these discharging events could be due to water evaporating off the sensing plate. The sensing plate was wiped down with a wet cloth prior to evacuation, with no glitches observed. None were observed when individual drops of water were deposited on the surface and the chamber was evacuated.

  – *Dust/gas charging*: Particulate matter was introduced into the chamber directly on and around the plate by rubbing my hands together above the chamber. No additional glitching was observed during pump down.

  – *Discharge near plates*: Introducing small pieces of plastic on the sensing plate (charged just by cutting from a larger structure) and pumping down induced discharge due to dielectric breakdown. Additionally, after a first pump down, subsequent evacuations did not produce these transients, presumably because the localized charges had already been discharged.

  Further investigations implicated the ceramic standoffs between the sensing plate and cube as a likely source of charge. The only way to induce these

transients in the test EFM was to slightly displace the sensing plate, in the process triboelectrically charging the standoffs. Of course, the geometry of this setup is very different from the main optic chamber, and additional tests would be required to confirm this. Fortunately, while it's clear that a measurable amount of charge is introduced when the chamber is open, tests with the full interferometer suggest that this all has a negligible effect on low-noise operation.

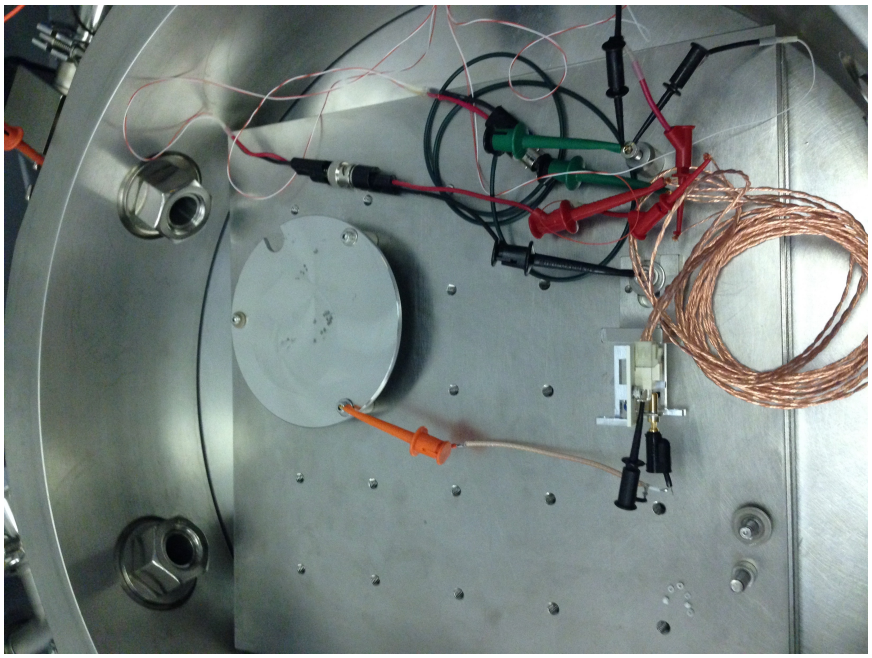

Figure 4-23: Tabletop setup for EFM testing. A spare calibration plate is used as a sensing plate and is connected to a set of EFM sensing electronics. The plate is isolated from the vacuum chamber by the same ceramic standoffs used for the real EFM.

4.8 Summary

In this chapter we have described work that:

- Shows the effect of large-scale stray electric fields on the test mass is negligible
- Fully characterized the electrostatic drive actuators, and estimated their contribution to the total noise budget

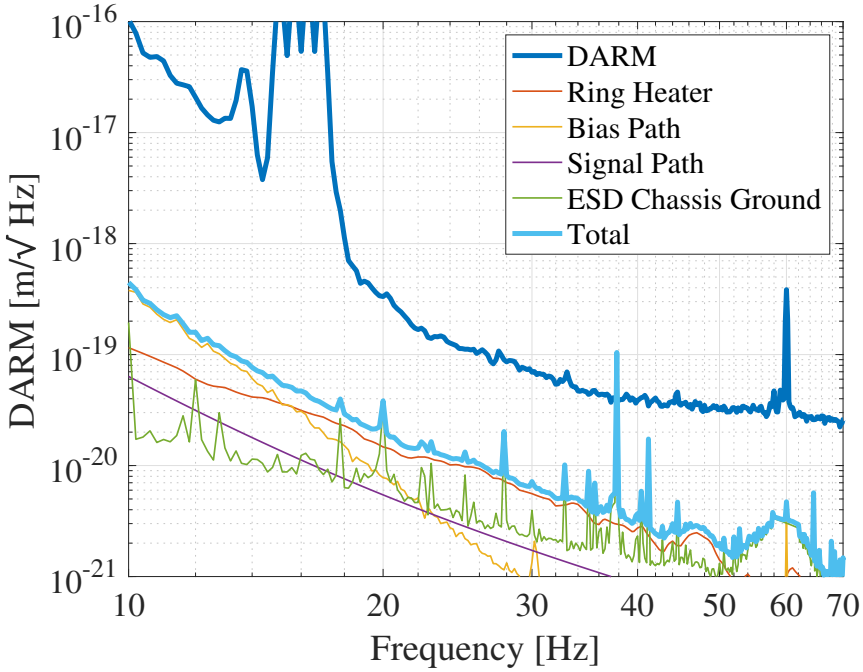

Figure 4-24: Summary of the known sources of stray electric field noise as compared with the current total interferometer sensitivity for an ETMY operating bias of 200 V.

- Eliminated a few sources of excess noise in certain actuator configurations, which will provide greater freedom for future observations
- Provided techniques to mitigate such noise in even more sensitive interferometers
- Investigated one possible origin of a type of transient ("glitch") in the gravitational-wave channel that negatively affects background rates

Figure 4-24 shows the summary of the above in terms of contribution to the interferometer noise spectrum. In the future, this noise may become non-negligible. We have presented a few techniques that can be used to decrease this coupling, such as setting the bias on unused optics to minimize the total noise and properly grounding the ESDs. There is ongoing work to develop conductive optical coatings [162], however, it remains challenging to produce coatings with satisfactory optical properties. Such coatings, if applied to the entire test mass, would eliminate charge separation on the optic surface and may mitigate the effect of net charge accumulation by spreading it on the test mass surface. Ultimately, we have determined that the effect of stray

electric fields and charging is not a limiting source for current gravitational-wave detectors.

Additional work is needed to understand the origin of glitches in the detectors. For various performance-related reasons, at LLO ETMY is used for DARM actuation, which is also the chamber that houses the EFM. Because the EFM is sensitive to actuation signals from the ESD, it will also see transient events that appear in DARM. Extended tests to look for transients in the EFM while instead using ETMX as an actuator may give hints to the origin of some transient events. While electrical discharge is not responsible for the vast majority of loud transients, the analysis outlined in section 4.5.1 provides one way to investigate the physical origin of other types of glitches.

Chapter 5

Fiber Amplifiers for Gravitational-Wave Detectors

High-power lasers are used for a wide variety of commercial and scientific applications including materials processing, manufacturing, lidar, and GW detectors [26, 14]. As laid out in chapter 2, the shot noise sensitivity of these detectors is improved with greater intracavity power. As such, it's important to begin with a stable, low-noise, high-power laser source. These detectors have particularly stringent requirements on the noise performance and stability of the primary laser [163]. The original design of the aLIGO laser uses a nonplanar ring oscillator (NPRO) master oscillator and free-space solid-state amplifiers to injection-lock a high-power oscillator [24, 164]. Together with frequency and amplitude feedback-stabilization systems and a filtering optical cavity, this system provides 165 W of pre-stabilized light for the interferometer.

Fiber amplifiers provide an attractive alternative to this arrangement [165]. Using fiber amplifiers in a master oscillator power amplifier (MOPA) configuration provides numerous advantages over the current solid-state injection-locked configuration. The guided structure makes the output beam size and divergence independent of operating power, enabling efficient mode-matching at all power levels and eliminating issues with thermal lensing that are present in the current laser system. Ytterbium-doped fiber amplifiers (YDFAs) have higher efficiency than Nd:YAG amplifiers because of the lower quantum defect and the high intensities that saturate the gain medium.

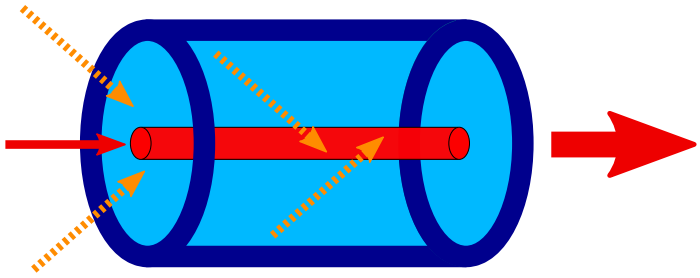

Figure 5-1: Double-clad active fiber, which consists of a doped core (red), inner cladding (light blue), and outer cladding (dark blue). The signal beam (red) is guided along the core, while pump light (dashed orange) is guided along the inner cladding to excite the lasing medium in the core. The output signal beam is amplified and retains excellent beam quality.

This reduces cooling requirements, allowing lower water flow rate, which can reduce vibrations that reduce GW detector sensitivity [166]. Figure 5-1 shows the principle of operation for these double-clad fiber amplifiers.

The removal of control electronics for injection-locking—or requiring more for coherent beam combining (CBC) [167]—is an additional advantage of this arrangement. Additionally, the development of high-power fiber amplifiers using erbium [168, 169, 170, 171], thulium [172, 173], or holmium [174, 175] is promising for use in third-generation gravitational-wave detectors [176], which may require lasers at 1.5 or 2 µm [177].

The improvement in efficiency and simplicity of optical fibers is not without downsides. Optical fibers can be significantly more susceptible to environmental coupling and non-linear effects, leading to phase shifts that can increase frequency noise or lead to beam pointing fluctuations [178]. To generate high powers, we would like to make the gain fiber long enough to amplify the incoming beam sufficiently. As we will see, nonlinear effects such as stimulated Brillouin scattering (SBS) limit the length of these fibers. Designing effective high-power fiber amplifiers requires an understanding of these different effects and engineering compromises to meet power and beam quality requirements.

In contrast to multi-kW continuous-wave (CW) YDFA systems, which use comparatively large linewidths to suppress instabilities caused by SBS [179, 180], interferometric GW detectors require single-mode, narrow-linewidth (kHz), polarized,

and stable laser sources. Narrow-linewidth YDFAs reaching >100 W have existed for some time [181], however, to achieve higher powers at narrow linewidths requires additional techniques to increase the SBS threshold while maintaining the requirements for GW detectors [182, 183], such as differential stress along the gain fiber [184], large temperature gradients [185], core pumping [186], or the use of photonic crystal fiber (PCF)/photonic-bandgap fiber (PBGF) [187, 188, 189].

In this chapter we present a YDFA MOPA that uses a specialty large-mode-area (LMA) fiber to increase the SBS threshold, and a passive output fiber to improve beam quality. The output power of the laser is as high as 178 W, and its noise performance is compatible with use in GW detectors. Section 5.4 discusses some of the additional requirements for reliable, long-term operation of these systems. This work was carried out in collaboration with MIT Lincoln Laboratory, where this fiber amplifier was constructed. Initial characterization was performed at Lincoln, with long-term measurements and noise characterization performed at the MIT LIGO Lab. This chapter is largely based on the results published in [190].

5.1 Stimulated Brillouin Scattering

High-power narrow-linewidth fiber amplifiers are in practice most limited by stimulated Brillouin scattering (SBS). A thorough discussion of the mechanism and observations of SBS is provided in [191]. We review the most relevant points here.

For high optical intensity in fiber, the incident light can induce phonons in the fiber, usually via electrostriction or less commonly, by thermal effects. The light will then scatter off these *stimulated* phonons. This process is coherent and is enhanced in longer fiber. To gain some intuition, we provide the expressions relevant for passive fiber (which also is relevant for the fiber ring experiment described in section 2.3.4). A derivation of these expressions can be found in [192]. The effect in pumped fibers is more complicated and requires more complex models.

The SBS gain factor can be approximated as a Lorentzian given by

$$g_{\mathrm{B}} = g_0 \frac{(\Gamma_B/2)^2}{(\Omega_B - \Omega)^2 + (\Gamma_{\mathrm{B}}/2)^2}, \tag{5.1}$$

with a center gain factor g_0 that depends on the material properties, $\Omega \equiv \omega_1 - \omega_2$, $\omega_{1(2)}$ is the frequency of the incident (scattered) beam, $\Omega_B = 2v\omega_1/(c/n)$ is the Brillouin frequency, v is the speed of sound in the material, and Γ_{B} is the Brillouin linewidth given by the inverse of the phonon lifetime. Table 5.1 provides values of these constants for silica fiber.

$\Omega_{\mathrm{B}}/2\pi$ [GHz]	$\Gamma_{\mathrm{B}}/2\pi$ [MHz]	g_0 [m/GW]
25.8	78	0.045

Table 5.1: SBS properties of bulk SiO_2 at 694 nm. These numbers can vary based on dopant concentration, wavelength, and other parameters. From [192].

As a result, SBS becomes particularly problematic for light sources with linewidths significantly smaller than the Brillouin linewidth. A common way to reduce the effect of SBS is to increase the source linewidth, but we require narrow-linewidths for interferometry.

Let $I_{1(2)}$ be the intensity of the incident (scattered) beam. Considering the interaction between the incident light, scattered light, and the induced phonons, the steady-state solution for constant incident power I_1 as a function of distance z is given by

$$I_2(z) = I_2(L) \exp\left[gI_1(L - z)\right]. \tag{5.2}$$

Hence, the scattered beam increases exponentially in the backward-propagating direction. This simple model assumes the incident beam is spatially invariant, however, for large enough intensities of the scattered beam one must consider the depletion of the incident beam. When allowing that intensity to vary, one finds that nearly all of the incident beam can be scattered into the Stokes wave. Additional complexity comes into play for pumped active fiber, but the basic concept remains: for a given length of uniform optical fiber, there is a maximum attainable intensity before SBS

dominates.

The approximate SBS threshold power for a narrow-linewidth source through passive fiber is given by [193, 194]

$$P_{\text{th}} \approx \frac{21 K A_{\text{eff}}}{g_0 L_{\text{eff}}}, \tag{5.3}$$

where K depends on polarization and is 1 for linearly polarized light, and A_{eff} is the effective mode area. For a given loss per unit length α, the effective length L_{eff} is given by

$$L_{\text{eff}} = \frac{1 - \exp(-\alpha L)}{\alpha}. \tag{5.4}$$

Fortunately, in active, pumped fibers, the changing temperature profile along the fiber provides a higher SBS threshold power.

SBS negatively affects the operation of the fiber amplifier in a few ways. First, because this process will scatter light and is highly nonlinear, this effectively sets a maximum power limit for a given fiber length and mode area. Second, because SBS beams are directed in the reverse direction, this process can reflect a damaging amount of power back toward low-power components. Lastly, this effect introduces broadband excess intensity noise. A number of explanations have been provided for the origin of this noise, one of which is due to the nonlinear phase shift due to SBS gain. This phase-modulated light interferes with the main beam to produce excess intensity noise. While these intensity fluctuations appear at MHz frequencies, far above the GW frequency band, they can be down-converted to this band due to modulation at these frequencies that are used to generate control signals (section 2.3.3).

Equation (5.3) provides guidance on how to limit the effect of SBS: the fiber length or the center SBS gain g_0 must be reduced or the effective mode area must be increased. The latter can be achieved by using optical fiber with larger mode field diameter (MFD), at the expensive of additional higher-order mode (HOM) beam propagation and degraded beam quality. The SBS gain can be reduced by using incident light with linewidths larger than the Brillouin linewidth Γ_B, though this is not a feasible option for gravitational-wave detectors, which require narrow linewidth

sources. The Brillouin frequency Ω_B can be made position-dependent by changing the material speed of sound v by applying varying stress loads or different temperatures along the fiber. The required length of fiber can be reduced by improving the pump absorption through core pumping. In principle the active dopant concentration can be increased, however at high doping fractions photodarkening and other negative effects become a problem.

The amplifier discussed in the following section provides one solution to the SBS problem by using a specialty LMA fiber designed to minimize the effect of HOMs.

5.2 Amplifier Design

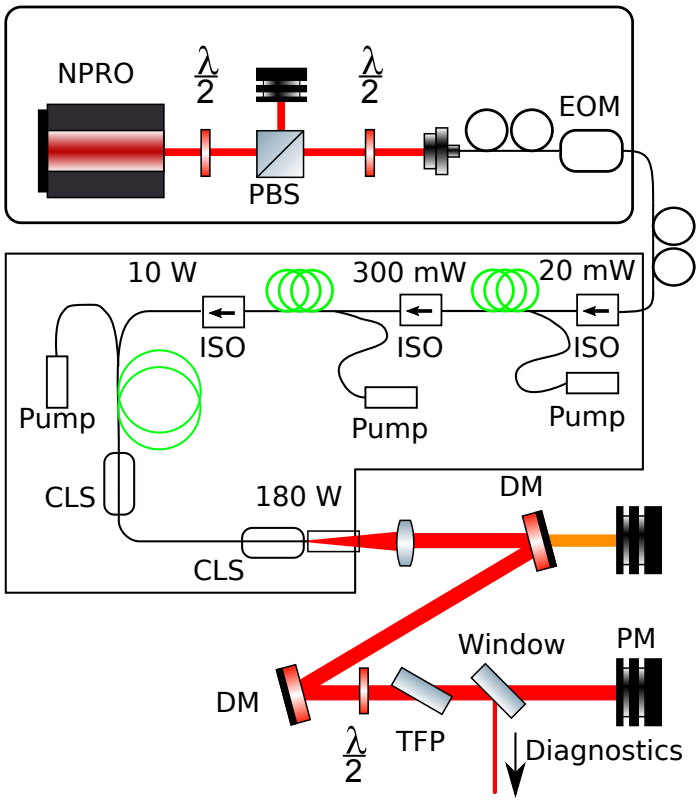

Figure 5-2: Amplifier layout. EOM: in-fiber electro-optic modulator, ISO: fiber isolator, CLS: cladding light stripper, DM: dichroic mirror, TFP: thin-film polarizer, PBS: polarizing beam splitter, PM: power meter.

Figure 5-2 shows the layout of the amplifier. The system is a MOPA-style fiber amplifier seeded by a JDSU 126 Nd:YAG NPRO emitting at 1064 nm (linewidth

<5 kHz). This is fiber-coupled and passed through a fiber phase modulator to generate sidebands for a Pound-Drever-Hall lock to an optical cavity [43, 44] for diagnostic purposes.

The amplifier uses three stages of Yb-doped, polarization-maintaining (PM) fiber, each forward pumped at 976 nm. It amplifies a seed power of 20 mW up to 178 W. To maximize the SBS threshold, the last stage uses a short length (1.2 m) of 35 µm core fiber (Yb MCOF-35/250-07-2.5-PM) from INO [195], which has been designed with a depressed cladding profile and confined-core doping to suppress higher-order-modes (HOMs). To further improve the output beam quality, the output of this fiber is spliced directly to 19 cm of passive PM fiber (Nufern PLMA-GDF-20/400-M); this splice incorporates an etched pump light stripper. The end of the passive fiber includes a high-index epoxy cladding light stripper before a fused silica endcap (2 mm $\varnothing \times$ 10 mm) to allow the beam to expand and reduce the intensity on the end face. A free-space collimating lens, dichroic mirrors, and an output polarizer provide a clean, collimated output beam.

The amplifier electronics and software incorporate interlocks that monitor loss of seed, low output power, excess SBS backscatter, and temperatures of individual stages. The output power is limited to 178 W by high fiber temperatures rather than SBS (-17 dB from Rayleigh backscatter peak), suggesting that better thermal management would enable operation at higher power.

5.3 Results

The performance requirements of the aLIGO input laser beam [163] are achieved with a suite of stabilization systems. This PSL is then injected into the interferometer vacuum envelope. The actuation scheme [24] acts only on the seed laser or the output beam and can thus be used in this type of system with minimal modification. As such, our demonstration that in most respects the free-running performance of this system matches or improves on the current setup will mean that the PSL requirements can easily be attained.

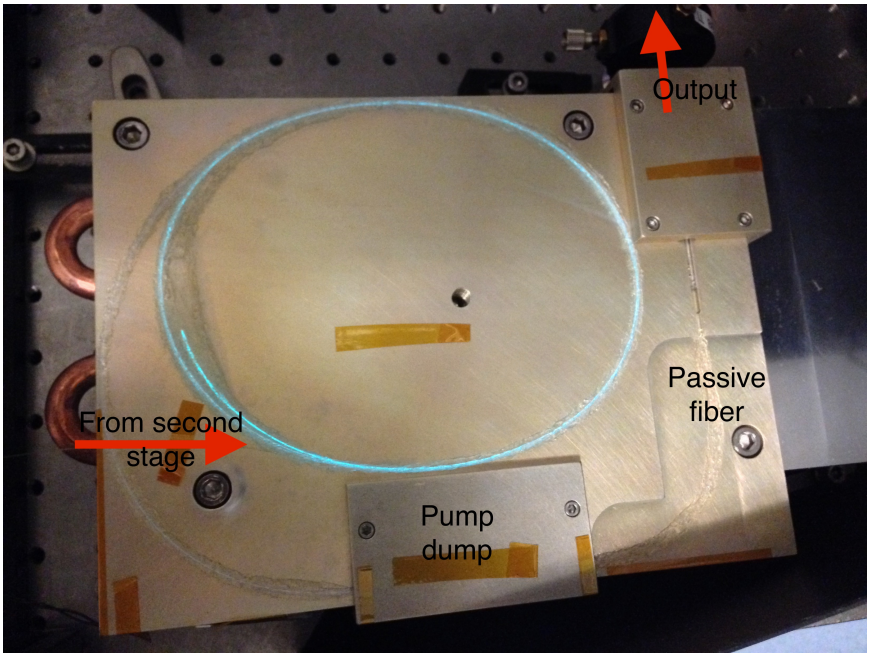

Figure 5-3: Third stage of fiber amplifier. The end cap is hidden below the output cover. The fluorescence of the gain fiber can be seen.

5.3.1 Power Stability

Gravitational wave detectors require high circulating optical powers to reduce the contribution from shot noise. Current designs require 165 W in the TEM_{00} mode [163], with stringent output power noise requirements to avoid radiation pressure effects and sensing noise. Maximum attained output power with this amplifier was 178 W at 1064 nm with 73% optical-to-optical efficiency. The relative power noise is shown in Figure 5-4. Above 100 kHz, the free-running relative power noise is below 10^{-7} $Hz^{-1/2}$, limited by sensor electronic noise. In the saturated regime, pump intensity noise will couple to the output [198], and indeed, significant coherence was observed above 10 Hz between the pump diode current and the output power (figure 5-5). The pump diode drivers used here were not designed to have particularly low noise, so improving the noise performance of the current driver is expected to improve the free-running performance. Alternatively, a power stabilization loop utilizing pump diode current feedback could be used to improve performance [94, 187].

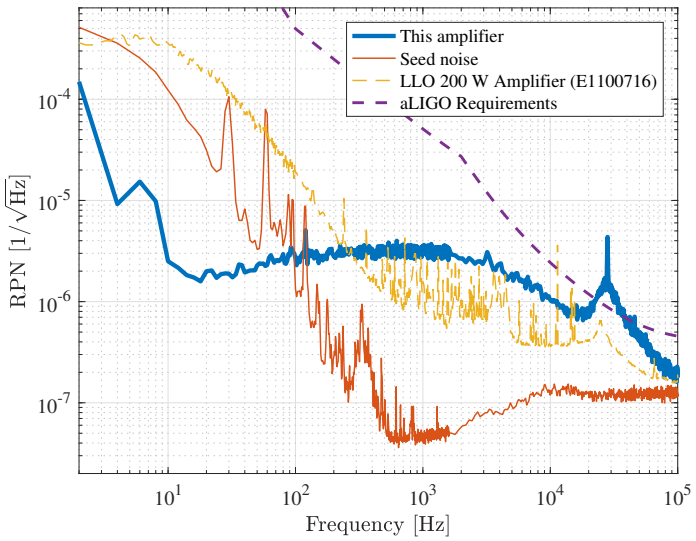

Figure 5-4: Output relative power noise (RPN). The peak at 28 kHz is from excess power noise in the second stage. Also included is the power noise of the LIGO Livingston observatory (LLO) high-power oscillator, from [196]. Free-running requirements are calculated from current power noise stabilization loop gains [197].

5.3.2 Frequency Noise

Previous measurements of fiber amplifier phase noise [199, 200] have observed negligible contribution as compared to a free-running NPRO laser. In contrast to those differential measurements, the total frequency noise of this laser was measured by coupling a beam sample into an acoustically isolated 20-m fiber ring resonator and locked using polarization spectroscopy as laid out in section 2.3.4 and [47]. Figure 5-6 confirms that the amplifier does not add to the free-running NPRO seed frequency noise, which is the same as the free-running noise of the aLIGO PSL.

5.3.3 Beam Quality

The output beam must mostly be a TEM_{00} mode, as this is the light that can be coupled into the interferometer and determines the ultimate shot-noise sensitivity. Figure 5-8 shows a mode scan of the output beam produced by sweeping the length of a triangular ring cavity [95]. In contrast to a simple M^2 measurement or a Gaussian profile fit, which does not provide a reliable measure of the fiber modal content [178], this method provides a lower limit of the total power in the TEM_{00} mode [201] of

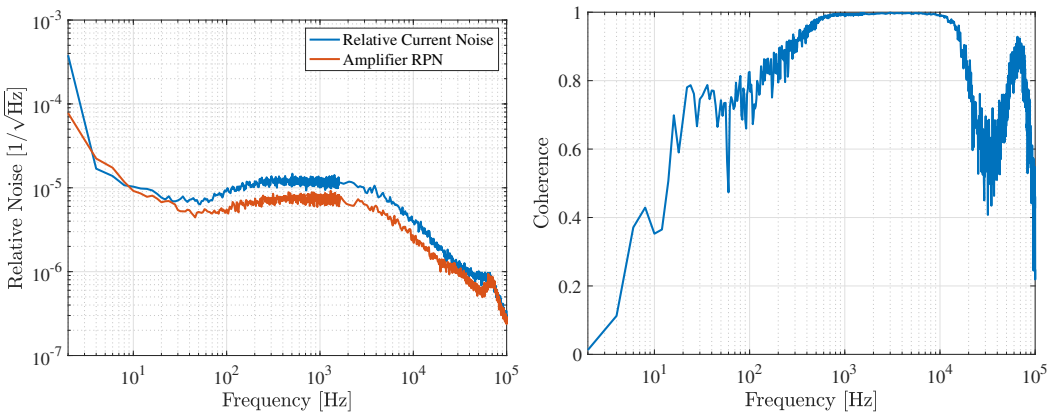

(a) (b)

Figure 5-5: (a) Relative current and power noise of the third-stage pump diodes and (b) magnitude squared coherence between current noise and output power fluctuations.

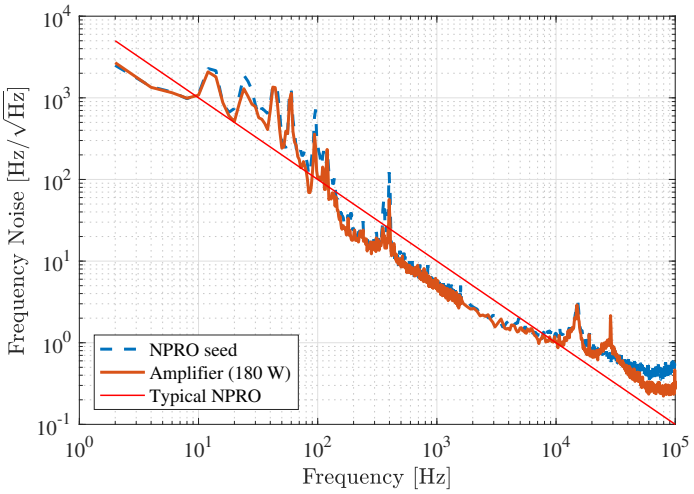

Figure 5-6: Frequency noise of the full amplifier compared to that of the seed, indicating the amplifier adds minimal phase noise. The peaks around 350 Hz are due to table vibrations from water flow, and the peak at 28 kHz is due to excess amplifier power noise. Below 100 Hz various mechanical resonances of the NPRO setup can be seen.

90.1%. The output beam was slightly astigmatic, so the use of cylindrical lenses may be able to improve this fraction. For completeness, we report M^2 values of $M_x^2 = 1.32$ and $M_y^2 = 1.24$ (figure 5-7). A polarization extinction ratio of 19 dB is obtained at full power.

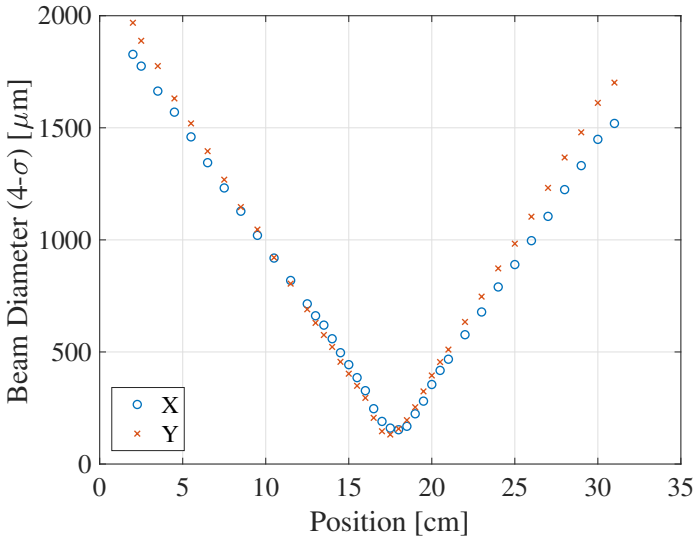

Figure 5-7: M^2 measurement of beam at full power.

5.3.4 Beam Pointing

In contrast to unguided amplifier systems, HOMs in optical fibers have a different propagation constant from the lowest-order LP_{01} mode. As a result, any change in refractive index, for example due to external stress or environmental drifts, will induce a different phase shift for the fundamental as compared to HOMs. This can lead to beam pointing shifts and will be more problematic in fibers such as these designed with large propagation constant splitting between the fundamental and HOMs [195]. Figure 5-9 shows that in spite of this effect, the intrinsic beam pointing noise is still within the aLIGO requirements. These data were taken by putting the beam directly on a quadrant photodetector (QPD), measuring first the translational noise in the collimated beam, followed by a measurement at the waist of a focused beam to measure the tilt component. In aLIGO, two optical cavities passively suppress this free-running jitter by a factor of $> 30,000$ before reaching the interferometer [24].

5.3.5 Long-term Stability

Long-term environmental drifts are observed to correlate with output laser power. This can be explained by environmentally-dependent phase shifts between the LP_{01} and LP_{11} modes in the INO gain fiber (having 35 µm core) changing the modal cou-

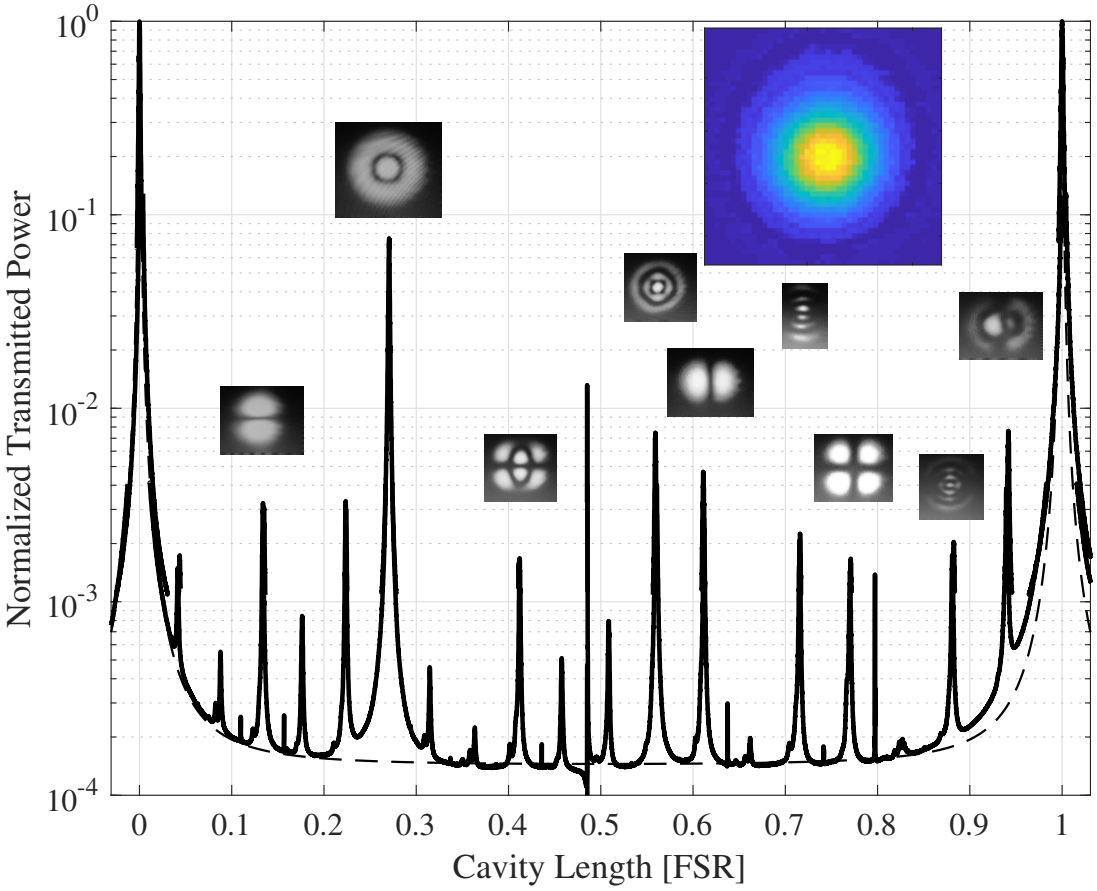

Figure 5-8: Transmission of an optical cavity as the length is swept through a full free spectral range (FSR). Included is the finite width of the TEM_{00} mode and a DC offset. Narrow peaks are due to leaked *s*-polarized light, which has a higher finesse and is split due to the intrinsic cavity birefringence. The baseline DC power is suspected to be due to scattered light and not amplified spontaneous emission (ASE), as it scales proportionally to total power. The high content of the Laguerre-Gaussian 10 mode is due to a beam astigmatism that could not be removed with spherical optics. Inset: near-field image of output beam.

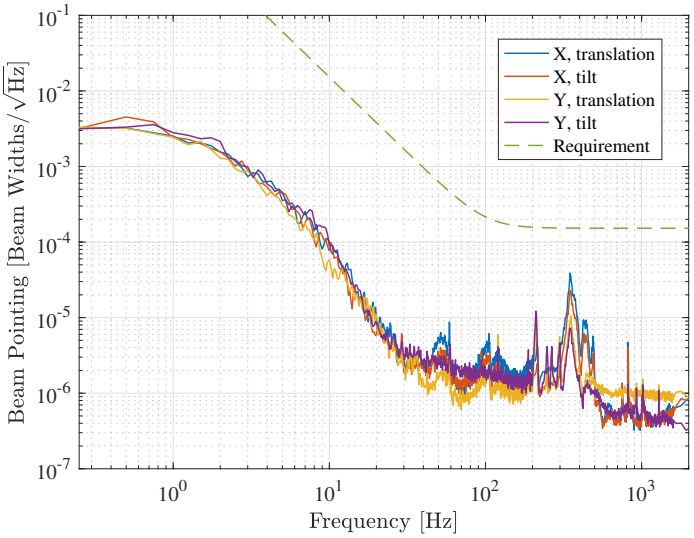

Figure 5-9: Beam pointing (jitter) noise as a fraction of beam waist. Below 20 Hz the noise is dominated by air currents, which can be reduced by using fan-less power supplies. Peaks near 350 Hz are due to table motion from water flow. The aLIGO requirements have been scaled to account for the passive jitter suppression from the finite finesse of the optical cavity.

pling into the output passive fiber (having 20 µm core) and producing large output power fluctuations $> 10\%$ over hour timescales because the smaller-core fiber filters out much of the HOM content. Despite the power variations, the output mode quality stays fairly constant, as can be seen by measuring the long-term transmission of a cavity locked to the TEM_{00} mode (Figure 5-10). This suggests the output passive fiber is largely insensitive to environmental fluctuations due to its larger 400 µm cladding. Tapering the splice from gain fiber to passive fiber will improve the coupling of LP_{01} modes between these fibers and should eliminate some of this environmental sensitivity. Additional coils on the output passive fiber can improve output beam quality; because this system is not SBS limited, this will not limit the output power.

5.4 Amplifier Damage

Unfortunately, many of the advantages of optical fiber systems also introduce major problems. The high energy densities require efficient cooling, without which the fiber can be thermally damaged. In our case, ensuring that such cooling systems do not

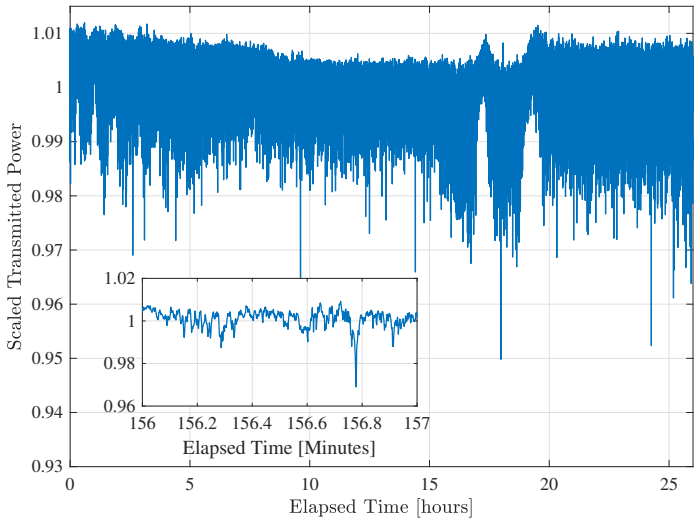

Figure 5-10: Long-term cavity transmission indicating relatively stable output mode quality. The laser is locked to the TEM_{00} mode of the optical cavity used in Figure 5-8. These have been scaled by total output power to decouple total power drifts from beam pointing and/or size fluctuations and normalized to the mean transmitted power. The bumps around 17 and 19 hrs are due to lab environmental fluctuations, which can change the cavity coupling via in-fiber relative phase-shifts or free-space alignment. Inset: Zoomed segment showing short-term fluctuations.

add additional phase or jitter noise is another challenge. The waveguide property of optical fibers simplifies many aspects, but it also makes these systems susceptible to damage due to amplification of undesired signals. Especially in CW fiber amplifiers, the production of unintended optical pulses can have damaging effects. For example, backward propagating (non-stimulated) Brillouin scattering off acoustic waves due to external vibrations can be powerful enough to damage the amplifier. Such impacts can be largely avoided via proper operator training, vibrational isolation, and appropriate electronic controls. Optical isolation between gain stages can eliminate some risk from backward-propagating beams, however, as we will show, these pulses can reach powers that can damage the gain fiber itself.

One possible issue for these systems is the inadvertent loss of the seed laser power, which can produce large pulses in the amplifier if the pumps do not shut off. Self-pulsing in unseeded fiber amplifiers (i.e., not an oscillator) is not a well-studied phenomenon [202], especially in these high-power applications where peak powers can easily surpass damage thresholds. A damaged laser can lead to long downtime for

the observatories, so reliable operation is paramount. Unfortunately, the fast nature of such phenomena, as well as the large currents and voltages at pump diodes and the need for stable current flow in normal operating conditions, make the design of appropriate interlocks challenging.

Understandably, such damage threshold tests are risky and potentially costly to undertake and the author is unaware of a systematic study of these damage thresholds *in situ*, in spite of the anecdotally common occurrence of this type of damage. Nevertheless, we can estimate the required speed for these electronic interlocks to shut off or shunt current from the pump diodes. Bulk measurements of the damage threshold of silica at 1064 nm have been carried out [203] and are used to estimate the damage threshold in fiber.

5.4.1 Fiber Damage Threshold

With the inadvertent loss of the seed laser, for a short time the gain fiber is still pumped. The excited gain-medium atoms are no longer stimulated downward and they remain in a high-energy state until a pulse seeded by spontaneous emission causes the atoms to relax. The resulting pulse has a power many times the CW power of the amplifier and can catastrophically damage the amplifier. What follows is a short discussion of damage mechanisms and rough (and pessimistic) numbers for required safety pump diode electronic shut off times.

For CW laser sources, the damage is often in the form of a fiber fuse [204], in which the fiber is locally heated to a temperature at which the absorption increases exponentially, and the fiber is then quickly heated to damaging temperatures (up to nearly 10,000 K). For our setup, in which a Q-switched pulse produces the damage, the most likely damage mechanism is electron avalanche due to the laser light electric field, which produces a plasma with high absorption, which then heats the fiber catastrophically [203]. For shorter pulses and/or shorter wavelengths photoionization becomes more important [203].

When the seed power is lost, the pump diodes are driven for a short time $\Delta\tau$ before the safety electronics engage and cut off power to the pumps. For this estimate, we

make the following assumptions:

1. $\Delta\tau$ is short enough that the gain fiber is not saturated.

2. All of the excited atoms amplify the pulse (very pessimistic).

This will then produce a pulse of width Δt with average power

$$P_{\text{pulse}} = \frac{P_{\text{pump}}\Delta\tau}{\Delta t}. \tag{5.5}$$

The small MFD of the fiber will produce very high intensities that can reach the silica damage threshold of I_d =4.75 kW/µm^2 [203]. The peak intensity for a Gaussian beam of diameter $d = 2w$ and power P is

$$I_{\text{peak}} = \frac{2P}{\pi w^2} = \frac{8P}{\pi d^2}, \tag{5.6}$$

which then requires

$$I_d > I_{\text{peak}} = \frac{8P_{\text{peak}}}{\pi d^2} = \frac{8}{\pi d^2}\frac{P_{\text{pump}}\Delta\tau}{\Delta t}. \tag{5.7}$$

For an amplifier similar to the third stage of the above amplifier, this requires a pump shut off time of

$$\Delta\tau < \frac{\pi}{8}\frac{d^2\Delta t}{P_{\text{pump}}}I_d = 40\,\mu\text{s}\left(\frac{\Delta t}{10\text{ ns}}\right)\left(\frac{300\text{ W}}{P_{\text{pump}}}\right)\left(\frac{d}{25\,\mu\text{m}}\right)^2\left(\frac{I_d}{4.75\,\text{kW}/\mu\text{m}^2}\right). \tag{5.8}$$

Even though the second stage has 20× lower output power, the smaller core size requires us to consider the damage threshold here as well. The second stage pump must turn off in time

$$\Delta\tau < 120\ \mu\text{s}\left(\frac{\Delta t}{10\text{ ns}}\right)\left(\frac{15\text{ W}}{P_{\text{pump}}}\right)\left(\frac{d}{10\ \mu\text{m}}\right)^2\left(\frac{I_d}{4.75\text{ kW}/\mu\text{m}^2}\right). \tag{5.9}$$

5.4.2 Self-Focusing Limit

Even if the pulse does not reach the damage threshold at the nominal fiber MFD, it may reach a power at which it will begin to self-focus due to either thermal lensing or the nonlinear index of refraction (optical Kerr effect) [192] and eventually reach the damage threshold. For completeness, we consider the required time to turn off the pump diodes due to the self-focusing limit even though it ends up being less stringent than the above limit.

Surprisingly, this self focusing limit depends only on the total power in the beam, not intensity[1]. The critical power for self-focusing is ≈4 MW in the 1 µm region for silica fiber. We require that the power of the pulse be smaller than the critical power:

$$P_{\text{pulse}} = \frac{P_{\text{pump}}\Delta\tau}{\Delta t} < P_{\text{crit}}, \tag{5.10}$$

which leads to

$$\Delta\tau < \frac{\Delta t P_{\text{crit}}}{P_{\text{pump}}} = 130\ \mu\text{s}\ \left(\frac{\Delta t}{10\ \text{ns}}\right)\left(\frac{P_{\text{crit}}}{4\ \text{MW}}\right)\left(\frac{300\ \text{W}}{P_{\text{pump}}}\right). \tag{5.11}$$

The self focus distance [192] is

$$z_{\text{sf}} = \frac{2n_0 w_0^2}{\lambda_0}\frac{1}{\sqrt{P/P_{\text{crit}} - 1}} \approx 2\ \text{mm} \tag{5.12}$$

at 5% above P_{crit} and a MFD of 24 µm. So even for powers just above the critical power, it is unlikely the pulse will be able to leave the fiber before causing catastrophic damage. Note that this neglects focusing or damage from thermal effects, which are possibly comparable or more important than Kerr self-focusing.

[1] Physically this is because while Δn depends on intensity (and thus will have a larger effect for smaller beams), the divergence of larger beams is smaller and requires a smaller Δn to counteract diffraction.

5.4.3 Discussion

Even with these rough numbers, a good safety margin on the pump electronics would require shutdown times around 10 µs. However, there are many unknowns with these numbers, and literature here is lacking. First, there are large uncertainties on the pulse width. The number used here is comparable to pulse widths achieved with Q-switched lasers, however, this is in an amplifier and not an oscillator. It's impossible to say which atom decays spontaneously to seed this pulse, which might give a pulse width of order $L/c \approx 3\,\mathrm{ns}$ for $L = 1\,\mathrm{m}$ section of gain fiber. Additional studies of self-pulsing in unseeded amplifiers are required to better estimate this number.

Other unknowns affect this estimate. While numerical simulations seem to suggest that P_{crit} is nearly the same when waveguide effects are considered (i.e., in fiber) as in bulk material [205, 206], the question is still up for debate [207]. Depending on the pulse shape, the peak power is of the form $P_p = \alpha E_0/\Delta t$, where $\alpha \lessapprox 1$ for most common pulse shapes. Finally, if the pulse is emitted in a HOM of the fiber, the peak intensity can be different and often will be higher than for a Gaussian beam.

As we have seen, excess noise on the pump diode current drivers can produce excess output intensity noise. Inductors could be added along this path to stabilize the current, however, this would prevent rapid current shutoff without some sort of current shunt. Additional work is required to develop electronics that are both low-noise and can be shut off quickly to prevent damage to the amplifier.

5.5 Conclusion

We present an amplifier with noise performance in line with free-running requirements for Advanced LIGO. Improved thermal management can allow for higher powers using this configuration and additional shielding of the output fiber to environmental fluctuations will ensure stable long-term performance. We demonstrate that this design can be used to produce a low-noise, diffraction-limited output beam, with noise performance consistent with current aLIGO requirements. The use of such a system can provide a high-power coherent light source for current and future gravitational

wave detectors. Additional study is required to understand the limitations to long-term reliable operation, though we have laid out some of the approximate numbers for these types of studies.

Chapter 6

Conclusion and Outlook

We have presented a few experiments aimed at reaching and understanding the radiation-pressure-dominated regime of gravitational-wave detectors and other interferometric systems. This work will lead the way to improved sensitivity of such interferometers as the circulating power is increased.

We described a small-scale Fabry-Pérot Michelson interferometer with gram-scale mechanical oscillators designed to reach the quantum-limited measurement regime in the audio frequency band. While we ultimately were not able to reach the QRPN-dominated regime, we have validated the thermal noise model previously published [81] and improved on it by nearly a factor of 5. Such design considerations will be helpful in other optomechanical experiments, where thermal noise is often a limiting noise source.

We investigated the effect of test mass charging and stray electric fields as impediments to reaching the QRPN-dominated regime of current aLIGO detectors, ultimately finding that these effects contribute negligibly to the current noise spectrum. This is the first comprehensive noise budget for these contributions. In the course of these investigations, we have also discovered ways to address these noise contributions in the future as the detectors become more sensitive. Additional work is required to understand the effect of discharging transients in the gravitational-wave channel, though we have proposed one way to investigate the origin of other large transients in the detector. Additionally, this investigation will help inform the types

of actuators to use in future gravitational-wave detectors.

Finally, we have characterized a narrow-linewidth high-power ytterbium-doped fiber amplifier (YDFA) that reaches the noise requirements for future use in gravitational-wave detectors. Additional work is required to improve the long-term reliability of these systems, but development of high-power laser systems, both fiber amplifiers and free-space amplifiers [208], continues for increased power in O4 and beyond. Research into high-power fiber amplifiers at other wavelengths is already underway, making them obvious candidates as sources for next-generation detectors.

Work continues on increasing the LIGO circulating power. Currently the detectors operate with approximately 200 kW circulating power in each arm, though challenges remain to increasing this [32]. In addition to being limited by the input laser power, point absorbers on the surface of the test masses produce thermal distortions that become larger at higher powers and limit the effective maximum circulating power. Solutions to these issues must be found before increasing power can be injected.

Even with increased circulating power, other improvements are required to reach the SQL and below[1]. First, the mirror coating thermal noise must be reduced. This can only be solved with an aggressive research campaign for low-loss dielectric coatings. Increased squeezing levels will cause QRPN to limit low-frequency sensitivity unless a combination of frequency-dependent squeezing and larger test masses are developed. Many of these improvements are already planned for aLIGO Plus (A+) [177, 209], while others are planned for third-generation gravitational-wave detectors like Cosmic Explorer (CE) [210, 211] and the Einstein Telescope (ET) [212]. Other interferometer geometries like speedmeters [213, 214], which, because they measure momentum instead of position, are immune to QRPN and offer an alternative to the current detector geometries. Instruments designed specifically for high-frequency sensitivity will require even higher circulating powers [215].

Increasing the sensitivity in the shot-noise-limited frequency range is particularly useful for studying the tidal effects in merging neutron stars, helping to understand

[1] In principle it will be possible to make a narrow-band sub-SQL measurement using squeezed light injected at an appropriate angle, though this configuration is not ideal for GW astronomy [79].

the behavior of nuclear-rich material, including the arrangement of atomic nuclei [216]. These interferometric GW detectors, which require large CW laser powers to measure sub-attometer displacements in kilogram-scale masses, allow us to study some of the densest objects in the universe. As the sensitivity of these detectors increases, so too will our understanding of the universe.

Appendix A

Derivations

A.1 Amplitude and Phase Derivative

Consider complex number $z = re^{i\varphi}$, with $r, \varphi \in \mathbb{R}$, that depends on some parameter q. The full derivative of z is given by

$$\begin{aligned}\frac{dz}{dq} &= \frac{\partial r}{\partial q}e^{i\varphi} + ire^{i\varphi}\frac{\partial \varphi}{\partial q} \\ &= \left(\frac{\partial r}{\partial q}\frac{1}{r} + i\frac{\partial \varphi}{\partial q}\right) z. \end{aligned} \tag{A.1}$$

In this form, it's clear this complex number is a sum of a radial component parallel to the original vector in the complex plane, which represents the changing amplitude, and a perpendicular term representing the phase change. We can then rewrite these components in terms of the full derivative z':

$$z'_r = \Re\left\{z'\frac{r}{z}\right\} = \Re\left\{\frac{z'|z|}{z}\right\} \tag{A.2}$$

$$z'_\varphi = \Im\left\{\frac{z'}{z}\right\} \tag{A.3}$$

A.2 Phase and Amplitude Modulation

As we have seen in section 2.3.3, we can write a small phase modulation in terms of sidebands at the modulation frequency. More generally, we can use the Jacobi-Anger expansion to write a phase-modulated light beam as

$$E_0 e^{i(\omega t + \beta \sin \Omega t)} = E_0 e^{i\omega t} \sum_{n=-\infty}^{\infty} J_n(\beta) e^{in\Omega t}, \tag{A.4}$$

where J_n are the Bessel functions of the first kind. Thus, in general there are an infinite number of frequency components. In most cases, this modulation will be small ($|\beta| \ll 1$), so we can expand the above in orders of β. In general, for non-negative integer n, we can write the Bessel functions as

$$J_n(\beta) = \sum_{m=0}^{\infty} \frac{(-1)^m}{m!\,(m+n)!} \left(\frac{\beta}{2}\right)^{2m+n} = \frac{1}{n!} \left(\frac{\beta}{2}\right)^n + \mathcal{O}\left(\beta^{n+2}\right). \tag{A.5}$$

For $n \in \mathbb{Z}$, we can use the relation

$$J_{-n}(\beta) = (-1)^n J_n(\beta) \tag{A.6}$$

to write equation (A.4) as

$$\begin{aligned} E_0 e^{i(\omega t + \beta \sin \Omega t)} &\approx E_0 e^{i\omega t} \left[1 - \frac{\beta^2}{4} + \frac{\beta}{2}(e^{i\Omega t} - e^{-i\Omega t})\right] \\ &= E_0 e^{i\omega t} \left[1 - \frac{\beta^2}{4} + i\beta \sin \Omega t\right], \end{aligned} \tag{A.7}$$

where we have expanded $J_0(\beta)$ to order $m = 1$ to show that the carrier power is transferred into the sidebands. Note that at $t = 0$ one sideband is in phase with the carrier while the other is 180° out of phase. When the sidebands have the same phase, they are 90° out of phase with the carrier. At no time are the sidebands and carrier all in phase; if this were so, we should observe a changing amplitude. The sidebands interfere to produce a fluctuating imaginary part of the amplitude. This is exactly in accordance with what we derived in section A.1 to be the phase component.

Now let's consider an amplitude modulated field. We will write the amplitude in a form that is more useful for devices that can only reduce the incident power, such as acousto-optic modulators (AOMs). We have

$$E_0 e^{i\omega t}\left[1-\frac{\beta}{2}(1-\sin\Omega t)\right] = E_0 e^{i\omega t}\left[1-\frac{\beta}{2}-i\frac{\beta}{4}(e^{i\Omega t}-e^{-i\Omega t})\right]. \tag{A.8}$$

Note that this expression is exact, so there are exactly two sidebands, even for full modulation ($\beta = 1$). In this case, whenever the sidebands are 180° out of phase, they are 90° out of phase from the carrier. Similarly, when they are in phase, they are either in phase or 180° out of phase with the carrier. This also follows the expectation following section A.1.

The similarity of these expressions to the amplitude and phase quadratures of equation (2.56) is not an accident[1]. Indeed, when calculating the total electric field in terms of the quantum operators, one finds $a_{1(2)}$ corresponds to the cosine (sine) quadrature [75, 76]. For a cosine carrier, these correspond to amplitude and phase modulation, respectively.

A.3 Thermal Noise Calculation via Normal-Mode Decomposition

To calculate the expected thermal noise of the new gram-scale suspensions (chapter 3), we use the normal-mode decomposition method originally devised by Gillespie and Raab [103]. In this technique, the mechanical system is decomposed into normal modes, each with some quality factor (and associated loss angle). The effect of a given mode is determined by the overlap of the mechanical mode and the laser beam. By scaling by the total oscillator mass, an "effective mass" for each mode can be calculated, and the thermal noise from each mode can then be summed together to estimate the total thermal noise.

[1] Conventionally the carrier is a cosine, whereas in these derivations we have used a sine modulation, which is why these are swapped.

Consider the displacement $\vec{u}_n(x,y)$ on the mirror surface S. This will induce a phase shift that will scatter the incident TEM_{00} mode with wavevector $\vec{k}$ into

$$\psi_{00} \rightarrow \psi_{00} e^{i2\vec{k}\cdot\vec{u}_n} \approx \psi_{00}\left[1 + i2\vec{k}\cdot\vec{u}_n - 2(\vec{k}\cdot\vec{u}_n)^2\right] \tag{A.9}$$

for mirror motion much smaller than the wavelength. In general, this will scatter light into other transverse modes, but (ideally) only ψ_{00} will resonate in the cavity. Thus we will only consider light that is scattered back into ψ_{00}. In this case, the imaginary term above represents a phase shift, while the last term represents the power scattered out of the fundamental mode. This phase shift represents an apparent length change of

$$\Delta l_n = \frac{1}{|\vec{k}|}\int_S \psi_{00}^* \psi_{00} \vec{k}\cdot\vec{u}_n(x,y)\, dS. \tag{A.10}$$

We can further normalize the contribution of each mode by calculating an "effective mass" of $\alpha_n m$ for each, which will be determined by the beam overlap with the mechanical mode. This is given by the ratio of the total energy in a mode U_n to the energy in the mode with a point-mass oscillator moving with amplitude Δl_n at frequency Ω_n:

$$\begin{aligned} \alpha_n &\equiv \frac{U_n}{\frac{1}{2}m\Omega_n^2\Delta l_n^2} \\ U_n &= \int_V \frac{1}{2}\rho(\vec{r})\Omega_n^2|\vec{u}_n(\vec{r})|^2\, dV \\ \Rightarrow m_{\mathrm{eff},n} = \alpha_n m &= \frac{\int_V \rho(\vec{r})|\vec{u}_n(\vec{r})|^2\, dV}{\Delta l_n^2}. \end{aligned} \tag{A.11}$$

From the fluctuation-dissipation theorem, the individual contribution of each mode to the thermal noise power spectrum is

$$G_{xx,n}(\Omega) = \frac{4k_BT}{\alpha_n m\Omega}\frac{\Omega_n^2\phi_n(\Omega)}{(\Omega^2-\Omega_n^2)^2 + \Omega_n^4\phi_n^2(\Omega)}. \tag{A.12}$$

The contribution from each mode can then be added together to get an estimate for the total thermal noise.

Appendix B

Schematics and Drawings

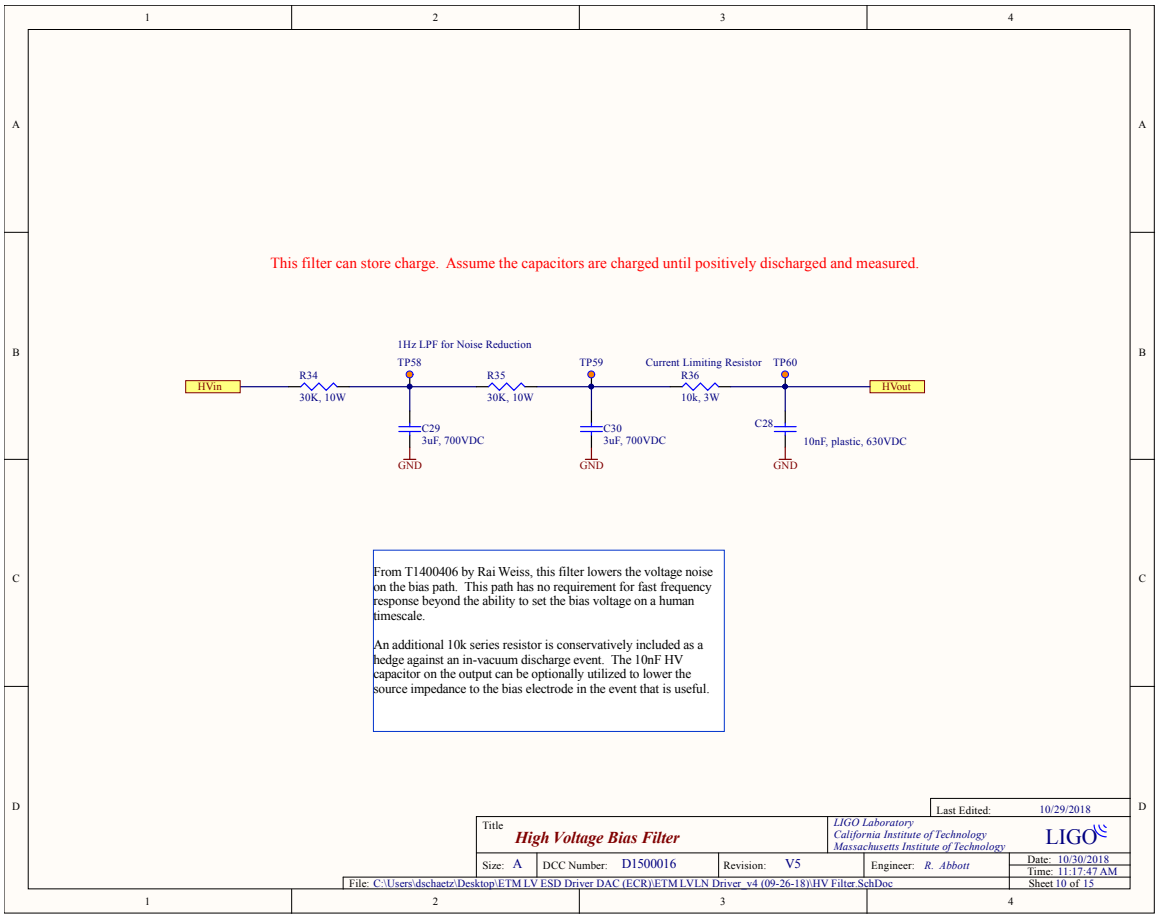

Figure B-1: Schematic showing the output of the high-voltage ESD driver. The output of this circuit is cabled into the vacuum chamber straight to the ESD bias trace. Note that the HV monitor is located prior to this low-pass filter, so the glitches discussed in 4.5.1 must occur after this filter.

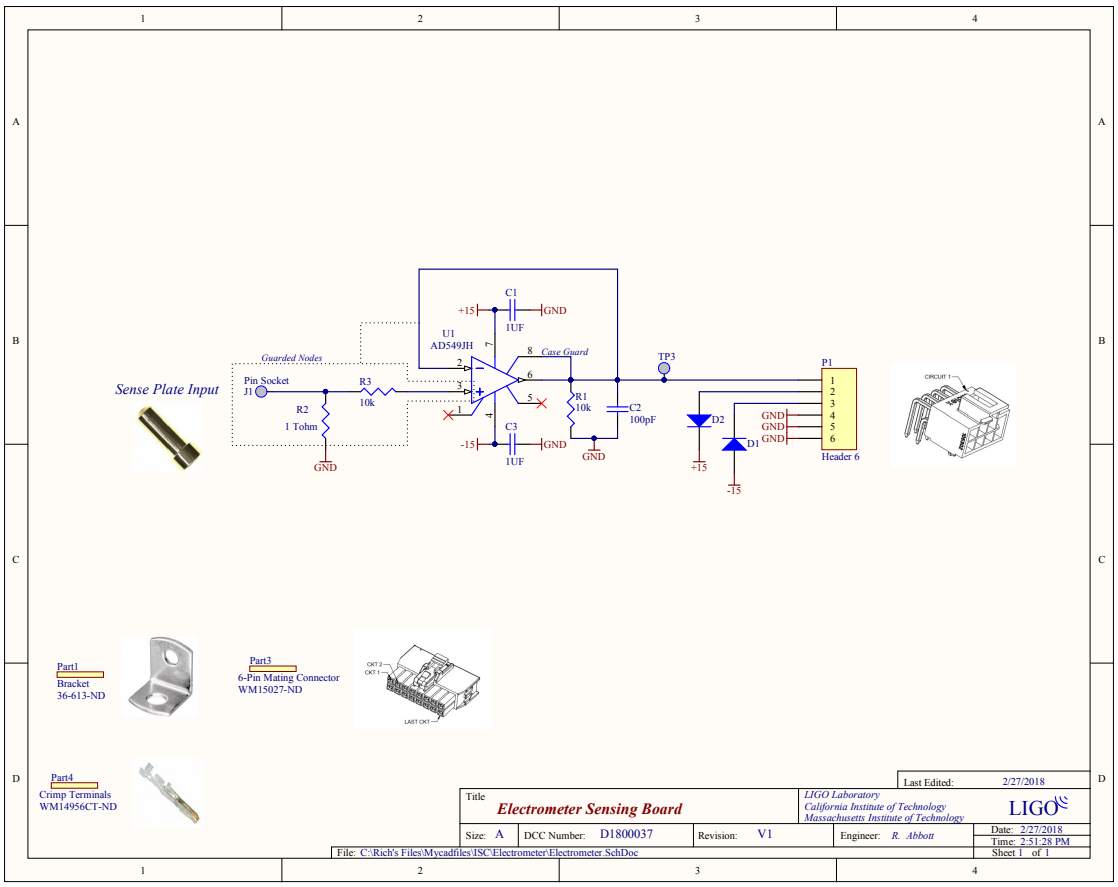

Figure B-2: Schematic of the sensing circuit for the EFM for each plate. The signals for a given axis are subtracted and converted to a differential output for readout in the digital system.

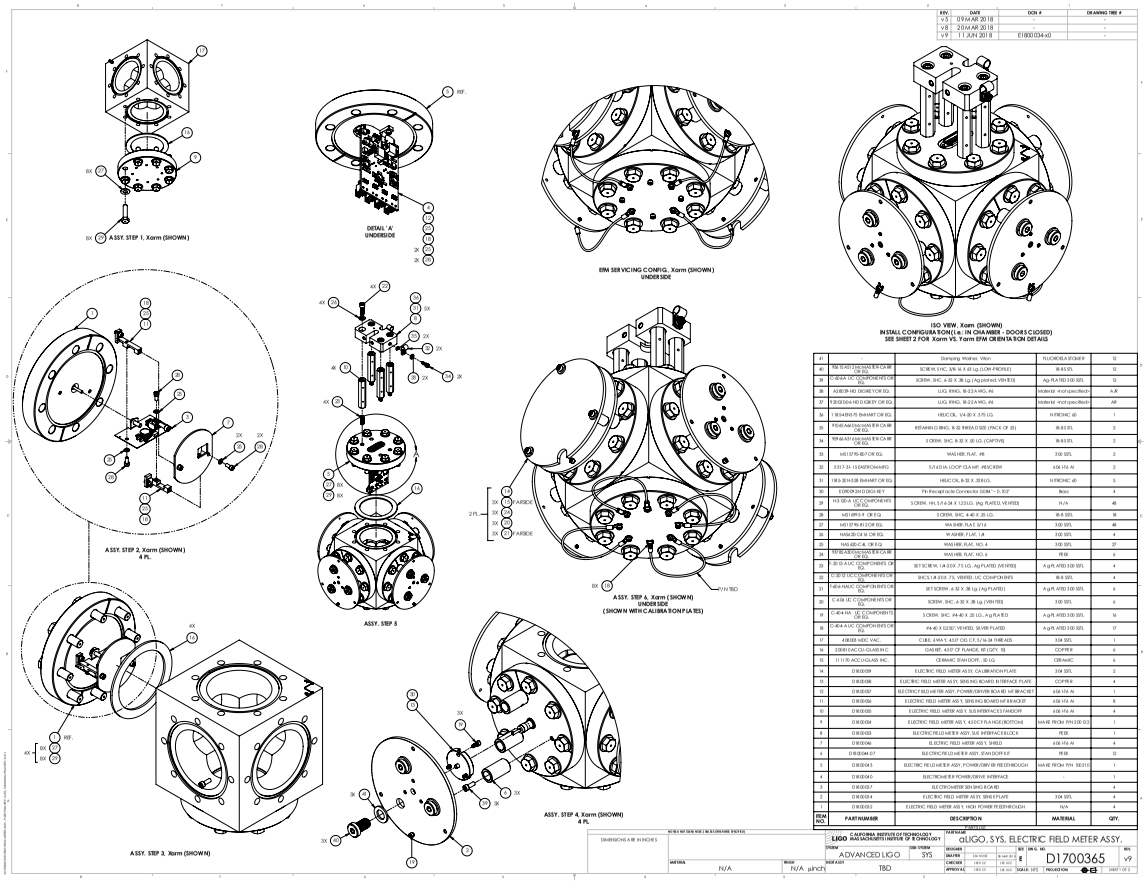

Figure B-3: Assembly drawings of the electric field meter. The sensing plates are mounted 1/2 inch away from the cube. Step 6 shows the EFM with additional calibration plates installed along one axis; see section 4.6.2 for details.

Bibliography

[1] J S Bendat and A G Piersol. *Random Data: Analysis and Measurement Procedures*. Wiley Series in Probability and Statistics. Wiley, 2011.

[2] Anthony E. Siegman. *Lasers*. University Science Books, 1986.

[3] Sydney Schreppler, Nicolas Spethmann, Nathan Brahms, et al. Optically measuring force near the standard quantum limit. *Science*, 344(6191):1486–1489, June 2014, 1312.4896.

[4] E. Oelker, R. B. Hutson, C. J. Kennedy, et al. Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks. *Nat. Photonics*, 13(10):714–719, October 2019.

[5] Aaron Chou, Henry Glass, H. Richard Gustafson, et al. The holometer: An instrument to probe Planckian quantum geometry. *Class. Quantum Gravity*, 34(6), 2017.

[6] Rana X. Adhikari. Gravitational radiation detection with laser interferometry. *Rev. Mod. Phys.*, 86(1):121–151, 2014, 1305.5188.

[7] D. V. Martynov, E. D. Hall, B. P. Abbott, et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93(11):112004, 2016.

[8] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *Am. J. Sci.*, s3-34(203):333–345, November 1887.

[9] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar PSR1913 + 16. *Nature*, 277(5696):437–440, 1979.

[10] B P Abbott, R Abbott, T D Abbott, et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116(6):61102, 2016, 1602.03837.

[11] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119(16):30–33, 2017, 1710.05836.

[12] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Multi-messenger observations of a binary neutron star merger. *Astrophys. J.*, 848(2):L12, 2017, 1710.05833.

[13] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs. *Phys. Rev. X*, 9(3):31040, 2019, 1811.12907.

[14] F. Acernese, M. Agathos, K. Agatsuma, et al. Advanced Virgo: A second-generation interferometric gravitational wave detector. *Class. Quantum Gravity*, 32(2), 2015, 1408.3978.

[15] C. Affeldt, K. Danzmann, K. L. Dooley, et al. Advanced techniques in GEO 600. *Class. Quantum Gravity*, 31(22), 2014.

[16] Yoichi Aso, Yuta Michimura, Kentaro Somiya, et al. Interferometer design of the KAGRA gravitational wave detector. *Phys. Rev. D*, 88(4):043007, August 2013, 1306.6747.

[17] M Tse, Haocun Yu, N Kijbunchoo, et al. Quantum-enhanced Advanced LIGO detectors in the era of gravitational-wave astronomy. *Phys. Rev. Lett.*, 123(23):231107, 2019.

[18] S. M. Aston, M. A. Barton, A. S. Bell, et al. Update on quadruple suspension design for Advanced LIGO. *Class. Quantum Gravity*, 29(23), 2012.

[19] Jorge L. Cervantes-Cota, Salvador Galindo-Uribarri, and George F. Smoot. A brief history of gravitational waves. *Universe*, 2(3):1–30, 2016.

[20] Sam Waldman, Peter Fritschel, and Rana Adhikari. Enhanced LIGO. Technical Report LIGO-T060156, November 2009.

[21] J. Aasi, J. Abadie, B. P. Abbott, et al. Characterization of the LIGO detectors during their sixth science run. *Class. Quantum Gravity*, 32(11):0–30, 2015, 1410.7764.

[22] S. Wen, R. Mittleman, K. Mason, et al. Hydraulic external pre-isolator system for ligo. *Class. Quantum Gravity*, 31(23):235001, December 2014.

[23] Tobin T Fricke, Nicolás D Smith-Lefebvre, Richard Abbott, et al. DC readout experiment in Enhanced LIGO. *Class. Quantum Gravity*, 29(6):065005, 2012, 1110.2815.

[24] P Kwee, C Bogan, K Danzmann, et al. Stabilized high-power laser system for the gravitational wave detector Advanced LIGO. *Opt. Express*, 20(10):10617, May 2012.

[25] Lisa Barsotti, Peter Fritschel, Matthew Evans, and Slawomir Gras. Updated Advanced LIGO sensitivity design curve. Technical Report LIGO-T1800044-v5, April 2018.

[26] J. Aasi, B P Abbott, R Abbott, et al. Advanced LIGO. *Class. Quantum Gravity*, 32(7):074001, April 2015, 1411.4547.

[27] Jocelyn S. Read, Charalampos Markakis, Masaru Shibata, et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys. Rev. D*, 79(12):124033, June 2009, 0901.3258.

[28] John A. Sidles and Daniel Sigg. Optical torques in suspended Fabry–Perot interferometers. *Phys. Lett. A*, 354(3):167–172, May 2006.

[29] Matthew Evans, Slawek Gras, Peter Fritschel, et al. Observation of parametric instability in Advanced LIGO. *Phys. Rev. Lett.*, 114(16):6–9, 2015.

[30] Aidan F. Brooks, Benjamin Abbott, Muzammil A. Arain, et al. Overview of Advanced LIGO adaptive optics. *Appl. Opt.*, 55(29):8256, October 2016.

[31] S. Biscans, S. Gras, C. D. Blair, et al. Suppressing parametric instabilities in LIGO using low-noise acoustic mode dampers. *Phys. Rev. D*, 100(12):122003, 2019, 1909.07805.

[32] Aaron Buikema, Craig Cahillane, Georgia Mansell, Carl Blair, et al. Sensitivity and performance of the Advanced LIGO detectors in the third observing run. *In prep.*

[33] F. Ya Khalili. Optimal configurations of filter cavity in future gravitational-wave detectors. *Phys. Rev. D*, 81(12):122002, June 2010.

[34] C Gerry, P Knight, and P L Knight. *Introductory Quantum Optics*. Cambridge University Press, 2005.

[35] Martin W Regehr. *Signal Extraction and Control for an Interferometric Gravitational Wave Detector*. PhD thesis, California Institute of Technology, 1995.

[36] Kiwamu Izumi and Daniel Sigg. Advanced LIGO: Length sensing and control in a dual recycled interferometric gravitational wave antenna. *Class. Quantum Gravity*, 34(1), 2017.

[37] M. Rakhmanov, R.L Savage, D.H Reitze, and D.B Tanner. Dynamic resonance of light in Fabry–Perot cavities. *Phys. Lett. A*, 305(5):239–244, December 2002.

[38] E Hecht. *Optics*. Pearson Education, 2008.

[39] John Bechhoefer. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.*, 77(3):783–836, August 2005.

[40] Jonathan Cripe, Nancy Aggarwal, Robinjeet Singh, et al. Radiation-pressure-mediated control of an optomechanical cavity. *Phys. Rev. A*, 97(1):1–6, 2018, 1710.04700.

[41] R. L. Barger, M.S. Sorem, and J.L. Hall. Frequency stabilization of a cw dye laser. *Appl. Phys. Lett.*, 22(11):573–575, June 1973.

[42] Benjamin S. Sheard, Malcolm B. Gray, Conor M. Mow-Lowry, David E. McClelland, and Stanley E. Whitcomb. Observation and characterization of an optical spring. *Phys. Rev. A*, 69(5):051801, May 2004.

[43] R. V. Pound. Electronic frequency stabilization of microwave oscillators. *Rev. Sci. Instrum.*, 17(11):490–505, 1946.

[44] R. W. P. Drever, J. L. Hall, F. V. Kowalski, et al. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B Photophysics Laser Chem.*, 31(2):97–105, June 1983.

[45] Eric D. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *Am. J. Phys.*, 69(1):79–87, 2001.

[46] TW Hänsch and B Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. *Opt. Commun.*, 35(3):441–444, December 1980.

[47] Adam Libson, Nicolas Brown, Aaron Buikema, et al. Simple method for locking birefringent resonators. *Opt. Express*, 23(3):3809, February 2015.

[48] J.M. Boon-Engering, W.E. van der Veer, E.A.J.M. Bente, and W. Hogervorst. Stabilization of an optical cavity containing a birefringent element. *Opt. Commun.*, 140(4-6):285–288, August 1997.

[49] Y. Honda, H. Shimizu, M. Fukuda, et al. Stabilization of a non-planar optical cavity using its polarization property. *Opt. Commun.*, 282(15):3108–3112, August 2009.

[50] Shigenori Moriwaki, Takumi Mori, Kohei Takeno, and Norikatsu Mio. Frequency discrimination method making use of polarization selectivity of triangular optical cavity. *Appl. Phys. Express*, 2(1):016501, January 2009.

[51] Peter Asenbaum and Markus Arndt. Cavity stabilization using the weak intrinsic birefringence of dielectric mirrors. *Opt. Lett.*, 36(19):3720–3722, October 2011.

[52] G. Gagliardi, S. De Nicola, P. Ferraro, and P. De Natale. Interrogation of fiber bragg-grating resonators by polarization-spectroscopy laser-frequency locking. *Opt. Express*, 15(7):3715–3728, 2007.

[53] D V Strekalov, R J Thompson, L M Baumgartel, I S Grudinin, and N Yu. Temperature measurement and stabilization in a birefringent whispering gallery mode resonator. *Opt. Express*, 19(15):14495–14501, July 2011, 1102.4227.

[54] A. Schliesser, R. Rivière, G. Anetsberger, O. Arcizet, and T. J. Kippenberg. Resolved-sideband cooling of a micromechanical oscillator. *Nat. Phys.*, 4(5):415–419, April 2008.

[55] Garrett D. Cole, Wei Zhang, Michael J. Martin, Jun Ye, and Markus Aspelmeyer. Tenfold reduction of Brownian noise in high-reflectivity optical coatings. *Nat. Photonics*, 7(8):644–650, July 2013.

[56] T. Kessler, C. Hagemann, C. Grebing, et al. A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity. *Nat. Photonics*, 6(10):687–692, September 2012.

[57] Feng Zhang and John W. Y. Lit. Direct-coupling single-mode fiber ring resonator. *J. Opt. Soc. Am. A*, 5(8):1347–1355, August 1988.

[58] B Lamouroux, B Prade, and A Orszag. Polarization effect in optical-fiber ring resonators. *Opt. Lett.*, 7(8):391–393, August 1982.

[59] K Iwatsuki, K Hotate, and M Higashiguchi. Eigenstate of polarization in a fiber ring resonator and its effect in an optical passive ring-resonator gyro. *Appl. Opt.*, 25(15):2606–2612, August 1986.

[60] Wei Wang and Junlei Xia. The characterizations of polarization in resonator integrated optic gyroscope. *Opt. Quantum Electron.*, 42(5):313–325, May 2011.

[61] Pierre-Henri Merrer, Olivier Llopis, and Gilles Cibiel. Laser stabilization on a fiber ring resonator and application to RF filtering. *IEEE Photonics Technol. Lett.*, 20(16):1399–1401, August 2008.

[62] T J Paul and E A Swanson. Narrow-linewidth diode laser using a fiber-optic ring resonator. *Opt. Lett.*, 18(15):1241–1243, August 1993.

[63] William F. Brinkman, Thomas L. Koch, David V. Lang, and Daniel P. Wilt. The lasers behind the communications revolution. *Bell Labs Tech. J.*, 5(1):150–167, August 2002.

[64] Martin Ostermeyer, Philip Kappe, Ralf Menzel, and Volker Wulfmeyer. Diode-pumped Nd:YAG master oscillator power amplifier with high pulse energy, excellent beam quality, and frequency-stabilized master oscillator as a basis for a next-generation lidar system. *Appl. Opt.*, 44(4):582–590, February 2005.

[65] H C Lefevre, R A Bergh, and H J Shaw. All-fiber gyroscope with inertial-navigation short-term sensitivity. *Opt. Lett.*, 7(9):454–456, September 1982.

[66] WK Burns. Polarization characteristics of single-mode fiber couplers. *IEEE Trans. Microw. Theory Tech.*, 30(10):1577–1588, October 1982.

[67] Zhang Xiao-Guang and Zheng Yuan. The number of least degrees of freedom required for a polarization controller to transform any state of polarization to any other output covering the entire poincaré sphere. *Chinese Phys. B*, 17(7):2509–2513, July 2008.

[68] E.P. Ippen. Stimulated Brillouin scattering in optical fibers. *Appl. Phys. Lett.*, 21(11):539–541, 1972.

[69] D A Shaddock, M B Gray, and D E McClelland. Frequency locking a laser to an optical cavity by use of spatial mode interference. *Opt. Lett.*, 24(21):1499–1501, November 1999.

[70] John Miller and Matthew Evans. Length control of an optical resonator using second-order transverse modes. *Opt. Lett.*, 39(8):2495–2498, April 2014.

[71] Euan Morrison, Brian J. Meers, David I. Robertson, and Henry Ward. Automatic alignment of optical interferometers. *Appl. Opt.*, 33(22):5041, 1994.

[72] Patrick Kwee, Frank Seifert, Benno Willke, and Karsten Danzmann. Laser beam quality and pointing measurement with an optical resonator. *Rev. Sci. Instrum.*, 78(7), 2007.

[73] L Barsotti, M Evans, and P Fritschel. Alignment sensing and control in Advanced LIGO. *Class. Quantum Gravity*, 27(8):084026, 2010.

[74] Markus Aspelmeyer, Tobias J. Kippenberg, and Florian Marquardt. Cavity optomechanics. *Rev. Mod. Phys.*, 86(December), 2014, 0712.1618.

[75] Carlton M. Caves and Bonny L Schumaker. New formalism for two-photon quantum optics. I. Quadrature phases and squeezed states. *Phys. Rev. A*, 31(5):3068–3092, 1985.

[76] Bonny L. Schumaker and Carlton M. Caves. New formalism for two-photon quantum optics. II. Mathematical foundation and compact notation. *Phys. Rev. A*, 31(5):3093–3111, 1985.

[77] Carlton Caves. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Phys. Rev. Lett.*, 45(2):75–79, July 1980.

[78] Vladimir B Braginsky, Farid Ya Khalili, and Kip S Thorne. *Quantum Measurement*. Cambridge University Press, 1992.

[79] Alessandra Buonanno and Yanbei Chen. Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors. *Phys. Rev. D*, 64(4):1–21, 2001, 0102012.

[80] H. J. Kimble, Yuri Levin, Andrey B. Matsko, Kip S. Thorne, and Sergey P. Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics. *Phys. Rev. D*, 65(2):1–31, 2001, 0008026.

[81] Abraham R. Neben, Timothy P. Bodiya, Christopher Wipf, et al. Structural thermal noise in gram-scale mirror oscillators. *New J. Phys.*, 14, 2012.

[82] Thomas Corbitt, David Ottaway, Edith Innerhofer, Jason Pelc, and Nergis Mavalvala. Measurement of radiation-pressure-induced optomechanical dynamics in a suspended Fabry-Perot cavity. *Phys. Rev. A*, 74(2):021802, August 2006.

[83] Thomas Corbitt, Yanbei Chen, Edith Innerhofer, et al. An all-optical trap for a gram-scale mirror. *Phys. Rev. Lett.*, 98(15):11–14, 2007, 0612188.

[84] Thomas Corbitt, Christopher Wipf, Timothy Bodiya, et al. Optical dilution and feedback cooling of a gram-scale oscillator to 6.9 mK. *Phys. Rev. Lett.*, 99(16):6–9, 2007, 0705.1018.

[85] G. S. Agarwal and Sumei Huang. Electromagnetically induced transparency in mechanical effects of light. *Phys. Rev. A*, 81(4):041803, April 2010, 0911.4157.

[86] T. Bodiya, V. Sudhir, C. Wipf, et al. Sub-hertz optomechanically induced transparency with a kilogram-scale mechanical oscillator. *Phys. Rev. A*, 100(1):013853, 2019.

[87] Timothy Paul Bodiya. Extreme Optomechanically Induced Transparency. Master's thesis, Massachusetts Institute of Technology, 2015.

[88] T. P. Purdy, R. W. Peterson, and C. A. Regal. Observation of radiation pressure shot noise on a macroscopic object. *Science*, 339(6121):801–804, February 2013, 1209.6334.

[89] Jonathan Cripe, Nancy Aggarwal, Robert Lanza, et al. Measurement of quantum back action in the audio band at room temperature. *Nature*, 568(7752):364–367, 2019.

[90] Min Jet Yap, Jonathan Cripe, Georgia L. Mansell, et al. Broadband reduction of quantum radiation pressure noise via squeezed light injection. *Nat. Photonics*, 2019, 1812.09804.

[91] Christopher Wipf. *Toward quantum opto-mechanics in a gram-scale suspended mirror interferometer*. PhD thesis, Massachusetts Institute of Technology, 2013.

[92] Thomas Randall Corbitt. *Quantum Noise and Radiation Pressure Effects in High Power Optical Interferometers*. PhD thesis, Massachusetts Institute of Technology, 2008.

[93] Rick L Savage, P J King, and S U Seel. A highly stabilized 10-watt Nd:YAG laser for the Laser Interferometer Gravitational-Wave Observatory (LIGO). *Laser Phys.*, 8(3):679–685, 1998.

[94] R. S. Abbott and P. J. King. Diode-pumped Nd:YAG laser intensity noise suppression using a current shunt. *Rev. Sci. Instrum.*, 72(2):1346–1349, 2001.

[95] B Willke, N Uehara, E K Gustafson, et al. Spatial and temporal filtering of a 10-W Nd:YAG laser with a Fabry-Perot ring-cavity premode cleaner. *Opt. Lett.*, 23(21):1704–6, November 1998.

[96] Seiji Kawamura and Janeen Hazel. Small optics suspension (SOS) final design (mechanical system). Technical Report LIGO-T970135, July 1997.

[97] Joseph Giaime, Partha Saha, David Shoemaker, and Lisa Sievers. A passive vibration isolation stack for LIGO: Design, modeling, and testing. *Rev. Sci. Instrum.*, 67(1):208–214, 1996.

[98] Theodore Welton and Herbert Callen. Irreversibility and generalized noise. *Phys. Rev.*, 83(1):34–40, 1951.

[99] J. B. Johnson. Thermal agitation of electricity in conductors. *Phys. Rev.*, 32(1):97–109, July 1928.

[100] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32(1):110–113, July 1928.

[101] A. L. Kimball and D. E. Lovell. Internal friction in solids. *Phys. Rev.*, 30(6):948–959, December 1927.

[102] S. Gras and M. Evans. Direct measurement of coating thermal noise in optical resonators. *Phys. Rev. D*, 98(12):122001, 2018.

[103] A. Gillespie and F. Raab. Thermally excited vibrations of the mirrors of laser interferometer gravitational-wave detectors. *Phys. Rev. D*, 52(2):577–585, July 1995.

[104] Yu Levin. Internal thermal noise in the LIGO test masses: A direct approach. *Phys. Rev. D*, 57(2):659–663, January 1998.

[105] Nancy Aggarwal. *A room temperature optomechanical squeezer*. PhD thesis, Massachusetts Institute of Technology, 2019.

[106] David J Ottaway, Peter Fritschel, and Samuel J. Waldman. Impact of upconverted scattered light on advanced interferometric gravitational wave detectors. *Opt. Express*, 20(8):8329, 2012.

[107] Jean-Yves Vinet, Violette Brisson, and Stefano Braccini. Scattered light noise in gravitational wave interferometric detectors: Coherent effects. *Phys. Rev. D*, 54(2):1276–1286, July 1996.

[108] Jean-Yves Vinet, Violette Brisson, Stefano Braccini, et al. Scattered light noise in gravitational wave interferometric detectors: A statistical approach. *Phys. Rev. D*, 56(10):6085–6095, 1997.

[109] Koji Arai, Sam Barnum, Peter Fritschel, Jeffrey Lewis, and Sam Waldman. Output mode cleaner design. Technical Report LIGO-T1000276-v5, January 2013.

[110] Safura Sharifi, Yaser Banadaki, Torrey Cullen, et al. Design of microresonators to minimize thermal noise below the standard quantum limit, 2019, 1911.02200.

[111] Kentaro Komori. *Optomechanical Torsion Pendulum for Measurement of Quantum Radiation Pressure Fluctuation*. PhD thesis, University of Tokyo, 2019.

[112] G. Manson and G. Hoffmann De Visme. The frequency spectrum of Barkhausen noise. *J. Phys. D. Appl. Phys.*, 5(8):1389–1395, 1972.

[113] Rainer Weiss. Collection of reports on Barkhausen noise. Technical Report LIGO-T0900061, 2009.

[114] J Miller. *On Non-Gaussian Beams and Optomechanical Parametric Instabilities in Interferometric Gravitational Wave Detectors*. PhD thesis, University of Glasgow, 2010.

[115] Carl Blair, Slawek Gras, Richard Abbott, et al. First demonstration of electrostatic damping of parametric instability at Advanced LIGO. *Phys. Rev. Lett.*, 118(15):1–7, 2017.

[116] Denis Martynov and Matthew Evans. LLO logbook: ESD actuation and charge. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=14853, September 2014.

[117] Jeffrey Kissel. LHO logbook: Charge measurement update. https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=24547, December 2015.

[118] John Miller and Matthew Evans. The ESD and the ring heater: musings on LLO14853. Technical Report LIGO-T1500076-v1, February 2015.

[119] L. G. Prokhorov and V. P. Mitrofanov. Space charge polarization in fused silica test masses of a gravitational wave detector associated with an electrostatic drive. *Class. Quantum Gravity*, 27(22), 2010.

[120] V Mitrofanov, L Prokhorov, K Tokmakov, and P Willems. Investigation of effects associated with variation of electric charge on a fused silica test mass. *Class. Quantum Gravity*, 21(5):S1083–S1089, March 2004.

[121] H. J. Wintle. Charge motion and trapping in insulators: Surface and bulk effects. *IEEE Trans. Dielectr. Electr. Insul.*, 6(1):1–10, 1999.

[122] Minora Tomozawa and Dong Wook Shin. Charge carrier concentration and mobility of ions in a silica glass. *J. Non. Cryst. Solids*, 241(2-3):140–148, 1998.

[123] Y. Awakuni and J. H. Calderwood. Water vapour adsorption and surface conductivity in solids. *J. Phys. D. Appl. Phys.*, 5(5):1038–1045, 1972.

[124] P. Ho, K. Lehovec, and L. Fedotowsky. Charge motion on silicon oxide surfaces. *Surf. Sci.*, 6(4):440–460, 1967.

[125] M. Hewitson, K. Danzmann, H. Grote, et al. Charge measurement and mitigation for the main test masses of the GEO 600 gravitational wave observatory. *Class. Quantum Gravity*, 24(24):6379–6391, 2007.

[126] D J Griffiths. *Introduction to Electrodynamics*. Prentice Hall, 1999.

[127] Rainer Weiss. Note on electrostatics in the LIGO suspensions. Technical Report LIGO-T960137, August 1996.

[128] Brian Lantz. Estimate of force coupling between charged optic face and optic cage. Technical Report LIGO-G1401179-v3, January 2016.

[129] P. Campsie, J. Hough, S. Rowan, and G. D. Hammond. A measurement of noise created by fluctuating electrostatic charges on dielectric surfaces using a torsion balance. *Class. Quantum Gravity*, 31(17), 2014.

[130] Leonid Prokhorov and Jeffrey Kissel. Actuation force of the ESD with charges and cage effects. Technical Report LIGO-T1500467-v1, August 2015.

[131] Sheila Dwyer, Jeffrey Kissel, and Nergis Mavalvala. ESD actuation strength measurements. Technical Report LIGO-T1700446-v1, September 2017.

[132] Shivaraj Kandhasamy and Arnaud Pele. LLO logbook: Tidal and ETMY actuation strength. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=31460, February 2017.

[133] Shivaraj Kandhasamy, Arnaud Pele, and Joseph Betzwieser. LLO logbook: Tidal control vs ESD actuation strength. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=49561, October 2019.

[134] D. Ugolini, C. Fitzgerald, I. Rothbarth, and J. Wang. Discharging fused silica optics occluded by an electrostatic drive. *Rev. Sci. Instrum.*, 85(3), 2014.

[135] Antonios Kontos, Rainer Weiss, Michael Zucker, Scott McCormick, and Harry Overmier. Photon emission spectrum of ion pumps. Technical Report LIGO-T1900858, November 2015.

[136] Friedemann Freund, Michael M. Masuda, and Minoru M. Freund. Highly mobile oxygen hole-type charge carriers in fused silica. *J. Mater. Res.*, 6(8):1619–1622, 1991.

[137] J F O'Hanlon. *A user's guide to vacuum technology*. A Wiley-Interscience Publication. Wiley, 1980.

[138] D. V. Koptsov, L. G. Prokhorov, and V. P. Mitrofanov. Effects of humidity on the interaction between a fused silica test mass and an electrostatic drive. *Phys. Lett. Sect. A Gen. At. Solid State Phys.*, 379(40-41):2535–2540, 2015.

[139] Alex Marin. Cosmic muon signature in LIGO. Technical Report LIGO-G970257, November 1997.

[140] P.E. Khramchenkov, Valery Mitrofanov, and Leonid Prokhorov. Test mass charge relaxation. Technical Report LIGO-G070204, March 2007.

[141] Aaron Buikema, Denis Martynov, Carl Blair, Rachel Gray, and Joseph Betzwieser. LLO logbook: ETMY ESD quadrant actuators not balanced and likely changing differently in time. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=43224, February 2019.

[142] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *Phys. Rev. D*, 95(6):062003, March 2017.

[143] Craig Cahillane, Joe Betzwieser, Duncan A. Brown, et al. Calibration uncertainty for Advanced LIGO's first and second observing runs. *Phys. Rev. D*, 96(10):1–16, 2017, 1708.03023.

[144] S. Karki, D. Tuyenbayev, S. Kandhasamy, et al. The Advanced LIGO photon calibrators. *Rev. Sci. Instrum.*, 87(11), 2016, 1608.05055.

[145] H W Ott. *Noise reduction techniques in electronic systems*. A Wiley-Interscience publication. Wiley, 1988.

[146] Robert Schofield, Anamaria Effler, and Terra Hardwick. LLO logbook: Multi-amp power grid currents in the LLO beam tubes are responsible for crab killer glitches and may contribute to beam tube corrosion. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=16036, December 2014.

[147] Carl Blair, Michael Laxen, and Aaron Buikema. LLO logbook: Electric field noise investigations. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=41707, November 2018.

[148] N. Christensen, K. Riles, J. D. Tasson, et al. Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97(8):82002, 2018.

[149] Sebastien Biscans and Peter Fritschel. 20-bit DAC noise update. Technical Report LIGO-G1500761, June 2015.

[150] Christopher Wipf, Michael Laxen, and Aaron Buikema. LLO logbook: Offset-dependent DAC noise not due to digital filtering effects. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=43268, February 2019.

[151] Aaron Buikema. LLO logbook: Ring heater coupling to DARM. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=42534, December 2018.

[152] Aaron Buikema. LLO logbook: Line cancellation with ESD bias. https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=43729, February 2019.

[153] V.P Mitrofanov, L.G Prokhorov, and K.V Tokmakov. Variation of electric charge on prototype of fused silica test mass of gravitational wave antenna. *Phys. Lett. A*, 300(4-5):370–374, August 2002.

[154] M Cabero, A Lundgren, A H Nitz, et al. Blip glitches in Advanced LIGO data. *Class. Quantum Gravity*, 36(15):155010, August 2019.

[155] M. Zevin, S. Coughlin, S. Bahaadini, et al. Gravity Spy: integrating advanced LIGO detector characterization, machine learning, and citizen science. *Class. Quantum Gravity*, 34(6):064003, March 2017.

[156] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Class. Quantum Gravity*, 21(20 SPEC. ISS.), 2004, 0412119.

[157] Shourov Keith Chatterji. *The search for gravitational wave bursts in data from the second LIGO science run.* PhD thesis, Massachusetts Institute of Technology, 2005.

[158] Ltspice spice simulator. www.analog.com/LTspice.

[159] Rich Abbott and Carl Blair. Understanding the EFM calibration. Technical Report LIGO-T1800464-v1, October 2018.

[160] S. Matsusaka, H. Maruyama, T. Matsuyama, and M. Ghadiri. Triboelectric charging of powders: A review. *Chem. Eng. Sci.*, 65(22):5781–5807, 2010, arXiv:1011.1669v3.

[161] Carl Blair, Rich Abbott, and Aaron Buikema. Note on electric field meter sense board glitch experiments. Technical Report LIGO-T1800454-v1, May 2019.

[162] A. S. Markosyan, R. Bassiri, R. M. Faris, et al. Al-doped ZnO amorphous films as conductive layers in ultra-low absorptive optical coatings. *Opt. InfoBase Conf. Pap.*, pages 30–32, 2014.

[163] Richard Savage, Peter Fritschel, Benno Willke, and Peter King. Pre-stabilized laser design requirements. Technical Report LIGO-T050036-v4, May 2010.

[164] L. Winkelmann, O. Puncken, R. Kluzik, et al. Injection-locked single-frequency laser with an output power of 220 W. *Appl. Phys. B*, 102(3):529–538, March 2011.

[165] Michalis N. Zervas and Christophe A. Codemard. High power fiber lasers: A review. *IEEE J. Sel. Top. Quantum Electron.*, 20(5):219–241, 2014.

[166] Guido Mueller. Beam jitter coupling in Advanced LIGO. *Opt. Express*, 13(18):7118, 2005.

[167] T Y Fan. Laser beam combining for high power, high-radiance sources. *IEEE J. Sel. Top. Quantum Eletron.*, 11(3):567–577, 2005.

[168] Omar De Varona, Willy Fittkau, Phillip Booker, et al. Single-frequency fiber amplifier at 15 μm with 100 W in the linearly-polarized TEM_{00} mode for next-generation gravitational wave detectors. *Opt. Express*, 25(21):24880, 2017.

[169] Eisuke Fujita, Yutaka Mashiko, Seira Asaya, Mitsuru Musha, and Masaki Tokurakawa. High power narrow-linewidth linearly-polarized 1610 nm Er:Yb all-fiber MOPA. *Opt. Express*, 24(23):26255, November 2016.

[170] Xiaolei Bai, Quan Sheng, Haiwei Zhang, et al. High-power all-fiber single-frequency erbium-ytterbium co-doped fiber master oscillator power amplifier. *IEEE Photonics J.*, 7(6), 2015.

[171] Daniel Creeden, Herman Pretorius, Julia Limongelli, and Scott D. Setzler. Single frequency 1560 nm Er:Yb fiber amplifier with 207 W output power and 50.5% slope efficiency. Number March 2016, page 97282L, March 2016.

[172] Ke Yin, Bin Zhang, Guanghui Xue, Lei Li, and Jing Hou. High-power all-fiber wavelength-tunable thulium doped fiber laser at 2 μm. *Opt. Express*, 22(17):19947, 2014.

[173] Alex Sincore, Joshua D Bradford, Justin Cook, Lawrence Shah, and Martin C Richardson. High average power thulium-doped silica fiber lasers: Review of systems and concepts. *IEEE J. Sel. Top. Quantum Electron.*, 24(3):1–8, May 2018.

[174] Alexander Hemming, Nikita Simakov, Alan Davidson, et al. A monolithic cladding pumped holmium-doped fibre laser. In *CLEO 2013*, page CW1M.1, Washington, D.C., 2013. OSA.

[175] Alexander Hemming, Nikita Simakov, John Haub, and Adrian Carter. A review of recent progress in holmium-doped silica fibre sources. *Opt. Fiber Technol.*, 20(6):621–630, 2014.

[176] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Exploring the sensitivity of next generation gravitational wave detectors. *Class. Quantum Gravity*, 34(4), 2017, 1607.08697.

[177] Brian Lantz, Stuart Reid, Riccardo Bassiri, et al. LSC instrument science white paper 2019. Technical Report LIGO-T1700231-v4, 2019.

[178] Stephan Wielandy. Implications of higher-order mode content in large mode area fibers with good beam quality. *Opt. Express*, 15(23):15402–15409, 2007.

[179] F. Beier, C. Hupel, S. Kuhn, et al. Single mode 43 kW output power from a diode-pumped Yb-doped fiber amplifier. *Opt. Express*, 25(13):14892, 2017.

[180] Pengfei Ma, Rumao Tao, Rongtao Su, et al. 1.89 kW all-fiberized and polarization-maintained amplifiers with narrow linewidth and near-diffraction-limited beam quality. *Opt. Express*, 24(4):4187, 2016.

[181] A. Liem, J. Limpert, H. Zellmer, and A. Tünnermann. 100-W single-frequency master-oscillator fiber power amplifier. *Opt. Lett.*, 28(17):1537, 2003.

[182] Michael Steinke, Henrik Tunnermann, Vincent Kuhn, et al. Single-frequency fiber amplifiers for next-generation gravitational wave detectors. *IEEE J. Sel. Top. Quantum Electron.*, 24(3), 2018.

[183] Felix Wellmann, Michael Steinke, Fabian Thies, et al. Characterization of the monolithic fiber amplifier engineering prototype for the next generation of gravitational wave detectors. In *Fiber Lasers XVI Technol. Syst.*, page 72. SPIE, March 2019.

[184] Lei Zhang, Shuzhen Cui, Chi Liu, Jun Zhou, and Yan Feng. 170 W, single-frequency, single-mode, linearly-polarized, Yb-doped all-fiber amplifier. *Opt. Express*, 21(5):5456, March 2013.

[185] V. I. Kovalev and R. G. Harrison. Suppression of stimulated Brillouin scattering in high-power single-frequency fiber amplifiers. *Opt. Lett.*, 31(2):161–163, 2006.

[186] Thomas Theeg, Christoph Ottenhues, Hakan Sayinc, Jörg Neumann, and Dietmar Kracht. Core-pumped single-frequency fiber amplifier with an output power of 158 W. *Opt. Lett.*, 41(1):9, January 2016.

[187] Jian Zhao, Germain Guiraud, Christophe Pierre, et al. High-power all-fiber ultra-low noise laser. *Appl. Phys. B*, 124(6):114, June 2018.

[188] Benjamin Pulford, Thomas Ehrenreich, Roger Holten, et al. 400-W near diffraction-limited single-frequency all-solid photonic bandgap fiber amplifier. *Opt. Lett.*, 40(10):2297, May 2015.

[189] Matthias Hildebrandt, Maik Frede, Patrick Kwee, Benno Willke, and Dietmar Kracht. Single-frequency master-oscillator photonic crystal fiber amplifier with 148 W output power. *Opt. Express*, 14(23):11071–11076, 2006.

[190] Aaron Buikema, Franklin Jose, Steven J. Augst, Peter Fritschel, and Nergis Mavalvala. Narrow-linewidth fiber amplifier for gravitational-wave detectors. *Opt. Lett.*, 44(15):3833, August 2019.

[191] Matthias Hildebrandt. *Brillouin scattering in high-power narrow-linewidth fiber amplifiers*. PhD thesis, Leibniz University Hannover, 2009.

[192] R W Boyd and D Prato. *Nonlinear Optics*. Elsevier Science, 2008.

[193] R. G. Smith. Optical power handling capacity of low loss optical fibers as determined by stimulated Raman and Brillouin scattering. *Appl. Opt.*, 11(11):2489, November 1972.

[194] M.O. van Deventer and A.J. Boot. Polarization properties of stimulated Brillouin scattering in single-mode fibers. *J. Light. Technol.*, 12(4):585–590, April 1994.

[195] Vincent Roy, Claude Paré, Pierre Laperle, Louis Desbiens, and Yves Taillon. Yb-doped large mode area fibers with depressed clad and dopant confinement. In John Ballato, editor, *Proc. SPIE*, volume 9728, page 97281W, April 2016.

[196] Benno Willke. Pre-stabilized laser subsystem testing and acceptance-L1. Technical Report LIGO-E1100716-v6, 2012.

[197] Jan Poeld, Peter King, and Benno Willke. Control loop design for the aLIGO second loop power stabilization. Technical Report LIGO-T1300867, 2013.

[198] Jian Zhao, Germain Guiraud, Florian Floissat, et al. Gain dynamics of clad-pumped Yb-fiber amplifier and intensity noise control. *Opt. Express*, 25(1):357, 2017.

[199] Steven J. Augst, T. Y. Fan, and Antonio Sanchez. Coherent beam combining and phase noise measurements of ytterbium fiber amplifiers. *Opt. Lett.*, 29(5):474, 2004.

[200] Michael Tröbs, Peter Weßels, and Carsten Fallnich. Power- and frequency-noise characteristics of an Yb-doped fiber amplifier and actuators for stabilization. *Opt. Express*, 13(6):2224, 2005.

[201] Patrick Kwee, Frank Seifert, Benno Willke, and Karsten Danzmann. Laser beam quality and pointing measurement with an optical resonator. *Rev. Sci. Instrum.*, 78(7), 2007.

[202] Steven X. Li, Scott A. Merritt, Michael A. Krainak, and Anthony Yu. Giant pulse phenomena in a high gain erbium doped fiber amplifier. In Hamid Hemmati and Don M. Boroson, editors, *Free. Laser Commun. Atmos. Propag. XXX*, volume 1052418, page 46. SPIE, February 2018.

[203] Arlee V Smith and Binh T Do. Bulk and surface laser damage of silica by picosecond and nanosecond pulses at 1064 nm. *Appl. Opt.*, 47(26):4812–4832, 2008.

[204] Raman Kashyap. The fiber fuse - from a curious effect to a critical issue: A 25th year retrospective. *Opt. Express*, 21(5):6422, 2013.

[205] Gadi Fibich and Alexander L. Gaeta. Critical power for self-focusing in bulk media and in hollow waveguides. *Opt. Lett.*, 25(5):335, 2000.

[206] Roger L Farrow, Dahv A V Kliner, G Ronald Hadley, and Arlee V Smith. Peak-power limits on fiber amplifiers imposed by self-focusing. *Opt. Lett.*, 31(23):3423, December 2006.

[207] Chujun Zhao, Ying Li, Dajun Lei, et al. Review of self-focusing of high power lasers in large-mode-area optical fibers. *J. Phys. Conf. Ser.*, 276(1):23–26, 2011.

[208] Fabian Thies, Nina Bode, Patrick Oppermann, et al. Nd:YVO_4 high-power master oscillator power amplifier laser system for second-generation gravitational wave detectors. *Opt. Lett.*, 44(3):719, February 2019.

[209] Lisa Barsotti, Lee McCuller, Matthew Evans, and Peter Fritschel. Output mode cleaner design. Technical Report LIGO-T1800042-v5, October 2018.

[210] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Exploring the sensitivity of next generation gravitational wave detectors. *Class. Quantum Gravity*, 34(4), 2017, 1607.08697.

[211] David Reitze, Rana X Adhikari, Stefan Ballmer, et al. Cosmic Explorer: The U.S. contribution to gravitational-wave astronomy beyond LIGO, 2019, 1907.04833.

[212] M. Punturo, M. Abernathy, F. Acernese, et al. The third generation of gravitational wave observatories and their science reach. *Class. Quantum Gravity*, 27(8), 2010.

[213] V. B. Braginsky and F. Ja Khalili. Gravitational wave antenna with QND speed meter. *Phys. Lett. A*, 147(5-6):251–256, 1990.

[214] E. Knyazev, S. Danilishin, S. Hild, and F.Ya. Khalili. Speedmeter scheme for gravitational-wave detectors based on EPR quantum entanglement. *Phys. Lett. A*, 382(33):2219–2225, August 2018, 1701.01694.

[215] Haixing Miao, Huan Yang, and Denis Martynov. Towards the design of gravitational-wave detectors for probing neutron-star physics. *Phys. Rev. D*, 98(4):44044, 2018, 1712.07345.

[216] Jorge Piekarewicz and Farrukh J. Fattoyev. Neutron-rich matter in heaven and on earth. *Phys. Today*, 72(7):30–37, 2019.