# Noise Cancellation for Gravitational Wave Detectors

Thesis by

Jennifer Clair Driggers

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2015
(Defended 15 May 2015)

# Acknowledgments

Many people have helped in numerous ways to support me, and have helped make this thesis what it is. Thank you all.

To my advisor, Rana Adhikari, thank you for everything. You've always pushed me to accomplish more than I thought that I could, which showed me how much I was actually capable of. I may never match your graphics arts style, but you have taught me an incredible amount, and I'm lucky to have been your student.

To my two favorite staff scientists, Koji Arai and Jan Harms, you guys have been my rocks that kept me sane, and helped me through it all. Very few people actually live the philosophy that "the only bad question is the one unasked", but you both do. I aspire to live up to your model. This thesis wouldn't exist without you, so thank you.

To all of my labmates and LIGO-Caltech friends, thanks for many great games of softball, many beers, and many good times all around. Steve V., Eric Q., Kiwamu I., Alberto S., Rob W., John M., Yuta M., Annalisa A., Jamie R., Joe B., Yoichi A., Bob T., Manasa T., Vivien R., Rory S., Haixing M., Dmass, Zach K., Evan H., Tara C., Pinkesh P., Den M., Aidan B., Nic S., Margot P., Ben A., Rich A., Alastair H., Larry P., Gabriele V., Matt A., Frank S., Kate D., Masha O., Eric G., Julie H., and everyone else, thank you for being excellent co-workers.

To everyone at the LIGO sites, thank you for being so welcoming when I visited, and for putting up with things like orange cones and sensors all over the floors.

To my "California family", you guys are the best. Seriously. Nick, Jonny, Andrew, Aaron, Kim, Tim, Elizabeth, Fujimi, Tara V., Becky, Nicole, and Harmony, thanks for convincing me to do things other than labwork sometimes.

To my parents, my grandparents, my brothers and the rest of my family, thank you for supporting me in everything that I do. All of this would have been much more difficult without you. In particular, to my parents, thank you for never suggesting that I *couldn't* do things like organize a Science Olympiad competition as an 8th grader, try to solve the Middle East conflict by having Palestinians and Israelis meet and talk in our living room, fix cars, program computers, aspire to be an astronaut, or do physics.

Butterball, you're a dog so you can't read this, but thanks for keeping my lap warm while I was

writing, and for making me smile on even the hardest days.

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This work was also supported by a National Science Foundation Graduate Research Fellowship, and a Caltech Millikan Graduate Fellowship.

This thesis has LIGO document number LIGO-P1500001.

# Abstract

The LIGO gravitational wave detectors are on the brink of making the first direct detections of gravitational waves. Noise cancellation techniques are described, in order to simplify the commissioning of these detectors as well as significantly improve their sensitivity to astrophysical sources. Future upgrades to the ground based detectors will require further cancellation of Newtonian gravitational noise in order to make the transition from detectors striving to make the first direct detection of gravitational waves, to observatories extracting physics from many, many detections. Techniques for this noise cancellation are described, as well as the work remaining in this realm.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

Einstein predicted the existence of gravitational waves as a consequence of general relativity in 1916 [1]. Unfortunately, gravitational waves couple very weakly to matter, so direct measurements require massive sources to create relatively large waves. The most promising sources are all astrophysical and cosmological, and include compact inspiralling binary systems and pulsars.

In the nearly 100 years since general relativity was first described, no direct measurements of gravitational waves have yet been made, although many projects have worked toward such detections. Projects such as Weber bar detectors (starting in the 1960's) [2] and Forward's interferometer (starting in the 1970's) [3] looked directly for gravitational waves in their output signals, while projects such as pulsar timing [4] look for periodic changes in the travel time of pulses from millisecond pulsars. LIGO, the Laser Interferometer Gravitational wave Observatory, built a 4 km long baseline interferometer and began collecting data in the 2000's, although no detections were made. Despite a lack of detections, each experiment provides useful technological experience for all subsequent projects, as well as sets upper limits on how strong gravitational waves from different sources could be. Utilizing this knowledge, Advanced LIGO, a follow-on of the Initial LIGO project, has constructed a world-class interferometer, and intends to make a first direct detection of gravitational waves within the next few years [5]. One of the ways that Advanced LIGO will meet its goals is the use of noise cancellation techniques developed and described in this thesis.

This thesis briefly introduces gravitational waves and the astrophysical motivations for measuring them in Chapter 2. Since the work in this thesis primarily improves the low-frequency sensitivity of the LIGO detectors, low-frequency gravitational wave sources are highlighted.

LIGO interferometers are the primary application of these techniques, and so they are discussed in Chapter 3. Also, the 40 m Prototype Laboratory on the Caltech campus is described in Chapter 4, as this is where much of the development of the techniques occurred.

One of the main results of this thesis is the global seismic noise cancellation that was developed and then deployed during the 2009-2010 5th LIGO Science Run. This work is discussed in Chapter 5, and draws from work published in [6] and [7]. In addition, adaptive variants of this noise cancellation are under development and testing at the 40 m Lab, and are described.

Another significant result presented in this thesis is the analysis of Newtonian gravitational noise, including measurements and simulations that show Newtonian noise will not be a significant issue for Advanced LIGO, but will be important for future generations of gravitational wave detectors. In preparation for the noise cancellation that will be required in the future, simulations were performed to determine the requirements for sensing Newtonian noise in order to suppress it below quantum noise levels. Chapter 6 discusses this work, and draws from work in a LIGO Technical Report [8], as well as work published in [9]. The final section of this chapter describes in detail future work that must be done before Newtonian noise cancellation can be applied to our interferometers.

Conclusions and future work are discussed in Chapter 7, an overview of basic control theory is given in Appendix A, a review of Pound-Drever-Hall locking is given in Appendix B and photos of the 40 m Lab's optical layout are included in Appendix C.

# Chapter 2

# Gravitational waves and astrophysical motivations

Various wavelengths of electromagnetic radiation have been used over the years to study many astrophysical events and phenomena. With the advent of new technology and telescopes that are able to observe in different wavelengths, humans have progressively been able to get an ever-clearer picture of the universe. One of the major restrictions of observations with electromagnetic radiation is the fact that light interacts strongly with matter; it will not pass unchanged through dense materials.

Gravitational waves interact very weakly with matter, and so are able to propagate through many materials, such as intergalactic dust, without significant distortion. Also, they are uniquely able to propagate unimpeded from the moment of the big bang to now, in contrast with electromagnetic waves that only propagate freely since the era of recombination, about 380,000 years after the big bang. Measuring gravitational waves and using them as another spectrum for observing the universe provides a whole host of new information. Gravitational wave astronomy will add a new dimension to our understanding of extreme conditions, including the core of neutron stars, black holes, and the moments immediately following the formation of the universe.

LIGO is a ground-based observatory, and so suffers from limitations at low frequencies, mostly due to seismic noise and thermal noise in the fibers required to suspend the test masses. Advanced LIGO expects to be sensitive to gravitational waves in the band between 10 Hz to a few kHz; inspiralling binary systems of neutron stars or solar mass black holes are some of the most promising sources in this frequency range.

LIGO aims to unveil the new era of gravitational wave astronomy with a first detection in the next few years [10]. Following that first detection, LIGO and companion gravitational wave observatories including VIRGO [11] and KAGRA [12] hope to measure many gravitational wave events with enough precision to determine parameters that describe the astrophysical event itself. With a network of detectors, we can infer the sky location of a source, allowing us to follow up a detection with electromagnetic telescopes, and garner even more information. Parameters such as distance to the source, mass ratio of the inspiralling objects, and spin of the objects, all depend on the makeup of the system creating the gravitational waves, and so extracting these parameters from the measured waveforms will provide invaluable information about the system.

While there are some uncertainties in the expected rates of detection of gravitational waves, due to uncertainty in population densities of various astrophysical systems, Advanced LIGO is expected to have many measurable signals per year. A summary of the best estimates for detection rates and the assumptions that are taken into account for binary systems is provided in [13]. As shown in Table 2.1, which shows pessimistic ($\dot{N}_{\rm low}$), realistic ($\dot{N}_{\rm re}$), optimistic ($\dot{N}_{\rm high}$) and maximal ($\dot{N}_{\rm max}$) rates of detections per year, Advanced LIGO expects to detect several tens of "standard" binary coalescences involving neutron stars (NS) and black holes (BH) per year. More unusual binary systems, such as an intermediate mass ratio inspiral merging into an intermediate mass black hole (IMRI into IMBH) or a pair of inspiralling intermediate mass black holes (IMBH-IMBH) have much lower expected rates.

| Interferometer generation | Source | $\dot{N}_{\rm low}$ yr$^{-1}$ | $\dot{N}_{\rm re}$ yr$^{-1}$ | $\dot{N}_{\rm high}$ yr$^{-1}$ | $\dot{N}_{\rm max}$ yr$^{-1}$ |
|---|---|---|---|---|---|
| Initial | NS-NS | $2 \times 10^{-4}$ | 0.02 | 0.2 | 0.6 |
| | NS-BH | $7 \times 10^{-5}$ | 0.004 | 0.1 | |
| | BH-BH | $2 \times 10^{-4}$ | 0.007 | 0.5 | |
| | IMRI into IMBH | | | $< 0.001$ | 0.01 |
| | IMBH-IMBH | | | $10^{-4}$ | $10^{-3}$ |
| Advanced | NS-NS | 0.4 | 40 | 400 | 1000 |
| | NS-BH | 0.2 | 10 | 300 | |
| | BH-BH | 0.4 | 20 | 1000 | |
| | IMRI into IMBH | | | 10 | 300 |
| | IMBH-IMBH | | | 0.1 | 1 |

Table 2.1: Table V from [13] comparing Initial LIGO's range of expected detection rates with Advanced LIGO. Pessimistic ($\dot{N}_{\rm low}$), realistic ($\dot{N}_{\rm re}$), optimistic ($\dot{N}_{\rm high}$), and maximal ($\dot{N}_{\rm max}$) rates of detections per year are shown for various types of compact binary inspirals. The source type indicates what the elements of the binary system are comprised of: a neutron star (NS), a black hole (BH), or an intermediate mass black hole (IMBH). IMRI into IMBH indicates a pair of compact objects with an intermediate mass ratio (between $10^{-2}$ and $10^{-3}$) coalescing into an intermediate mass black hole.

The detection and measurement of gravitational waveforms will open a new window on the universe. The work described in this thesis will facilitate the commissioning of the latest generation of gravitational wave detectors, and thus help bring this new era of gravitational wave astronomy into being.

## 2.1 Gravitational waves

Gravitational waves are perturbations of the spacetime metric. We assume that the dominant term in the metric of the universe is a flat Minkowski metric, $\eta$. With a perturbation term $h(t)$ to describe the gravitational waves, we have a total metric of the form

$$g(t) = \eta + h(t). \tag{2.1}$$

This follows from general relativity, which assumes that moving massive objects change the curvature of spacetime. Any accelerating mass quadrupole moment (mass moving with a lack of perfect spherical symmetry) will create time-dependent changes in spacetime, which we refer to as gravitational waves. Because of conservation of mass and momentum, the lowest order multipole term describing $h(t)$ will be the quadrupole moment of the expansion of the post-Newtonian solutions to Einstein's equations.

One of the quintessential examples of an astrophysical source that should be seen by Advanced LIGO is the inspiral of two compact objects, such as black holes or neutron stars. In the simple case of a pair of compact objects before the merger time, with the system optimally pointing toward Earth, the perturbation term in Equation 2.1 will look like

$$h_{\text{inspiral}}(t) = -\left(\frac{G\mathcal{M}}{c^2 D_{\text{eff}}}\right)\left(\frac{t_0 - t}{5G\mathcal{M}/c^3}\right)^{-1/4} \cos\left(2\theta_0 + 2\theta(t)\right). \tag{2.2}$$

Here, $G$ is Newton's gravitational constant, $c$ is the speed of light in a vacuum, and $t_0$ is the time of coalescence as measured at the detector. $\mathcal{M}$ is the chirp mass,

$$\mathcal{M} = M\rho^{(3/5)}, \tag{2.3}$$

where $M = m_1 + m_2$ is the total mass of the system and $\rho$ is the symmetric mass ratio,

$$\rho = \frac{m_1 m_2}{(m_1 + m_2)^2}. \tag{2.4}$$

$D_{\text{eff}}$ is the effective distance to the binary system. Since we cannot (with a single detector) disentangle the true distance from the orientation of the system, they are both included in this variable. $\theta_0$ is related to the phase of the wave at the time of coalescence, and $\theta(t)$ is the orbital phase of the binary [14]. $\theta(t)$, up to the second post-Newtonian order, is given by Equation 7 of [15] as

$$2\theta(t) = -\frac{1}{\rho}\left[\Theta^{5/8} + \left(\frac{3715}{8064} + \frac{55}{96}\rho\right)\Theta^{3/8} - \frac{3\pi}{4}\Theta^{1/4} + \left(\frac{9275495}{14450688} + \frac{284875}{258048}\rho + \frac{1855}{2048}\rho^2\right)\Theta^{1/8}\right], \tag{2.5}$$

where

$$\Theta = \frac{c^3 \rho}{5GM}(t_0 - t). \tag{2.6}$$

Figure 2.1 shows an example of the inspiral phase of a waveform for a pair of 1.4 solar mass (1.4 $\mathrm{M}_\odot$) neutron stars 1 Mpc from Earth. The inspiral phase is followed by a merger lasting a fraction of a second, which is followed by a ringdown of the final compact object.

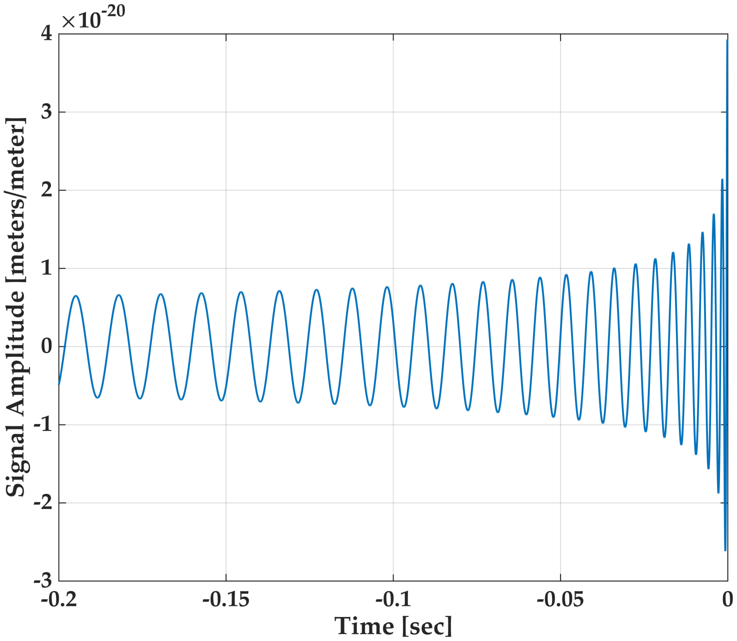

Figure 2.1: An example waveform described by Equation 2.2 for a pair of 1.4 $\mathrm{M}_\odot$ neutron stars, optimally oriented, 1 Mpc from Earth.

Since gravitational waves are changing spacetime, we should be able to detect these perturbations by observing distance changes between freely falling test masses. While much more detail on the theory behind gravitational waves and the way they perturb spacetime can be found in books such as *Gravitation* [16], here I sketch a derivation of how light can be used to measure gravitational waves. This motivates the use of interferometric measurement devices that can precisely measure the phase of light as a function of time for this purpose.

Let $\phi$ be the phase of an electromagnetic wave of frequency $\omega_0$. Assuming that the light begins with a phase of 0 and travels at the speed of light $c$, after it has gone a distance $L$ it will have a phase $\delta\phi$. We can express this as a distance integral over the spacetime metric,

$$\delta\phi = \frac{\omega_0}{c} \int_0^L g \, \mathrm{d}x, \tag{2.7}$$

where $g$ is defined by Equation 2.1. Plugging Equation 2.1 into Equation 2.7 gives

$$\delta\phi(t) = \frac{\omega_0}{c} \int_0^L (\eta + h(t))\, \mathrm{d}x. \tag{2.8}$$

Since the Minkowski metric $\eta$ is time-independent and we are only interested in the time-dependent perturbations, we can drop the first term, which leaves us with

$$\delta\phi(t) = \frac{\omega_0}{c} \int_0^L h(t)\, \mathrm{d}x. \tag{2.9}$$

We would now like to evaluate the integral to determine how the phase of the light changes with the spacetime perturbation. We declare that the perturbation of spacetime has a frequency $f_{\mathrm{gw}}$, and that for the purposes of this derivation we are dealing in the long wavelength limit where the light travel time is much smaller than the timescale of variations in the metric,

$$\frac{L}{c} \ll \frac{1}{f_{\mathrm{gw}}}. \tag{2.10}$$

Under these assumptions, we have

$$\delta\phi(t) = \frac{\omega_0 L}{c} h(t). \tag{2.11}$$

This tells us that if we allow some light of frequency $\omega_0$ travel a distance $L$, in the presence of a gravitational wave that is described by $h(t)$, we will see the phase of the light change by $\delta\phi$ as given in Equation 2.11. To maintain the condition defined in Equation 2.10, we require that the length of time an average photon is in the detector be smaller than the period of a gravitational wave. A gravitational wave with $f_{\mathrm{gw}}$ of 5 kHz will have a period of 200 $\mu$s. A photon travelling 4 km to an end mirror in a simple Michelson interferometer and reflecting back will take about 13 $\mu$s. If the condition in Equation 2.10 is not met, then the analysis problem of inferring $h(t)$ from $\phi(t)$ is slightly more complicated, since the integral in Equation 2.9 must be evaluated at all points.

## 2.2 Motivations for low frequency sensitivity improvement

Advanced LIGO's low frequency sensitivity limit will be approximately 10 Hz; why is it useful to suppress noise at lower frequencies? Several reasons exist. Section 2.2.1 through Section 2.2.3 will describe astrophysical events that will be easier to detect with improved low frequency sensitivity. Section 2.2.4 will discuss why reducing noise at frequencies below the LIGO low frequency limit can prevent nonlinear up-conversion of noise that significantly contaminates data for transient gravitational wave searches.

### 2.2.1 Early part of neutron star / neutron star inspiral

One of the primary methods of extracting gravitational waveforms from the interferometer's data stream is known as matched filtering. The idea is to compare a numerical version of Equation 2.2 with measured data. If the template filter matches the data well, then the parameters chosen for that template are correct for the astrophysical system. The complication is that a template must be created for every set of parameter values to be checked, and we must be able to differentiate which template matches the data with the highest fidelity.

To determine if a template matches the data, we evaluate

$$z(t) = 4 \int_0^\infty \frac{\tilde{s}(f)\, \tilde{h}^*_{\text{template}}(f)}{S_{\text{n}}(f)} e^{i2\pi f t}\, \mathrm{d}f, \tag{2.12}$$

where $S_{\text{n}}(f)$ is the one-sided strain noise power spectral density of the interferometer, $\tilde{h}^*_{\text{template}}(f)$ is the Fourier transform of a waveform template, and $\tilde{s}(f)$ is the calibrated output of the detector [14]. This implies that improving the detector sensitivity will increase the SNR of a candidate signal. In particular, if the sensitivity at low frequencies is improved, template waveforms can be much longer, which helps to constrain the astrophysical parameters that describe the signal in the data.

Another way to look at the improvement to be gained by lowering the low-frequency detection limit is to determine how long a gravitational wave will be measurable in our detectors. If we define a time to coalescence, $\tau$, as the time it will take for a binary system orbiting currently at a frequency $f_L$ to inspiral to the point of merger, we can write $\tau$ up to second post-Newtonian order [14] as

$$\tau = \frac{5}{256\rho} \frac{GM}{c^3} \left( v_L^{-8} + \left( \frac{743}{252} + \frac{11}{3}\rho \right) v_L^{-6} - \frac{32\pi}{5} v_L^{-5} + \left( \frac{3058673}{508032} + \frac{5429}{504}\rho + \frac{617}{72}\rho^2 \right) v_L^{-4} \right), \tag{2.13}$$

where

$$v_L = \left( \frac{GM}{c^3} f_L \right)^{1/3}. \tag{2.14}$$

If $f_L$ is the lowest frequency that can be confidently detected, $\tau$ is the amount of time the signal will be measurable (assuming that the merger happens within the bandwidth of the interferometer – see Section 2.3). For the Advanced LIGO detection limit of $f_L \sim 10\,\text{Hz}$, the inspiral of a pair of $1.4\,\text{M}_\odot$ neutron stars will be present in the detector band for the next 258 seconds. In contrast, Initial LIGO with a low frequency detection limit of $f_L \sim 40\,\text{Hz}$ would only see the same inspiral for the final 6.6 seconds before merger.

This increase in the time to coalescence is particularly useful in cases where we would like to perform an electromagnetic counterpart follow-up. The sooner we can be confident of a detection candidate, the sooner we can ask electromagnetic observatories to look at a localized point in the sky. Low latency electromagnetic followup is especially useful for sources such as gamma ray bursts, supernova, etc.

### 2.2.2 Pulsars

Another interesting source for Advanced LIGO is pulsars. Pulsars emit beamed radiation, and as they spin the beam is periodically pointed at Earth. The gravitational wave from a pulsar will be a simple sine wave, although the signal detected will be modulated by the motion of the earth. At this time, millisecond pulsars are the primary class of pulsars to which LIGO is sensitive. They are often found in binary systems, where the partner star helps to spin the pulsar up to a frequency detectable by LIGO. The Hulse-Taylor pulsar in system PSR1913+16 [17] was the first pulsar to be found in a binary system.

While there are quite a few millisecond pulsars, there are many more that rotate at a slower rate, with expected gravitational wave frequencies just below the LIGO band. Improving LIGO's sensitivity at low frequencies will allow us to detect a large quantity of gravitational waves from pulsars. Since pulsars are an "always on" continuous source of gravitational waves, they can be very useful for testing general relativity. If the measured gravitational wave (after correcting for the amplitude modulation due to the detectors' antenna patter and the fact that the detectors are moving through space) do not match the expected waveform, we can investigate the deviations as potentially indicating modifications required to general relativity.

### 2.2.3 Intermediate mass black holes

Very few intermediate mass black holes (with masses in the range 50-$10^6$ solar masses) have been observed, although there is evidence for one in Omega Centauri [18, 19]. Neutron stars have an upper limit on their mass of about 3 $M_\odot$ [20], so if we detect an inspiralling system with masses larger than this, they must be black holes. Stellar black holes are on the order of 5-50 $M_\odot$, whereas observed supermassive black holes are larger than about $10^6\,M_\odot$, so black holes in between this range are classified as intermediate mass black holes. Since supermassive black holes are thought to come from mergers of intermediate-size objects, intermediate mass black holes are expected to exist, despite the lack of observations.

Almost all intermediate mass black holes in a binary system will emit gravitational waves at frequencies less than 10 Hz. However, since a gravitational wave detection with component masses estimated to be on the order of an intermediate mass black hole would constitute a first observation of a theoretically postulated astrophysical object, it is interesting to attempt to make the detectors as sensitive as possible at these frequencies.

### 2.2.4 Better background for unmodeled burst searches

While LIGO is certainly searching for binary inspirals and continuous wave sources, we are also looking for unmodeled transient events. Such events could come from sources such as core-collapse supernovae in our galaxy [21], soft gamma ray repeaters, neutron star collapse into a rotating black hole [22], or others [23]. Glitches in the data stream from non-Gaussian events significantly contaminate the background for these searches.

Below some frequency, which for Advanced LIGO will be around 10 Hz, it is not feasible to improve the sensitivity enough to make detections at lower frequencies. However, it is useful to suppress noise at frequencies lower than this limit in order to reduce the effect of upconverted noise sources. Many non-linearities exist in physical systems, and they can be the cause of technical glitches in the data. A common example of this from Initial and Enhanced LIGO is Barkhausen noise in the ferromagnets used to actuate on LIGO's mirrors [24]. Barkhausen noise occurs when ferromagnets experience a change in the surrounding magnetic field [25]. Even though the external magnetic field changes may be at low frequencies, the sudden flipping of magnetic domains causes broad band high frequency changes in the amount of force between the magnet and a coil used to induce the external field. If, in the absence of control, the mirrors' displacement were reduced, a smaller amount of control force will be required to hold the mirrors in position. This reduces the required change in field surrounding the ferromagnets, which directly reduces the Barkhausen effect.

Another up-conversion mechanism that is difficult to avoid is the spurious coupling of scattered light back into the main interferometer beam. Because no surface is perfect, some amount of light is scattered from each mirror surface. Some of this light can reflect off of vibrating surfaces nearby and re-enter the readout path. The original mirror motion may have been at a low frequency, but it will be bilinearly upconverted, which creates a broadband increase in the noise, up to a cutoff frequency determined by the relative velocity between the mirror and the secondary surface. Reducing the overall motion of the mirrors will reduce the amount of changing scattered light in the system, and thus mitigate this broadband noise increase.

## 2.3 Motivations for mid-frequency sensitivity improvement

In the high-sensitivity region (which for Advanced LIGO will be roughly a few tens of Hz to a few kHz), the interferometers are expected to be sensitive enough to detect gravitational waves, however improving the sensitivity will enhance the confidence of any detection, and will also allow more accurate estimates of the astrophysical parameters of the source.

For any typical binary system of compact objects, the highest frequency gravitational waves will occur just before merger, when the objects are at their innermost stable circular orbit (ISCO). The frequency of this orbit is

$$f_{\mathrm{ISCO}} = \frac{c^3}{6\sqrt{6}\pi GM}, \tag{2.15}$$

where $M$ is the total mass of the system. For a pair of $1.4\,\mathrm{M}_\odot$ neutron stars, $f_{\mathrm{ISCO}} \sim 1.57\,\mathrm{kHz}$. However, for a pair of $10\,\mathrm{M}_\odot$ black holes, $f_{\mathrm{ISCO}} \sim 220\,\mathrm{Hz}$.

# Chapter 3

# LIGO interferometers

As discussed in Section 2.1, we can use laser interferometers as transducers to convert gravitational wave displacement to phase information encoded in a laser beam. We then use various readout techniques to extract this phase information. This chapter focuses on the Laser Interferometer Gravitational wave Observatory (LIGO) methodology for these measurements, building up from the core Michelson interferometer to the details of the design of Advanced LIGO. Chapter 4 describes the LIGO 40 m Prototype Laboratory on the Caltech campus, which is used as a proof-of-principle system for much of the work described later in this thesis.

## 3.1 Measuring gravitational waves

In principle, we could measure a gravitational wave's effect on two test masses by monitoring the change in distance between them. However, it is very difficult to manufacture such a system with low enough noise as to actually measure gravitational waves. We use several techniques to reduce the effect of technical noise sources.

### 3.1.1 Michelson interferometers

One of the first techniques that we use is a differential Michelson interferometer. The laser beam (which will have a large amount of noise relative to the quantities that we want to measure) is incident on a beam splitter (BS) that allows half of the light to transmit and reflects the other half. The light travels in each direction some distance $L_{[x,y]}$, and is entirely reflected by a mirror called an end test mass (ETM [X,Y]) back to the beam splitter. The notation of two variables in square brackets indicates equivalent quantities for each of the two directions. The light is recombined at the beam splitter and sent to the photodetector (PD). This setup is sketched in Figure 3.1. Note that in practice, all of the optics will have losses and imperfections, so the beam splitter will not be a perfect 50/50, and the end mirrors will not reflect 100 % of the light.

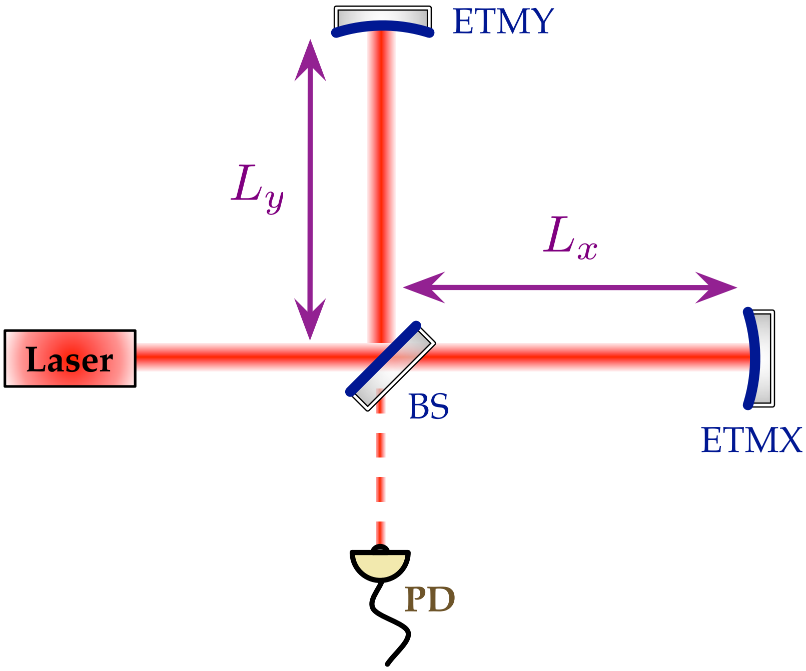

Figure 3.1: Cartoon of Michelson interferometer.

The light sent down a single direction (or "arm") will have to travel distance $2L_{[x,y]}$, and will accumulate some amount of phase $\phi_{[x,y]}$ during the round trip from the beam splitter to the end mirror and back. The original laser beam will also have some noise due to the laser itself and any optical components before the beam splitter, which we can denote by $\phi_{\text{noise}}$. Thus, the total phase from each arm just before recombination at the beam splitter will be

$$\phi_{[x,y]\text{ total}} = \phi_{[x,y]} + \phi_{\text{noise}}. \tag{3.1}$$

Gravitational waves are quadrupolar, so for a particular orientation relative to the interferometer they will change the lengths of the arms differentially, stretching one arm while squeezing the other. Half a cycle later the stretched and squeezed arms will be switched. At all but a few special moments, the path lengths of the Michelson arms will be unequal,

$$L_x \neq L_y, \tag{3.2}$$

which implies that the accumulated phase $\phi$ in the arms will also be unequal,

$$\phi_x \neq \phi_y. \tag{3.3}$$

As derived in many introductory textbooks, such as [26], the power measurable at the antisymmetric photodiode (PD in Figure 3.1) will only be related to the difference in phase accumulation between the arms

$$P^{\text{PD}} \propto sin^2(\Delta\phi). \tag{3.4}$$

Since the $\phi_{\text{noise}}$ is common to both arms, it is cancelled, leaving us only with the quantities of interest, $\phi_{[x,y]}$, which can be directly related to the change in optical path length of the arms. This elimination of common mode noises is the prime reason for choosing to use a Michelson interferometer. Note that while the gravitational waves can also make common mode length changes, we are sacrificing our ability to measure this orientation to gain the significant benefit of common mode noise rejection.

In practice, we will use a control servo to force the differential length of the arms such that the power detected at the photodiode is near zero. The force required to hold the mirrors in this state is directly related to how much the mirrors would have moved in the absence of control.

### 3.1.2 Fabry-Pérot arm cavities

Since the gravitational wave signal is related to the strain $\frac{\delta L}{L}$ of the distance $L$ between the test masses rather than just the absolute difference, we directly increase the sensitivity of the detectors by increasing the length of the arms. However, it is cost prohibitive to make extremely long arms. Instead, we increase the optical path length without increasing the physical distance between the beam splitter and the end mirrors.

We do this by inserting an extra mirror in each arm between the beam splitter and the end mirror. In Figure 3.2 we label the new optics input test masses (ITM [X,Y]). Here, the distance between the new mirrors and the previous end mirrors is denoted $L_{[x,y]}$, while the shorter distance between the beam splitter and the new mirrors is denoted $l_{[x,y]}$.

The addition of these extra mirrors creates Fabry-Pérot optical cavities in the arms of the Michelson. The Fabry-Pérot cavities increase the optical length of the arms by a factor

$$\frac{L_{\text{FP}}}{L_{\text{Mich}}} = 2\mathcal{F}, \tag{3.5}$$

where $\mathcal{F}$ is defined as the finesse of the Fabry-Pérot cavity,

$$\mathcal{F} = \frac{f_{\text{FSR}}}{2f_{\text{cav pole}}} \approx \frac{\pi}{2}\frac{\sqrt{(r_i r_e)}}{1 - r_i r_e}. \tag{3.6}$$

In Equation 3.6, $f_{\text{FSR}}$ is the frequency spacing between transmission peaks of the cavity, known as the free spectral range, and $f_{\text{cav pole}}$ is the cavity pole frequency, which is defined as the half-width of the transmission peak, at half of the maximum buildup. The free spectral range of the cavity

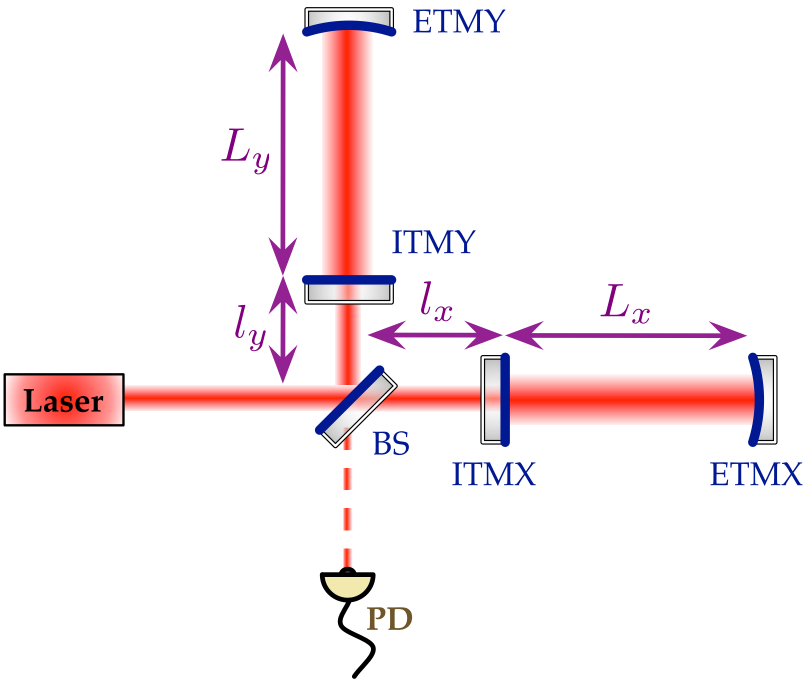

Figure 3.2: Cartoon of Michelson interferometer with Fabry-Pérot arms.

depends on the geometry of the system,

$$f_{\mathrm{FSR}} = \frac{c}{2L}, \tag{3.7}$$

where $c$ is the speed of light and $L$ is the distance between the cavity mirrors. The linewidth of the cavity (which is twice the cavity pole) depends on the amplitude reflectivities ($r_{[i,e]}$) of the input and end mirrors used,

$$f_{\mathrm{cav\ pole}} = \frac{c}{4L}\frac{1 - r_i r_e}{\pi\sqrt{r_i r_e}}. \tag{3.8}$$

This shows that we can dramatically increase the effective length of the detector arms by making the Fabry-Pérot cavity mirrors nearly perfectly reflective. Note that Equation 3.5 is based on the increased storage time of photons in the Fabry-Pérot cavities, which is an average time. Some photons may escape after only a few round trips, while others may remain in the cavity for much longer. This means that the increased optical length of Equation 3.5 is an average over all photons incident on the cavity.

As might be expected, there are some drawbacks to using this technique, so any interferometer design must balance these against the benefit of the increased path length. One of the primary requirements is that the arm cavities must be held on (or very close to) resonance. The light reflected from the cavity back toward the beam splitter is only linearly related to the length of the cavity when it is close to resonance. To hold the cavities on resonance, we use the canonical

Pound-Drever-Hall technique, which requires the addition of phase modulated radio frequency (RF) sidebands to the incident light [27]. See, for example, [28] or Appendix B, for a derivation of this technique for a simple 2-mirror cavity.

Instead of individually controlling the 2 Fabry-Pérot arm cavities, we chose to make a basis change and control the arms in differential and common modes. From Figure 3.2, we define the differential arm length (DARM) as

$$L_- = \frac{L_x - L_y}{2} \tag{3.9}$$

and the common arm length (CARM) as

$$L_+ = \frac{L_x + L_y}{2}. \tag{3.10}$$

Similarly, the short Michelson has two similar degrees of freedom,

$$l_- = \frac{l_x - l_y}{2} \tag{3.11}$$

and

$$l_+ = \frac{l_x + l_y}{2}, \tag{3.12}$$

where $l_-$ is called the Michelson (MICH) degree of freedom, and $l_+$ is uncontrolled.

We put a static offset in the MICH length $l_-$, to take advantage of the Schnupp modulation technique [29]. This allows some amount of light (particularly the RF sidebands) to always be present at the antisymmetric port of the interferometer. If the arm cavity lengths change differentially (for example, when a gravitational wave passes through the detector), the main carrier light that is reflected from the arms will have a slightly different phase than if the cavities were perfectly resonant. This phase modulation of the carrier light can be thought of as the addition of "signal sidebands" on the light, analogous to the original RF sidebands.

The carrier field (including the signal sidebands) will exit the interferometer through the antisymmetric port, and beat against the static RF sideband fields. As with the more simple 2-mirror cavity case, this beat creates amplitude modulations on the light, which are detected by the photodiode. This signal is then demodulated at the original RF frequency, which results in a signal with components only at the signal sideband frequencies [30]. In this way, we are able to hold the mirrors stationary and infer what the length change due to a gravitational wave signal would have been in the absence of control (assuming no noise sources). In contrast to the technique that will be described in Section 3.2.4, this is a heterodyne technique.

Another drawback of the addition of Fabry-Pérot cavities to the interferometer is the decreased detection bandwidth. The storage time of the cavities puts a high frequency limit on the signal sidebands that the detector is sensitive to. Derivations such as that shown in Chapter 2 of [31] give the frequency response of an interferometer with Fabry-Pérot arms as

$$G(\omega) = \frac{2L\Omega}{c}\left(\frac{\sin(\frac{L\omega}{c})}{\frac{L\omega}{c}}\right)\left(\frac{t_i^2 r_e}{1 - r_i r_e}\right)\left(\frac{e^{\frac{-iL\omega}{c}}}{1 - r_i r_e e^{\frac{-i2L\omega}{c}}}\right), \tag{3.13}$$

where $\Omega$ is the angular frequency of the laser.

When the interferometer includes Fabry-Pérot arm cavities, we typically use the difference in cavity lengths as our gravitational readout signal.

Ideally, we would like to measure just the differential arm length $L_-$ at the photodiode (PD in Figure 3.2), but any motion in the short Michelson's $l_-$ will also show up at this photodiode. We can calculate this coupling by determining what the electric field will be at the photodiode, and then taking the derivative relative to either $l_-$ or $L_-$. We compare these ratios to determine the relative importance of a differential arm length change versus a differential Michelson length change.

Explicitly, we will use the calculation method described in [32]. Figure 3.3 is a copy of Figure 3.2, but with the electric field at specific points in the interferometer labeled. Using the matrix formalism for determining the fields everywhere in the matrix, we must write down the transfer matrix $M_{(\text{from, to})}$, where "from" and "to" are the indicies indicated in Figure 3.3. For each element in this matrix, we need only write down how the electric field propagates from one point to another, in the absence of all other fields. For simplicity, we assume that both input mirrors are identical, and also that both end mirrors are identical.

We define the amplitude reflectivity $r$ as

$$r = \sqrt{R}, \tag{3.14}$$

where $R$ is the fraction of power reflected from a mirror. Similarly, we define the amplitude transmission $t$ as

$$t = \sqrt{T}, \tag{3.15}$$

where $T$ is the fraction of power transmitted through the optic. Both $r$ and $t$ can be complex numbers, and by convention we declare that the electric field will pick up a negative sign when reflecting off of a highly reflective surface. In general, some amount of light is lost when a beam is incident on a mirror; however, for simplicity here we assume lossless mirrors. To conserve the total energy, we note that

$$R^2 + T^2 = 1. \tag{3.16}$$

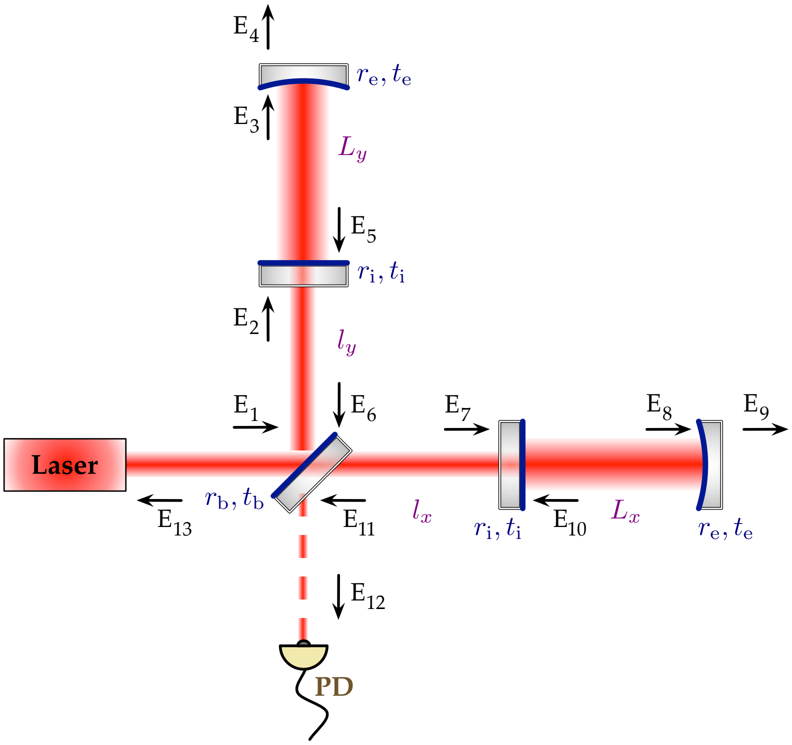

Figure 3.3: Cartoon of Fabry-Pérot Michelson with notations for electric field at various points in the interferometer, for use in calculating the coupling between length changes at different points in the interferometer. $E$ indicates the electric field at a given point, while $r$ and $t$ indicate the amplitude reflectivity and transmissivity of the mirrors. Subscript $b$ indicates the beam splitter, subscript $i$ indicates one of the (identical) input test masses and subscript $e$ indicates one of the (identical) end test masses.

We also note that light that travels a distance $L$ will accumulate phase

$$\phi = kL, \tag{3.17}$$

where $k$ is the wave number of the laser light,

$$k = \frac{2\pi}{\lambda}. \tag{3.18}$$

For the case shown in Figure 3.3, the transfer matrix will be

$$M = \begin{pmatrix}
0 & -r_b e^{-ikl_y} & 0 & 0 & 0 & 0 & t_b e^{-ikl_x} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & t_i e^{-ikL_y} & 0 & 0 & r_i e^{-ikl_y} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & t_e & -r_e e^{-ikL_y} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -r_i e^{-ikL_y} & 0 & 0 & t_i e^{-ikl_y} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & t_b & -r_b \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & t_i e^{-ikL_x} & 0 & 0 & r_i e^{-ikl_x} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & t_e & -r_e e^{-ikL_x} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & -r_i e^{-ikL_x} & 0 & 0 & t_i e^{-ikl_x} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & r_b & t_b \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}. \tag{3.19}$$

We assume that the only electric field input is from the laser at $E_1$, so we construct a vector of initial electric fields

$$E_{\text{input}} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \tag{3.20}$$

The steady-state resulting electric field everywhere will be given by

$$E_{\text{steady-state}} = \left(\mathbb{1} - M^T\right)^{-1} \cdot E_{\text{input}}, \tag{3.21}$$

simplified by changing the basis from $(x, y)$ to $(+, -)$ by utilizing Equation 3.9 through Equation 3.12. We are actually only interested in looking at the field at the detection port, which we have defined as element number 12 in Figure 3.3. To determine the response of the field at $E_{12} \equiv E^{PD}_{\text{steady-state}}$ due to length changes, we will take the derivative with respect to $L_-$ and also with respect to $l_-$. Taking the ratio of these sensitivities gives us the coupling factor between MICH and DARM,

$$\text{Coupling} = \frac{\text{MICH Sensitivity}}{\text{DARM sensitivity}} = \frac{\frac{\partial}{\partial l_-} E^{\text{PD}}_{\text{steady state}}}{\frac{\partial}{\partial L_-} E^{\text{PD}}_{\text{steady state}}} = \frac{(1 - r_i r_e)(r_i - r_e)}{r_e \left(1 - r_i^2\right)}. \tag{3.22}$$

For input test masses of transmission $T_i = 1.4\%$ and end test masses of transmission $T_e = 15\,\text{ppm}$, this coupling ratio is approximately $3.5 \times 10^{-3}$.

Note that the Michelson has a flat frequency response, but the DARM degree of freedom will have a cavity pole ($f_{\text{cav pole}}$), so the frequency-independent coupling factor in Equation 3.22 is valid from DC up to the cavity pole frequency.

As with the simple Michelson, we do not use the CARM degree of freedom to measure gravitational waves. However, since we assume that the long arm cavities provide a stable length reference, we can use the CARM error signal to feedback to and stabilize the laser frequency.

### 3.1.3 Power recycling

To further improve the sensitivity of the interferometer to gravitational waves, LIGO also uses "power recycling". An extra mirror is placed between the laser and the beam splitter at the symmetric port, as shown in Figure 3.4. We refer to this power recycling mirror as the PRM, and define the length of this cavity as

$$l_{\text{PRCL}} = l_p + l_+, \tag{3.23}$$

where $l_+$ is defined in Equation 3.12 and $l_p$ is the distance between the PRM and the beam splitter. The idea behind power recycling is that any true differential-mode gravitational wave signal will exit through the antisymmetric port to the photodetector, and all other light will be reflected back to the laser at the symmetric port. Instead of dumping and losing this light, we recycle it by sending it back into the interferometer. This has the effect of illuminating the beam splitter with more light, and is roughly equivalent to increasing the laser power.

If, at high frequencies, the sensitivity of the interferometer is limited by photon counting shot noise, power recycling will improve the signal-to-noise ratio. The increase in light power in the detector implies that, for the same gravitational wave signal, more photons will arrive at the photodetector. Averaging over more photons decreases the overall noise level due to shot noise.

Obviously we cannot change the interferometer's topology without altering the effect of signals at the various locations in the detector. However since power recycling affects the entire interferometer in the same way (i.e., it is a common-mode change), it will only reduce the bandwidth of the symmetric mode, leaving the antisymmetric bandwidth unchanged, as can be seen in Figure 3.5. As discussed in Section 3.1.1, we chose to only look at the gravitational wave signal at the antisym-

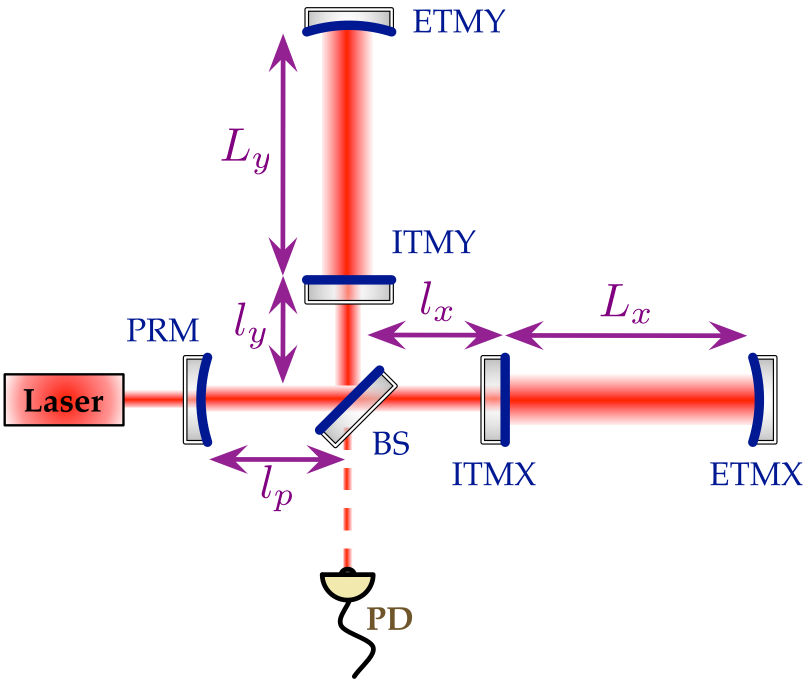

Figure 3.4: Cartoon of power recycled Fabry-Pérot Michelson interferometer.

metric port, so this loss of common mode bandwidth does not affect the ability to detect signals at higher frequencies. Overall, the addition of the power recycling mirror allows us to reduce the sensitivity to laser and other common mode noises while maintaining the same sensitivity to gravitational waves (and any differential mode noises).

Due to the non-zero difference in the path between the beam splitter and the two input test masses ($l_x \neq l_y$ and $r_{\mathrm{ITMX}} \neq r_{\mathrm{ITMY}}$), there will be some amount of signal that appears at the photodiode as a result of PRM motion. The DARM degree of freedom will see this coupling as a noise, so it is beneficial to understand the level of the noise. We can use the same method described in Section 3.1.1; however, as the calculation becomes more complicated it becomes useful to utilize programs such as Optickle [33] which will also include the frequency response of the system. For the parameters that are listed in Table 4.3 and Table 4.4 for the Advanced LIGO design, the response to either DARM motion or PRCL motion at the photodiode is shown in Figure 3.6.

### 3.1.4 Initial and Enhanced LIGO noise

Overall, these things made Initial LIGO meet design sensitivity, as proposed in 1989 [34]. Several enhancements within the then-current infrastructure improved the sensitivity by a factor of about 2 at high frequencies. These upgrades included higher laser power, better thermal adaptive optics, and an optical cavity to filter the output beam.

The laser power was increased in order to further improve the high frequency sensitivity of

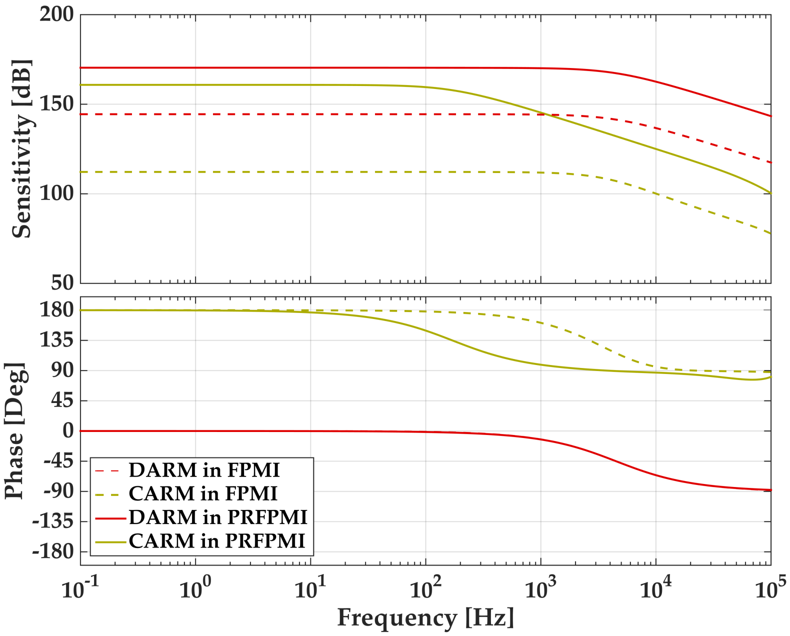

Figure 3.5: Comparison of the common mode (CARM) and differential mode (DARM) transfer functions between ETM motion and the detection photodiode in the Fabry-Pérot Michelson (FPMI) configuration versus the power recycled Fabry-Pérot Michelson (PRFPMI) configuration. The common mode bandwidth is reduced, while the differential mode bandwidth is unchanged, although the overall sensitivity is increased for both. Note that these traces reflect 40 m Lab parameters, although the relative change from FPMI to PRFPMI is similar for Advanced LIGO.

the interferometer, which is limited by the photon counting noise as described in Section 3.1.3. Since our mirrors are not perfectly reflective, this increase in laser power (by about a factor of 3.5 in design, although only a factor of 2 in practice) required an improvement in the thermal compensation system. The mirrors absorb some fraction of the energy in the laser beam, which heats up the mirrors. This causes an alteration in the physical shape of the mirrors, as well other optical properties of the mirrors [35], which changes the phase accumulated by the light as it circulates in the arm cavities. These effects are mitigated with the Thermal Compensation System (TCS) [36].

The final change between the Initial and Enhanced generations of the LIGO interferometers was the move from RF sensing to DC sensing of the power at the antisymmetric port. This required the addition of an optical cavity at the output port known as the Output Mode Cleaner (OMC), which will be described in more detail in Section 3.2.4.

The overall sensitivity of the LIGO detectors, for both the Initial and Enhanced generations is shown in Figure 3.7, as compared to the original design sensitivity. The Initial LIGO sensitivity is represented by the curve from the 5th Science Run (S5), taken in 2006. The Enhanced LIGO sensitivity is represented by the curve taken in 2010 from the 6th Science Run.

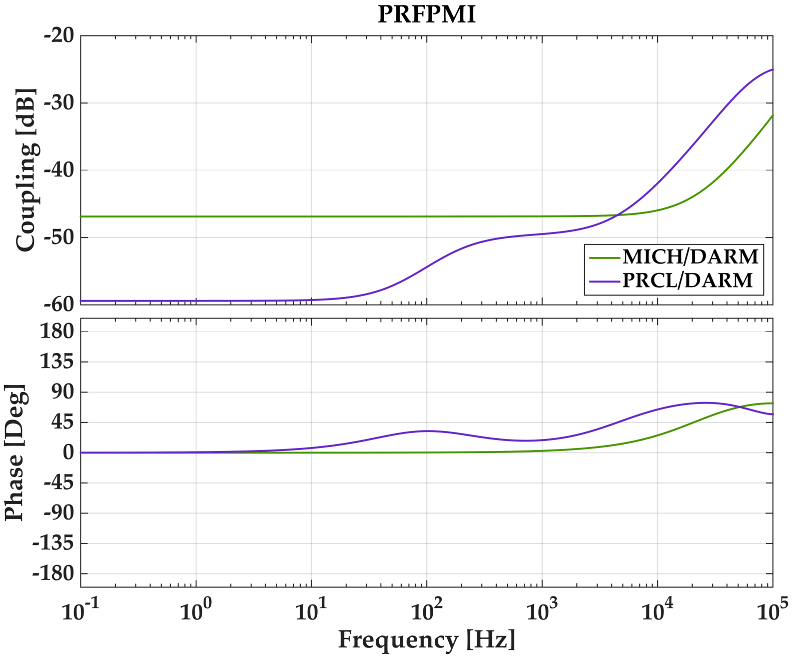

Figure 3.6: Coupling of PRCL and MICH degrees of freedom to DARM, in the PRFPMI configuration. Note that these traces reflect 40 m Lab parameters, although the overall shape is similar for Advanced LIGO.

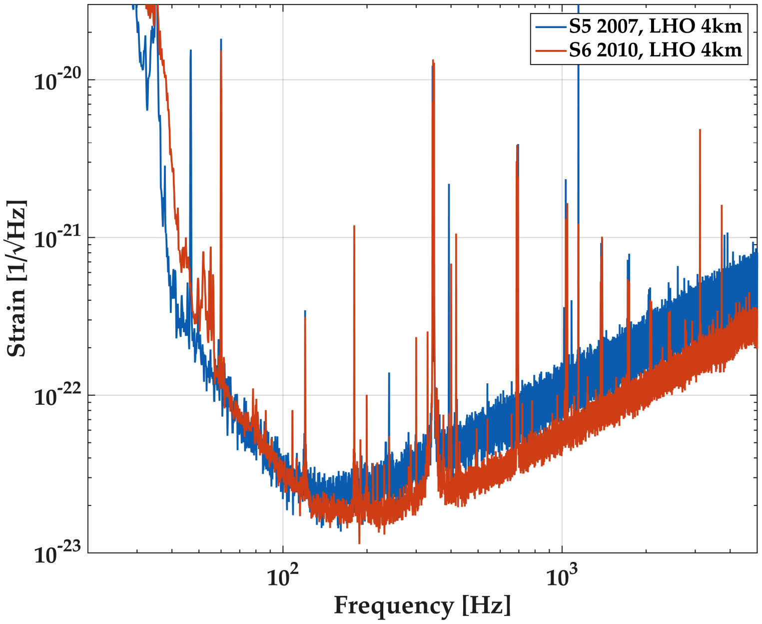

Figure 3.7: Measured sensitivity of Initial LIGO (“S5”) and Enhanced LIGO (“S6”). Data from [37].

## 3.2 Increasing measurement sensitivity: Advanced LIGO

After LIGO's sixth science run was completed in 2010, the interferometers were decommissioned for the installation of new Advanced LIGO hardware. Advanced LIGO is comprised of a set of improvements that are implemented in the original LIGO infrastructure, and are expected to provide a broad-band sensitivity improvement of about a factor of 10 compared to Initial LIGO [38].

### 3.2.1 Signal recycling

One of the most dramatic changes between Initial / Enhanced LIGO and Advanced LIGO is the addition of a so-called signal recycling mirror (SRM). This auxiliary mirror is placed at the detection port of the interferometer, as shown in Figure 3.8. In the case where a signal recycling cavity (SRC) is added to an interferometer that already includes power recycling, the interferometers are referred to as "dual recycled" Fabry-Pérot Michelson interferometers. The length of this new cavity is defined as

$$l_{\mathrm{SRCL}} = l_s + l_+, \tag{3.24}$$

where $l_s$ is the distance between the SRM and the beam splitter, and $l_+$ is as defined in Equation 3.12.

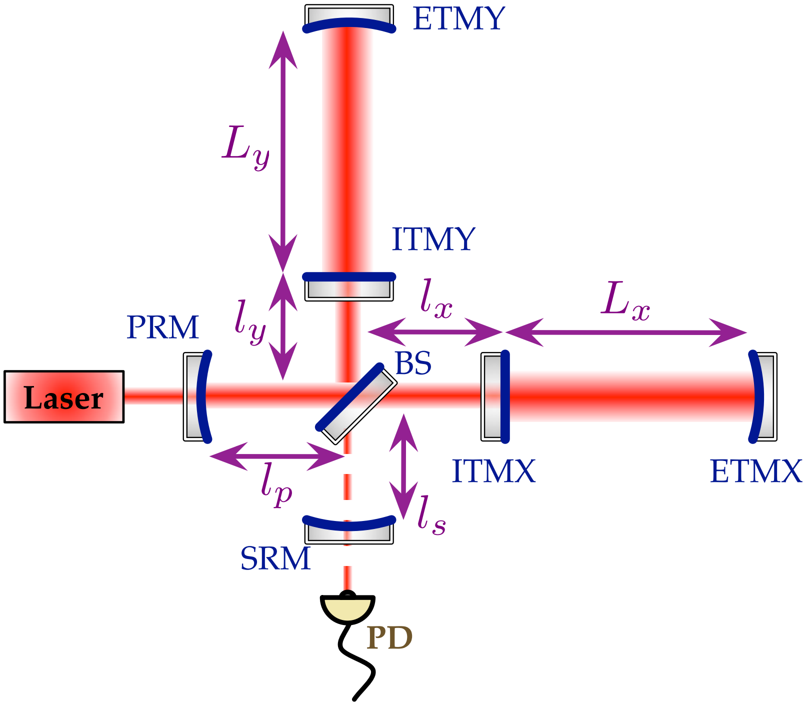

Figure 3.8: Cartoon of dual (power and signal) recycled Fabry-Pérot Michelson interferometer.

This signal recycling cavity can be utilized in one of two ways, either for true signal recycling, or for resonant sideband extraction. In the signal recycling case, the SRC is controlled such that the main carrier light is anti-resonant. This has the effect of amplifying the low frequency components of the gravitational wave signal, but comes at the cost of reducing the detector bandwidth. LIGO does not use this technique.

Rather, we utilize resonant sideband extraction (RSE), where the SRC is resonant for the carrier light. This reduces slightly the low frequency detection gain, but increases the detector's bandwidth. Figure 3.9 shows the effect of adding a resonant signal recycling cavity to an already power recycled interferometer. A key benefit of this technique is that this bandwidth improvement does not require increased light power incident on the mirrors, which reduces the need for thermal compensation. See [39] for a more thorough discussion of the RSE technique.

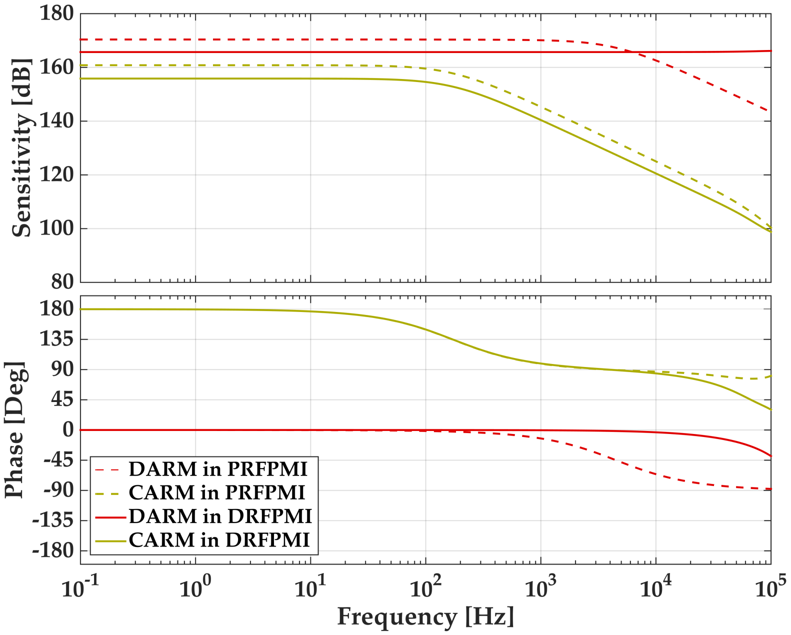

Figure 3.9: Comparison of the common mode (CARM) and differential mode (DARM) transfer functions between ETM motion and the detection photodiode in the power recycled Fabry-Pérot Michelson (PRFPMI) configuration versus the dual recycled Fabry-Pérot Michelson configuration (DRFPMI). The common mode bandwidth is unchanged, although the differential mode bandwidth is increased. Note that these traces reflect 40 m Lab parameters, although the relative change from PRFPMI to DRFPMI is similar for Advanced LIGO.

Regardless of whether the SRC is resonant or anti-resonant for the carrier light, the addition of an extra optical cavity requires that we add an extra set of radio frequency sidebands for control. Also, this auxiliary length degree of freedom will couple to the gravitational wave detection channel, much like Michelson and power recycling coupling described in Section 3.1.2 and Section 3.1.3. Figure 3.10 shows the coupling factors between all of the short length degrees of freedom to the differential arm, in the dual recycled RSE configuration. As in Figure 3.6, these are the ratios of sensitivity at the detection port to change in the length of one of the short degrees of freedom to the differential arm length change.

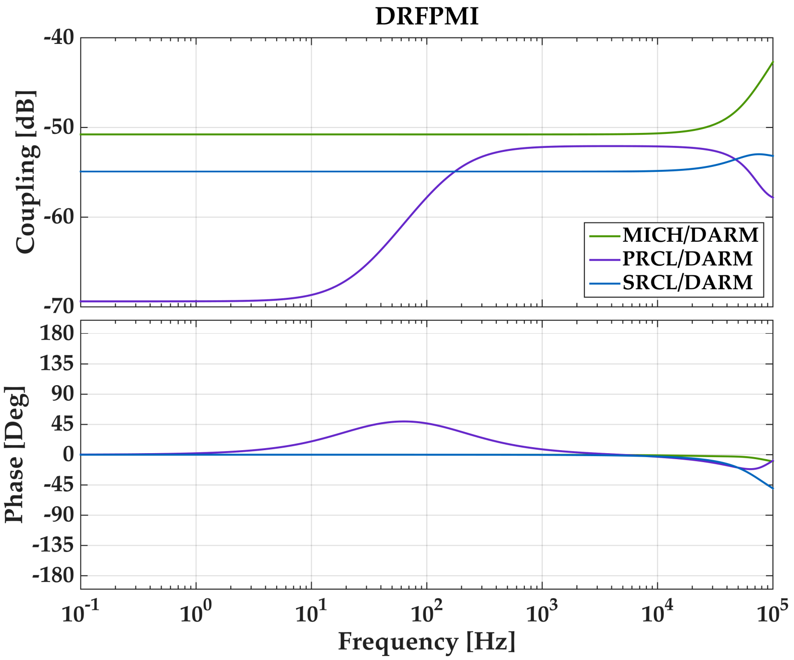

Figure 3.10: Coupling of PRCL, SRCL and MICH degrees of freedom to DARM, in the DRFPMI configuration. Note that these traces reflect 40 m Lab parameters, although the overall shape is similar for Advanced LIGO.

### 3.2.2 Improved seismic isolation

One of the most significant infrastructure changes between Initial LIGO and Advanced LIGO is the improvement in low frequency seismic isolation. The limiting noise at low frequency for ground based gravitational wave detectors has always been ground motion, so improvement in the isolation will be directly related to lowering the low frequency detection limit.

While the Initial LIGO mirrors were all suspended with single stage pendula, Advanced LIGO's test mass mirrors are all quadruple stage pendula, one of which is shown on the left in Figure 3.11. Some auxiliary mirrors, such as the beam splitter, will only use triple stage pendula, since their motion does not couple as strongly to the detection port. Since each pendulum stage increases the isolation by a factor of $\frac{1}{f^2}$ above the resonance frequency, the additional stages significantly improves the sensitivity of the detector. The right side of Figure 3.11 illustrates this improvement for a set of simplified multi-stage pendula. Note that at the resonant frequencies of the additional stages some isolation is lost, although at 10 Hz the quadruple pendulum offers a factor of $10^6$ more isolation than the single pendulum. While it would be conceptually simpler to increase the length of a single pendulum rather than adding extra stages, a single pendulum would need to be a factor of $10^{12}$ longer to achieve the same isolation improvement, which is clearly impractical.

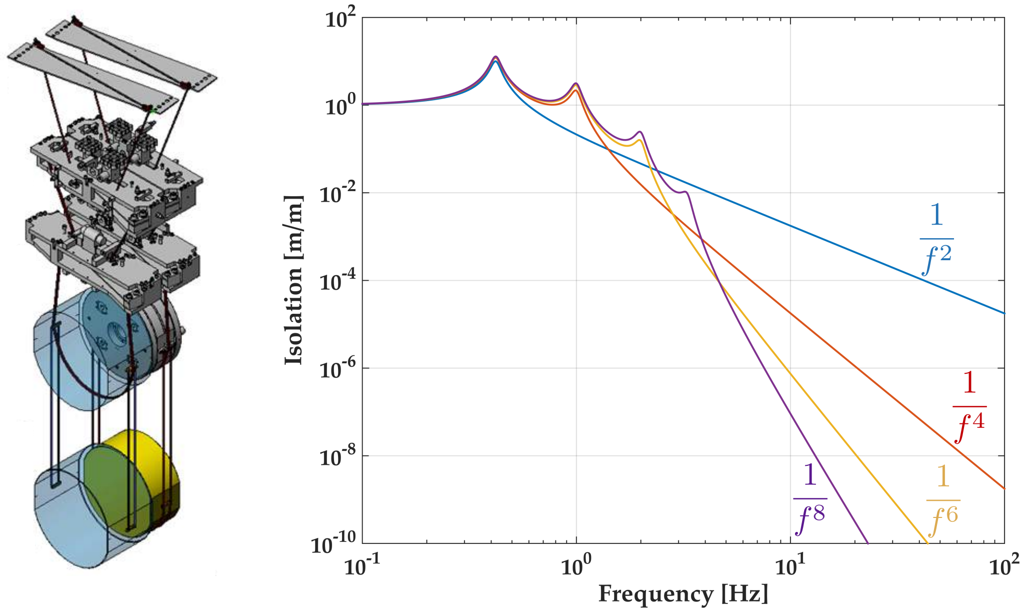

Figure 3.11: Left: Diagram of Advanced LIGO quadruple pendulum and reaction chain (Figure 1 from [40]). Right: Comparison of frequency response of example single stage, double stage, triple stage and quadruple stage pendula. The resonant frequencies for these example pendula have been chosen to be the same. Note that each additional pendulum stage changes the slope of the response at high frequencies (above 10 Hz) by an extra factor of $\frac{1}{f^2}$.

In addition to the passive seismic isolation provided by pendula, Advanced LIGO also incorporates active isolation. Initial LIGO in-vacuum optical tables were comprised of stacks of rubber and metal, which provided passive isolation. However, Advanced LIGO optical tables are supported by internal seismic isolation (ISI) tables that are loaded with sensors and actuators. These give the tables the ability to provide local damping beyond what the passive tables could achieve. As a prototype, one of these active tables was included in Enhanced LIGO for the output optics, but for Advanced LIGO every table is now equipped with the suite of sensors and actuators (see, for example, J. Kissel's thesis [41]).

In addition, the out of vacuum supports for the optical tables have also been upgraded. Initially, each table was supported by a set of piezo actuators; however, for Advanced LIGO these have been changed to longer-range hydraulic actuators [42]. Note that the ground motion at the Livingston, Louisiana site was large enough that these hydraulic external pre-isolators (HEPI) were installed during Initial LIGO, however the Hanford, Washington site just installed them for Advanced LIGO.

This set of active seismic isolation equipment primarily isolates the optics between 0.1 – 10 Hz, which adds significantly to the overall immunity to ground noise.

### 3.2.3 Increased laser power

The fundamental quantum noises that most closely limit the LIGO sensitivity curve are both related to the power of the main laser (more specifically, the power stored in the Fabry-Pérot arm cavities, which is related to the main laser power). At higher frequencies, the quantum limit is dominated by photon shot noise, which is a consequence of the discrete nature of photons. The gravitational wave signal is directly proportional to $N_{\text{photons}}$, but because the photons follow Poisson statistics, the uncertainty will go as $\sqrt{N_{\text{photons}}}$, so the overall signal to noise ratio at frequencies dominated by shot noise will improve with $\sqrt{N_{\text{photons}}}$.

A drawback to increasing the laser power is that the force due to radiation pressure exerted by the photons bouncing off the mirrors scales directly with the laser power,

$$F_{\text{rad press}} = \frac{2P}{c}, \tag{3.25}$$

where $P$ is the power of the laser incident on the mirror. The fluctuation of photon number (so, power incident) with time creates a frequency-dependent noise spectrum. This is particularly significant for angular control of the cavities, where the torque due to radiation pressure is

$$\tau_{\text{rad press}} = \frac{2Pd}{c}, \tag{3.26}$$

where $d$ is the distance that the beam spot is mis-centered on the optic. This is counter-acted by the natural pendulum torque,

$$\tau_{\text{pend}} = \alpha I \omega_{\text{pend}}^2, \tag{3.27}$$

where $\alpha$ is the angle of the pendulum with respect to its resting position, $I$ here is the moment of inertia of the mirror and $\omega_{\text{pend}}$ is the angular resonance frequency of the angular mode of the pendulum. Large ratios of $\frac{\tau_{\text{rad press}}}{\tau_{\text{pend}}}$ can cause large changes in the resonant frequency $\omega_{\text{pend}}$, which can make angular control of the interferometer more difficult [43].

Increasing the laser power also affects the radius of curvature and other optical properties of the mirrors, due to the finite absorption of energy of the mirrors [35].

Figure 3.12 shows the trade-off between radiation pressure noise at low frequencies and shot noise at higher frequencies as the laser power is increased (dashed traces). We must make a choice to balance the positive and negative effects of increasing the laser power, and design all other elements of the detector (where possible) to introduce less technical noise than the fundamental quantum noise. Note that at low frequencies the overall detector sensitivity (solid traces) is not limited by the radiation pressure noise. As we will see in Section 3.2.5, seismic and other noise sources dominate the overall sensitivity below about 20 Hz.

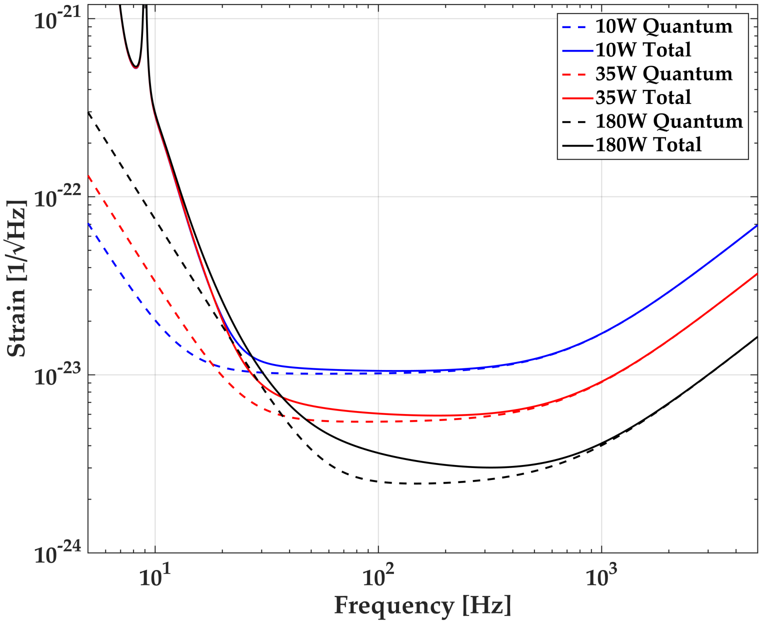

Figure 3.12: Advanced LIGO sensitivity as a function of input laser power, calculated using the LIGO Collaboration's Gravitational Wave Interferometer Noise Calculator, GWINC.

### 3.2.4 Output mode cleaner

Advanced LIGO also uses a significantly different method for reading out the gravitational wave information from the interferometer than Initial LIGO. While Initial LIGO utilized a heterodyne readout scheme (described in Section 3.1.2), Advanced LIGO will use a homodyne configuration. Although originally planned for Advanced LIGO, this system was implemented in Enhanced LIGO and proved to be very successful, as discussed in N. Smith [30] and T. Fricke's [44] theses.

Rather than comparing the effect that the gravitational wave has on the arm cavities with the non-resonant RF sidebands, we use the light reflected from the arm cavities as the reference. The DARM degree of freedom is detuned very slightly to allow some carrier light to leak to the antisymmetric port. When a gravitational wave passes through the detector, the amount of light reflected from the arm cavities will be modified. We detect this change in light power at the antisymmetric port directly, without having to demodulate the photodiode signal (which is required in the heterodyne detection scheme).

In addition to simplifying the readout electronics, this method has several other advantages. The carrier light is resonant in the arm cavities, and so is filtered by the CARM coupled cavity pole, which eliminates high frequency noise. This filtered light is used as the reference for comparison with the gravitational wave signal, so the detection will be lower noise by a factor of the cavity transfer function. Also, in order to get the maximum amount of signal at the photodiode, the reference beam and the signal beam must be perfectly co-aligned. With the heterodyne readout, the RF sidebands are not resonant in the arm cavities, and so can have slightly different alignment than the main carrier light. However, in the homodyne scheme, since we are comparing the arm cavity light with itself, the beams are inherently co-aligned.

While the advantages to the new readout scheme are significant, additional hardware is required to realize the full benefit of the method. Since we are just detecting the power differences at the photodiode, any light which is not the main carrier light directly causes extra noise. This can include extra RF sideband light that does not contain information about the gravitational wave signal, but does contribute to the shot noise level. It can also include higher order transverse mode light at the main carrier frequency, which is not resonant in the arms, so it again contributes only to the shot noise level.

To eliminate these noises, an extra filter cavity is added between the signal recycling mirror and the photodiode. This output mode cleaner (OMC) is resonant only for the main carrier light, and so removes the RF sidebands and any higher order transverse modes of the light. This additional cavity requires another set of length and angular controls, but is critical to the success of the homodyne readout scheme [30, 44].

### 3.2.5 Total Advanced LIGO noise

As part of making our design choices for all of the systems described in this chapter, we calculate what effect we expect each component to have on the final detector sensitivity. Figure 3.13 shows a plot of the result of these calculations made using the LIGO collaboration's Gravitational Wave Interferometer Noise Calculator (GWINC), for the dominant noise sources in the Advanced LIGO design. The purple trace is the quantum noise, as discussed in Section 3.2.3. The brown curve is the effect of direct seismic vibrations after passing through all of the stages of seismic isolation, including the suspensions described in Section 3.2.2. The green trace is the estimated noise contribution from Newtonian gravitational forces, which will be discussed in much more detail in Chapter 6. The blue and red curves are the effects of thermal noise on the suspension wires supporting the mirrors, and on the reflective optical coatings. Finally, the black trace is the total expected noise for Advanced LIGO.

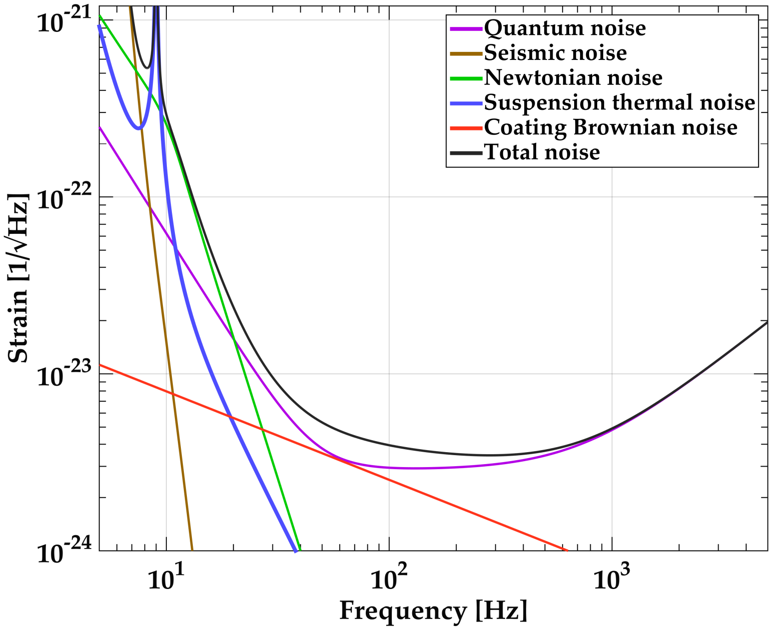

Figure 3.13: Advanced LIGO noise budget calculated by GWINC.

# Chapter 4

# The 40 m Prototype Interferometer

The 40 m interferometer was first constructed on the Caltech campus in the early 1980's. It was the first fully operational large-scale gravitational wave interferometer, although it never successfully detected gravitational waves. The 40 m underwent several upgrades throughout the years, often to improve the ability to prototype new technologies for future gravitational wave detectors.

In the 2000's, the 40 m added a signal recycling mirror, and became the first suspended interferometer to operate with dual recycling (signal and power recycling) in addition to Fabry-Pérot arms. Tabletop work had been pioneered previously, as described in K. Strain, et al. [45]. See R. Ward's thesis [39] for details on the configuration of the 40 m's detuned resonant sideband extraction interferometer.

In 2010, the 40 m was upgraded once again with dichroic optics, to test a new controls scheme, known as Auxiliary Length Stabilization for the arm cavities. With this configuration, we were able to show that this "green locking" method was promising for Advanced LIGO [46, 47]. Other changes made, to reflect the updated Advanced LIGO design, included the use of lower frequency RF modulations, and operating with broadband resonant sideband extraction. The lower modulation frequencies necessitated longer power and signal recycling cavities. Fitting these longer cavities into the existing vacuum system required the use of folding mirrors. The consequences of adding these folding mirrors will be discussed in Section 4.3.

Particularly after the challenges discovered with the addition of the signal recycling mirror, the 40 m Interferometer has focused on prototyping controls schemes for Advanced LIGO-like configurations. Our goal has been to commission and learn how to acquire full resonance and control for a dual recycled Fabry-Pérot Michelson interferometer, and to use the facilities as test bed for new controls techniques.

While the 40 m interferometer is designed to be Advanced LIGO-like in its controls and complexity thereof, it is not a full noise sensitivity prototype, and does not have all of the features of a full Advanced LIGO interferometer. Most obvious, the arm length of the 40 m is roughly 1 % that of the 4 km LIGO sites'. However, that is not the only significant difference. The 40 m has single loop optic suspensions, which sit on passive seismic isolation tables. The single loop suspensions are the same as several input and auxiliary optics in Initial LIGO [48]. The passive seismic isolation stacks were discussed in J. Giaime's thesis [49]. This combination of isolation implies that the 40 m is not as well seismically-isolated as Advanced LIGO.

The configuration of the 40 m interferometer is illustrated in Figure 4.1. The beam from the main laser (far left of diagram) is injected into the vacuum system (not depicted) and resonates in the input mode cleaner, the triangular ring cavity comprised of mirrors MC1, MC2, and MC3. Photodiodes MC REFL and WFS provide error signals, while MC TRANS provides information on the power buildup in the cavity. Light transmitted through the mode cleaner passes through a Faraday isolator, and then through mode matching and input steering optics (not shown) before approaching the power recycling mirror (PRM). Light reflected from the interferometer will be directed by the Faraday isolator to the reflection port of the interferometer where four resonant RF photodiodes are located. These photodiodes are labeled "REFL" and differentiated by their resonant frequency, $f_1 = 11\,\text{MHz}$, $f_2 = 55\,\text{MHz}$, $3f_1 = 33\,\text{MHz}$, and $3f_2 = 165\,\text{MHz}$. These photodiodes will provide error signals for many of the length degrees of freedom for the interferometer.

After the mode matching and input steering optics, the main beam passes through the PRM, reflects off of two folding mirrors (PR2 and PR3) and is incident on the main beam splitter (BS). Some amount of light leaks through PR2 and goes to the pickoff port of the power recycling cavity (known as "POP"). At the POP port we have several photodiodes as well as a DC quadrant photodiode. The "POP 2f" photodiodes will provide information on the RF sideband buildup in the power recycling cavity, independent of the resonant condition of the main laser frequency in any cavity.

The main beam is split at the beam splitter and half of the light is directed to each long Fabry-Pérot arm cavity. For each arm cavity, the beam transmits through an input test mass (ITMX or ITMY) and propagates toward the end test mass (ETMX or ETMY). Light reflected off of each ITM goes to a pickoff photodiode (POX or POY), which can be used to generate a length error signal for their respective individual Fabry-Pérot cavities. Light that leaks through the end test masses is directed toward DC photodiodes to monitor the power buildup in the arm cavities.

Light returning from the arm cavities toward the beam splitter will either go to the REFL port or to the anti-symmetric side of the BS. Light on this side of the BS reflects off of two folding mirrors in the signal recycling cavity (SR2 and SR3) and then is incident on the signal recycling mirror (SRM). Light transmitted through the SRM is incident on a set of photodiodes at the anti-symmetric (AS) port. AS $f_2 = 55\,\text{MHz}$ provides error signals for differential length degrees of freedom, while AS $2f_2 = 110\,\text{MHz}$ provides information on the RF sideband buildup in the signal recycling cavity.

At the ends of the interferometer, auxiliary laser beams are passed through second harmonic generation (SHG) crystals which are used to generate light at half the original wavelength, creating green 532 nm light from the infrared 1064 nm pump light. The green beams are locked to the arm cavities with error signals from the photodiodes REFL X and REFL Y for each arm. Note that the green beam resonant in the cavity will be colinear with the infrared beam, although they are shown slightly translated for clarity in the diagram. The main beam splitter is mostly transparent to the green light, as are the folding mirrors PR3 and SR3. This allows green light transmitted from each arm cavity to pass out of the vacuum system to a broadband photodiode (either BEAT X or BEAT Y). A small amount of light is also picked off from the main laser and frequency-doubled to generate green light, which is then used to create a beatnote with the light from either arm cavity.

Further details on the use of each group of components can be found in the text of Chapter 4.

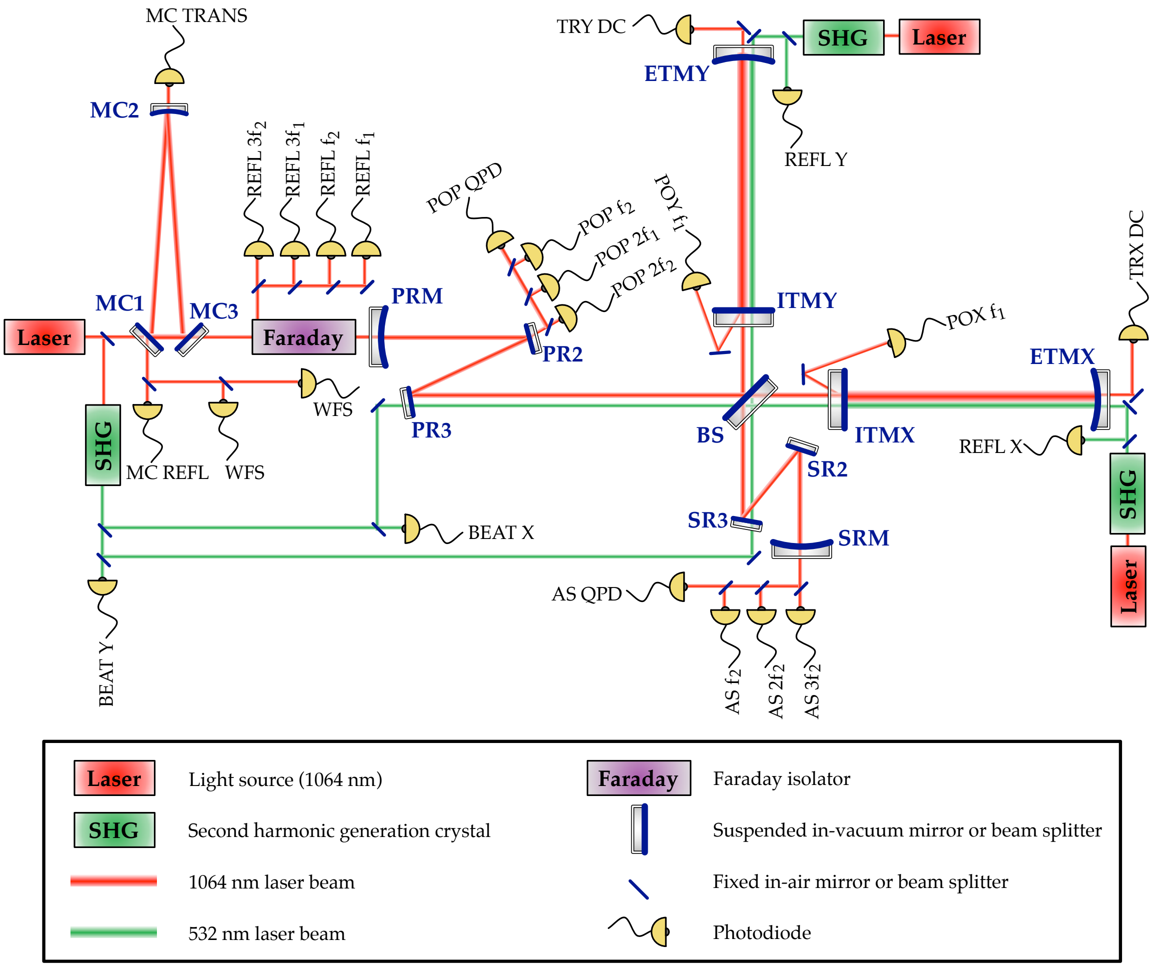

Figure 4.1: Schematic (not to scale) layout of the 40 m Lab's optical components. Note that the topology is similar to Advanced LIGO (and Figure 3.8), and has many of the same sensors and actuators. To-scale configuration of the 40 m Lab, including all optical components, is shown in Appendix C. Parameters for the 40 m interferometer are in Table 4.3 and Table 4.4.

## 4.1 Input optics

The light source for the 40m is a 2 W, 1064 nm Innolight Nd:YAG laser (nicknamed "Edwin"), labelled "Laser" in Figure 4.1. The laser beam transmits through a pre-mode cleaner (PMC, not shown), which rejects non-fundamental transverse modes of the beam. The PMC at the 40m is the same as those used in Initial LIGO, with round trip length 42 cm, and cavity finesse 700 [50, 51].

Following the PMC but prior to entering the vacuum system, the beam passes through an electro-optic modulator (EOM, not shown). An oscillatory voltage is applied across the crystal, which then introduces phase modulation to the beam at the frequency of the voltage applied. At the 40m, we apply modulations at 3 frequencies with this EOM, using a custom designed resonant circuit to amplify the frequencies 29.5 MHz, 11 MHz and 55 MHz. The 11 MHz, and 55 MHz frequencies will be used to acquire and maintain control of the interferometer, as discussed in Section 4.4.

The 29.5 MHz frequency is used to control the first in-vacuum cavity, the input mode cleaner (MC). The MC is a triangular ring cavity consisting of 2 flat mirrors (MC1 and MC3 in Figure 4.1) and one curved mirror (MC2). The round trip length of the MC cavity is approximately 27 m and it has a finesse of 1460. The precise value of the 11 MHz sideband frequency is chosen to be maximally transmitted through the MC to the rest of the interferometer.

The cavity length error signal is detected using the same type of PDH signal mentioned in Section 3.1.2. The 29.5 MHz sideband frequency is anti-resonant and so completely reflected from the cavity. Any carrier light that is reflected from the cavity is compared to this local oscillator, and the phase difference between them is detected using the "MC Refl" photodiode, and contains information about how far the cavity length is from resonance.

We also apply angular control to the mode cleaner cavity. We utilize two radio frequency resonant wavefront sensors (WFS in Figure 4.1) as well as a DC quadrant photodiode in transmission (MC Trans). These three diodes are used to feed back to all 3 mirrors, which holds the alignment of the cavity stable, and thus the input beam to the interferometer is also stabilized.

In addition to further filtering the transverse mode of the laser beam (by only resonating and transmitting the fundamental mode), the MC is also used for frequency stabilization of the laser. The MC is in vacuum on passive seismic isolation stacks, and each mirror is suspended, so the length of the cavity should be changing very little, particularly above a few Hz. Rather than actuating on the cavity mirrors to maintain laser resonance, the frequency of the laser is changed. At low frequencies we change the temperature of the laser crystal, which allows for long range slow actuation. At higher frequencies we use a PZT to actuate on the length of the laser crystal. This is capable of much faster actuation, but has a smaller range than the temperature actuation.

Following the mode cleaner, the laser beam passes through the Faraday isolator, which is used to pick off the backwards travelling light returning from the interferometer. This picked-off light is sent to the "Refl" suite of photodiodes as shown in Figure 4.1, which will be discussed in more detail in Section 4.4.

### 4.1.1 Input mode matching

The laser beam profile required for resonance in the input mode cleaner is much different from what is required for the full interferometer. Here we describe the use of Monte Carlo simulations in the design of a telescope to transform the beam exiting the input mode cleaner to one that will resonate in the main interferometer. The Monte Carlo is used to determine which potential solution will be most insensitive to errors in the placement of optics, or in the radii of curvature of the non-flat optics. In a situation with poor matching between the actual beam and the ideal beam incident on the interferometer, it would be much easier for a higher-order transverse mode of the beam to become resonant in the cavities, which can cause significant distortions to the length error signals. Thus, a mode matching solution that is robust against errors is highly desirable.

The profile of a beam can be described at any point $z$ along its path by the complex beam parameter

$$q(z) = z + iZ_R, \tag{4.1}$$

where $Z_R$ is the Rayleigh range of the beam and $z = 0$ is the location of the waist of the beam. The Rayleigh range is defined as

$$Z_R = \frac{\pi w_0^2}{\lambda}, \tag{4.2}$$

where $\lambda$ is the wavelength of the light, $w(z)$ is the radius at which the amplitude of the electric field of the beam is $\frac{1}{e}$ times the axial value, and $w_0 = w(0)$ is the minimum beam radius, known as the "waist" of the beam. For the ideal profile that will resonate in a Fabry-Pérot cavity with only one curved mirror we can find that

$$Z_R = L\sqrt{\frac{\text{RoC}}{L} - 1}, \tag{4.3}$$

where $L$ is the length of the cavity and RoC is the radius of curvature of the non-flat mirror. From $q(z)$ we can extract the radius $w(z)$ of the beam and the radius of curvature of the phase front $R(z)$ at any point along the path,

$$\frac{1}{q(z)} = \frac{1}{R(z)} - \frac{i\lambda}{\pi\,[w(z)]^2}. \tag{4.4}$$

A mode matching telescope was designed to match the $q_{\text{incident}}$ of the beam coming from the mode cleaner with the ideal $q_{\text{cavity}}$ of the arm cavities. To do this the ray tracing "ABCD" matrices [52] describing the optical path from the mode cleaner to the ITMs were expressed, with the distance elements as variables. Where possible, realistic descriptions of the optical components were used, e.g., the propagation through mirrors' substrates, and changes in the effective radius of curvature of interfaces that have different indices of refraction. See Chapter 15 of Siegman [52] for details on realistic ray tracing matrices. For a fixed choice of two curved mirrors to be placed in the path, the distances between the optics were optimized (using Matlab's standard *fminsearch*) to minimize the difference between $q_{\text{incident}}$ and $q_{\text{cavity}}$.

A goal of the new mode matching telescope (in contrast to the one which was used for the previous detuned RSE experiment) was to utilize off-the-shelf spherical mirrors rather than specialty off-axis parabolic mirrors. Since, for spherical mirrors, the effective radius of curvature of the mirror is modified by the angle of incidence of the laser beam differently for the sagittal and tangential beam axes, we would like to minimize the potential astigmatism by minimizing the angle of incidence of the beam on these mirrors. This implies that the distance between adjacent mirrors should be relatively long.

For every combination of a pair of off-the-shelf available mirrors, the distances between mirrors was optimized. Cases where the distances required would not realistically fit within the 40m's vacuum envelope were discarded. All non-excluded cases were put through a Monte Carlo simulation to estimate the sensitivity of the mode matching to misplacement of the mirrors or imperfections in the radius of curvature of any curved optic.

The variable parameters are each allowed to change in the Monte Carlo by some amount chosen from a normally distributed random number. For the telescope curved mirrors' radii of curvature, one standard deviation is set to be ± 0.5 %, the specification given by the manufacturer. The radius of curvature of the custom power recycling mirror (which the beam must pass through) is given a standard deviation of ± 5 m. The ideal positions of each optic is also allowed to vary, with a standard deviation of ± 4 mm.

For each set of perturbed radii and positions, the mode overlap integral between the beam incident from the mode cleaner through the optical path and the ideal arm cavity mode is calculated. The mode mismatch $\eta$ resulting from this integral is given by

$$\eta = 4\sqrt{\frac{Z_{R;\mathrm{x}}\, Z_{R;\mathrm{y}}\, Z_{R;\mathrm{a}}^2}{\left(z_{w;\mathrm{x}}^2 + (Z_{R;\mathrm{a}} + Z_{R;\mathrm{x}})^2\right)\left(z_{w;\mathrm{y}}^2 + \left(Z_{R;\mathrm{a}} + Z_{R;\mathrm{y}}\right)^2\right)}}, \tag{4.5}$$

where $Z_R$ is the Rayleigh range defined in Equation 4.2, for either the astigmatic beam (x and y for the tangential and sagittal axes) or the ideal arm mode and $z_w$ is the difference in locations of the waist position in the incident beam versus the ideal arm mode. We consider the two axes of the incident beam separately in order to account for the astigmatism caused by the non-normal angle of incidence on the curved telescope mirrors. See K. Arai's document [53] for a derivation of Equation 4.5.

The Monte Carlo simulation for each pair of curved telescope mirrors performs 30,000 iterations, and creates a histogram of the mode mismatch, $(1-\eta) = \frac{\delta\text{Power}}{\text{Power}}$. An example histogram is shown in Figure 4.2.

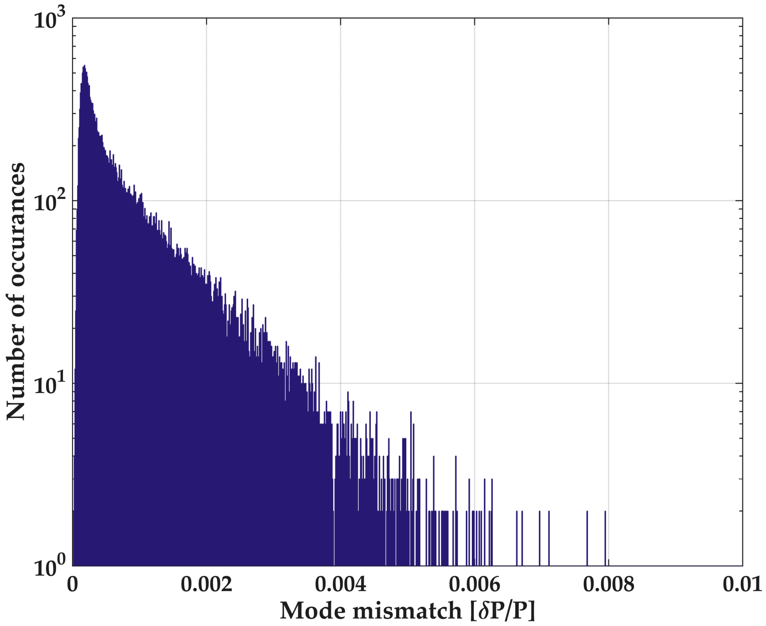

Figure 4.2: Example mode mismatch histogram from Monte Carlo simulation. Narrower histograms clustered toward the $\frac{\delta P}{P} = 0$ point indicate excellent mode matching despite errors in optic placement or optic radius of curvature.

This Monte Carlo simulation method is similar to numerically determining the diagonal elements of the Hessian matrix for $\eta$. Visually, we can create a hyper-ellipsoid where each axis represents $\eta$'s sensitivity to a different parameter. In the end, we will choose the solution with the smallest volume ellipsoid, indicating that it is robust against errors in optic placement or radius of curvature. We note from this analysis that the most sensitive parameter for this situation is the distance between the two curved mirrors, which indicates that special care should be taken to accurately place the optics the correct distance apart. Conversely, the mode matching is very insensitive to the location along the optical path that the telescope is placed, so we have more freedom in this dimension.

We conclude that we may misplace any optic by as much as 10mm, and still have a high probability of having excellent mode matching. We expect that we should be able, by using reasonable rulers, to place the optics within ~2mm or so, so we should have no trouble getting very good mode matching into the interferometer.

Figure 4.3 shows the beam radius profile expected from our chosen mode matching solution, for which the value of $\eta$ is calculated to be 99.9%. The vertical line indicates the ideal waist location for the arm cavity profile. For this final configuration, the first curved mirror (MMT1) has a radius of curvature of -5 m located 1.9282 m from the mode cleaner, while the second curved mirror (MMT2) has a radius of curvature of +8 m and is located 1.8760 m from MMT1.

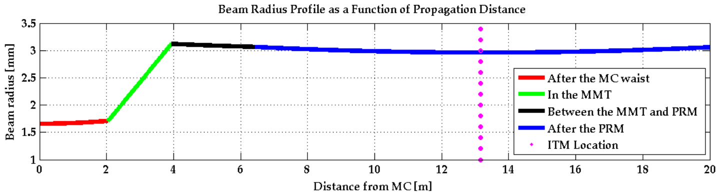

Figure 4.3: Ideal beam profile plot for our final design choice. Here 0 m is defined as the location of the mode cleaner waist.

Figure 4.4 shows the measured profile of the beam after the second curved mode matching mirror, as compared to the ideal beam profile. Here, the ideal beam profile does not include transmission through the power recycling mirror or other following optics, since the measurement was performed before those optics in the path. With this measured beam profile, the estimated mode matching to the interferometer is $\eta = 99.3\%$.

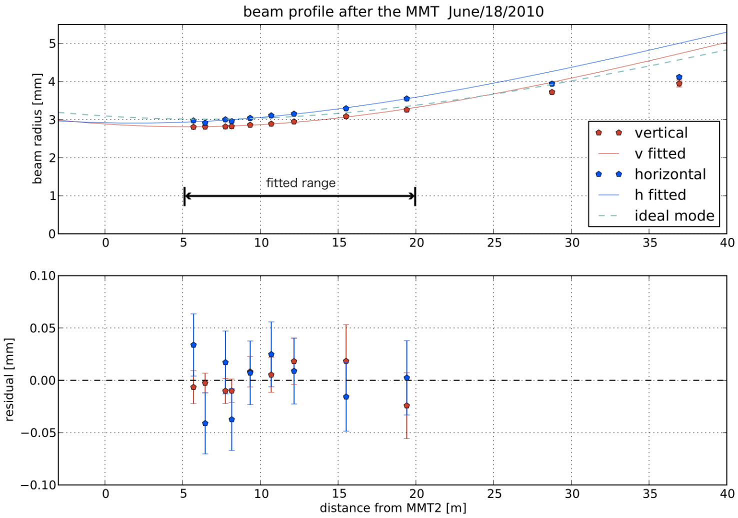

Figure 4.4: Measured mode profile of beam after the mode matching telescope, before power recycling mirror. Points at 27 m and 37 m are excluded from the fit, since they were clipping the aperture of the beam scanner, and so the beam radius measurement was not reliable [54]. Estimated mode matching to the interferometer is $\eta = 99.3\%$. Here 0 m is defined as the location of the second curved mirror in the mode matching telescope.

### 4.1.2 Input pointing tilt versus translation orthogonality

When determining how to lay out the optical path from the mode cleaner to the interferometer, another consideration was the ability to steer the input beam. In order to take advantage of our careful mode matching, we must match the laser beam's axis to that of the long arm cavities.

The axes each have four degrees of freedom: angle and translation, in both the tangential and sagittal planes. Ideally, our input steering would be able to actuate on the angle and translation of the input beam independently, however this is difficult to do in practice. To maximize the separation between these two degrees of freedom, we want the matrix describing the sensitivity of beam motion to angular actuation to be non-degenerate so that it can be inverted to create a control matrix.

Intuitively, we know that the steering optics should be far apart in the optical chain. Since the steering optics will be flat mirrors, we are free to place them where convenient within our vacuum envelope. For promising candidate solutions from Section 4.1.1 and maximally separated steering mirrors (within the constraints of the in-vacuum tables), we calculate the geometric ray-tracing response vector $(d, \alpha)$ of the beam propagated through the optical chain's ABCD matrix. For each steering mirror, we determine the response to angular actuation, and normalize the components by the beam waist, $w(0)$, or beam divergence angle, $\gamma = \lambda / (\pi * w_0)$, as appropriate. We then define an angle

$$\beta_{\mathrm{SMi}} = \arctan\left(\frac{\frac{d}{w_0}}{\frac{\alpha}{\gamma}}\right) \tag{4.6}$$

for the response of the beam motion to each of the steering mirrors. The difference between these response angles for each of the mirrors

$$\beta_{\mathrm{orthog}} = |\beta_{\mathrm{SM1}} - \beta_{\mathrm{SM2}}| \tag{4.7}$$

is our "orthogonality angle", which we want to be as close to 90° as possible.

To check that our convenience-chosen positions of the steering mirrors gave us the maximum available angle versus translation orthogonality, surface plots were generated showing $\beta_{\mathrm{orthog}}$ as a function of steering mirror positions. It was determined that the solution which fit most conveniently in our vacuum envelope resulted in an orthogonality angle of roughly 17°, and that moving the steering mirrors by many centimeters would only increase this number by about 0.5°. While 17° is not particularly orthogonal, it is sufficient to allow us to match the incident beam axis to the interferometer axis, and marginal changes to this $\beta_{\mathrm{orthog}}$ were not worth the effort of finding ways to mount mirrors between existing in-vacuum tables.

To help seismically isolate the beam motion, the steering mirrors were mounted in suspensions similar to those described in [55], which are used for steering in Advanced LIGO. These mirrors have electromagnetic actuators, but no local position sensing. Figure 4.5 shows one of these suspensions before placement in the vacuum system.

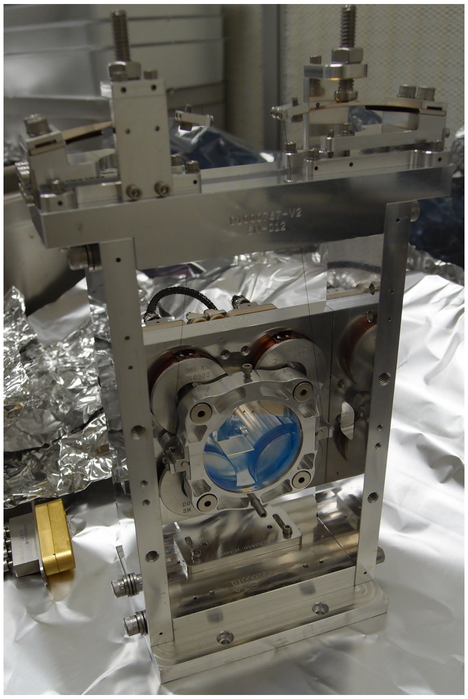

Figure 4.5: Input steering mirror.

## 4.2 Suspended optics

The 40 m lab has used the Initial LIGO style “small optic suspensions” [56] for the beam splitter, the recycling mirrors, and the input mode cleaner since the beginning of the detuned RSE experiment. However, with the upgrades to match the newer Advanced LIGO design, we replaced most of the mirrors, including making the arm cavity test mass mirrors smaller mass, so that all the core optics in the 40 m lab are this form factor with single loop steel wire suspensions.

Each optic has six degrees of freedom (three translation, and three rotation), which are described in Figure 4.7. Table 4.1 lists the resonant frequencies for each of the suspended optics.

| | Pitch | Yaw | Pos | Side | Bounce | Roll |
|---|---|---|---|---|---|---|
| MC1 | 0.672 Hz | 0.807 Hz | 0.968 Hz | 0.995 Hz | 16.31 Hz | 23.91 Hz |
| MC2 | 0.748 Hz | 0.820 Hz | 0.970 Hz | 0.993 Hz | 16.56 Hz | 24.06 Hz |
| MC3 | 0.763 Hz | 0.845 Hz | 0.980 Hz | 0.971 Hz | 16.35 Hz | 23.95 Hz |
| PRM | 0.612 Hz | 0.833 Hz | 0.993 Hz | 0.999 Hz | 16.38 Hz | 23.92 Hz |
| SRM | 0.567 Hz | 0.808 Hz | 0.962 Hz | 0.972 Hz | 16.44 Hz | 23.99 Hz |
| BS | 0.747 Hz | 0.792 Hz | 0.957 Hz | 0.995 Hz | 16.05 Hz | 23.51 Hz |
| ITMX | 0.514 Hz | 0.831 Hz | 0.966 Hz | 0.986 Hz | 16.09 Hz | 23.49 Hz |
| ITMY | 0.601 Hz | 0.855 Hz | 0.988 Hz | 1.003 Hz | 16.09 Hz | 23.49 Hz |
| ETMX | 0.829 Hz | 0.909 Hz | 0.951 Hz | 1.038 Hz | 16.42 Hz | 24.03 Hz |
| ETMY | 0.860 Hz | 0.894 Hz | 0.982 Hz | 1.016 Hz | 16.39 Hz | 23.99 Hz |

Table 4.1: Measured suspension resonance frequencies for the 40 m Lab’s optics [57].

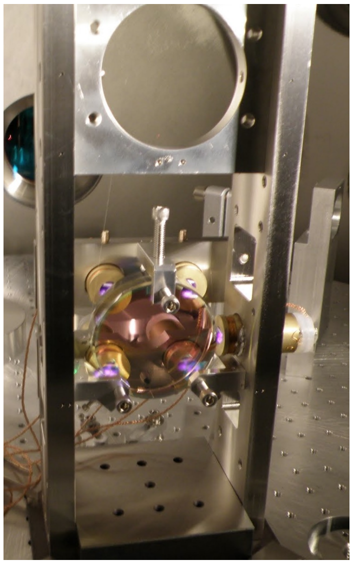

Figure 4.6: Photo of small optic suspension.

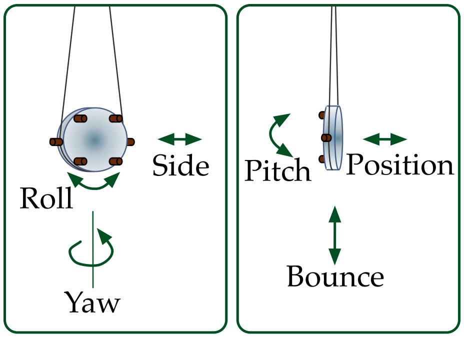

Figure 4.7: Suspended optic, viewed from the back (left figure) and side (right figure). Positions of permanent magnets are shown, as well as the canonical basis of six degrees of freedom for suspended optics.

Every optic has four permanent magnets glued to the back surface, and one on the side. The positions of these magnets is shown in Figure 4.7. Each of these is matched with an optical shadow sensor and electromagnetic actuator ("OSEM"), which are described in Section 4.1.8 of [58]. OSEMs are only designed to measure and control translational motion of the optic at the magnet attachment point. By combining the four sensors on the back face of the optic, we can make a change of basis so that we sense the pitch, yaw, and positional motion of the optic. The side-to-side swinging is sensed by the side OSEM, while the bounce and roll modes are (ideally) not sensed.

The naive matrix for the basis change between the individual sensors and the degrees of freedom shown in Figure 4.7 is

$$\begin{array}{c} \\ \mathrm{UL} \\ \mathrm{UR} \\ \mathrm{LL} \\ \mathrm{LR} \\ \mathrm{SD} \end{array} \begin{array}{c} \begin{array}{ccccc} \mathrm{Pit} & \mathrm{Yaw} & \mathrm{Pos} & \mathrm{Null} & \mathrm{Side} \end{array} \\ \begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 1 & -1 & 1 & -1 & 0 \\ -1 & -1 & 1 & 1 & 0 \\ -1 & 1 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array} . \tag{4.8}$$

By looking at the peak heights at the resonant frequencies for each mode in each sensor, we can create an "adjustment" matrix that will completely diagonalize the system. The condition number for the adjustment matrix tells us how far it is from a unitary matrix such as the identity matrix. So the condition number tells us how far our actual matrix is from the matrix in Equation 4.8.

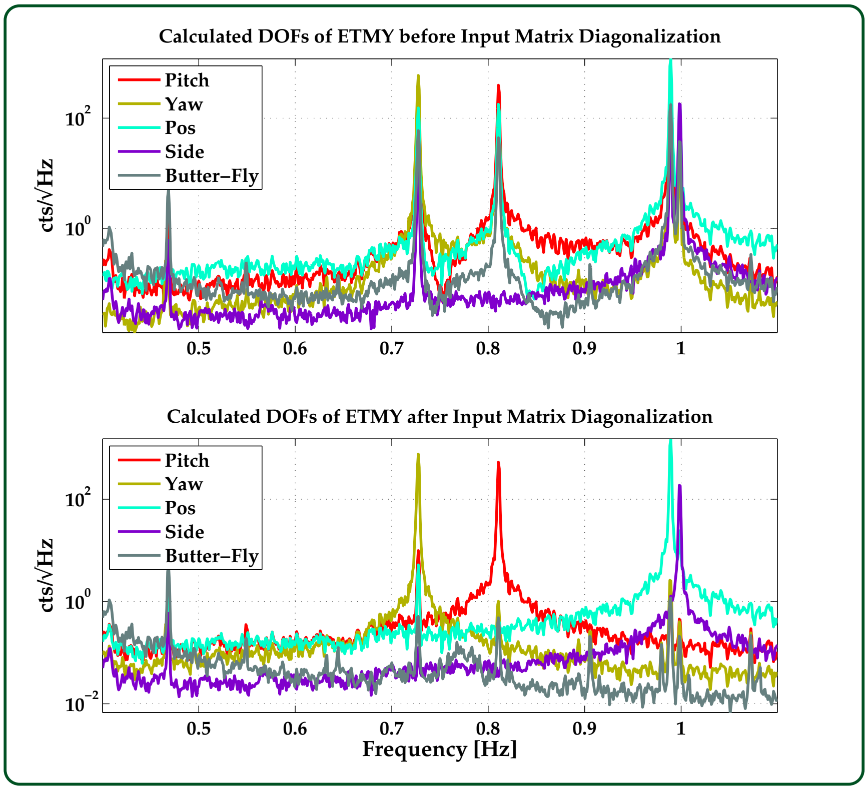
Calculated DOFs of ETMY before Input Matrix Diagonalization

Calculated DOFs of ETMY after Input Matrix Diagonalization

Figure 4.8: Diagonalization and basis rotation of suspension sensors. Top shows the sensors transformed to pitch, yaw, position, nullstream (butterfly), and side using the matrix in Equation 4.8. Bottom shows the sensors transformed using the naive matrix times the adjustment matrix. Notice in the bottom plot that each resonance peak is dominated by a single trace, indicating a lack of coupling. The condition number for the adjustment matrix used here was 4.2.

In order to know what effect a given amount of force will have on each optic, we must calibrate each suspension actuator in the "pos" direction. We can first calibrate the input test masses and the beam splitter by looking at the effect that pushing on them has on the Michelson-only error signal. Since the anti-symmetric port of the Michelson will change from maximally dark to maximally bright every $\lambda/2$ meters, we can look at the peak-to-peak values of the free-swinging anti-symmetric port photodiode to calibrate it from arbitrary sensor counts to meters. If we control the positions of the mirrors (to keep them within the linear range of the sensors) only at low frequencies, and then measure the actuation response only at high frequencies where the control loop gain is very small, we can avoid the need for loop compensation, and can directly use our measurements. All other optics can be calibrated in reference to the input test masses, so that we have an absolute calibration for each mirror. Table 4.2 lists the measured actuator calibrations.

| Optic | Position calibration |
|---|---|
| MC2 | $(15 \pm 1) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [59] |
| PRM | $(19.6 \pm 0.3) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [60] |
| SRM | $(19.0 \pm 0.7) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [61] |
| BS | $(20.7 \pm 0.1) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [62] |
| ITMX | $(4.70 \pm 0.02) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [62] |
| ITMY | $(4.66 \pm 0.02) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [62] |
| ETMX | $(13.31 \pm 0.21) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [59] |
| ETMY | $(13.59 \pm 0.20) \times 10^{-9} \left(\frac{\mathrm{Hz}}{f}\right)^2 \frac{\mathrm{m}}{\mathrm{ct}}$ [59] |

Table 4.2: Measured actuator calibrations for 40 m Lab's optics. These values are applicable above the pendulum resonant frequency.

## 4.3 Recycling cavities

One of the most significant differences between the current 40 m lab configuration and the previous one is the length of the recycling cavities. To simplify the radio frequency electronics, Advanced LIGO modified the design modulation frequencies. To match this new Advanced LIGO design, the 40 m is using lower frequencies than before. In order to resonate these frequencies, we require much longer recycling cavities than during the detuned resonant sideband extraction experiment. Previously, the 40 m power recycling length ($l_{\mathrm{PRC}}$) was approximately 2.3 meters. Now it is 6.753 m. Fitting this inside our pre-existing vacuum system required "folding" the cavity with two additional mirrors (PR2 and PR3 in Figure 4.1). Similarly the signal recycling cavity length ($l_{\mathrm{SRC}}$) increased from about 2.2 m to 5.4 m. See Table 4.4 for details on the recycling cavity parameters.

The folding mirrors are suspended on purely passive short single loop suspensions, as seen in Figure 4.9. They are similar in design to those described in [55] and shown in Figure 4.5, although they do not have any active sensing or actuation. We utilize compliant eddy current dampers with damping rings suspended behind the suspended optic, so that they do not directly inject seismic noise onto the optic. The lack of sensing and actuation on our folding mirrors has turned out to be a poor design choice. The suspensions are shorter than our test masses, so the pendulum resonant frequency is higher and they pass more seismic noise into the cavity than do the other optics.

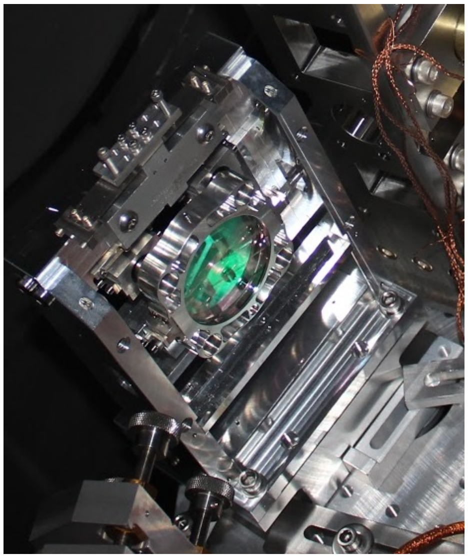

Figure 4.9: Recycling cavity folding mirror (SR2). Suspension point is isolated in the vertical direction at the top with "blade" springs. Mirror is held in an aluminum ring, which is suspended from the blade springs. The corners of the ring hold magnets, which along with the steel blocks separately suspended from the frame, induce eddy currents that are used to passively damp the optic's motion.

In particular, this has caused excess angular motion in our recycling cavities, which has proved problematic for locking and control. As in Figure 4.10, we can infer how much horizontal ground motion affects the angular motion of the individual optics in our cavity. This coupling estimate takes into account the passive seismic isolation stack (which is the same for both types of optic), and then calculates the position to pitch transfer function, following the derivation in [63]. The estimated coupling for each type of suspension is then multiplied by measured seismic ground motion. Recall that we have two folding mirrors per recycling cavity, and the beam reflects off of each of those twice per round-trip, so the seismic motion due to the folding mirrors is roughly $2\sqrt{2}$ times that of a single mirror. Also shown in Figure 4.10 is a measurement of the cavity axis pitch motion (multiplied by an arbitrary calibration factor), as seen by the quadrant photodiode at the POP port. Looking at the relative shapes of the angular motion spectra in Figure 4.10 makes it clear that a majority of the cavity axis motion is due to the folding mirrors. Areas where the cavity axis trace is much larger than either mirror contribution indicates that the quadrant photodiode signal is measuring noise unrelated to true cavity axis motion (such as vibrations of the out-of-vacuum table or mounting apparatus, air currents, etc.).

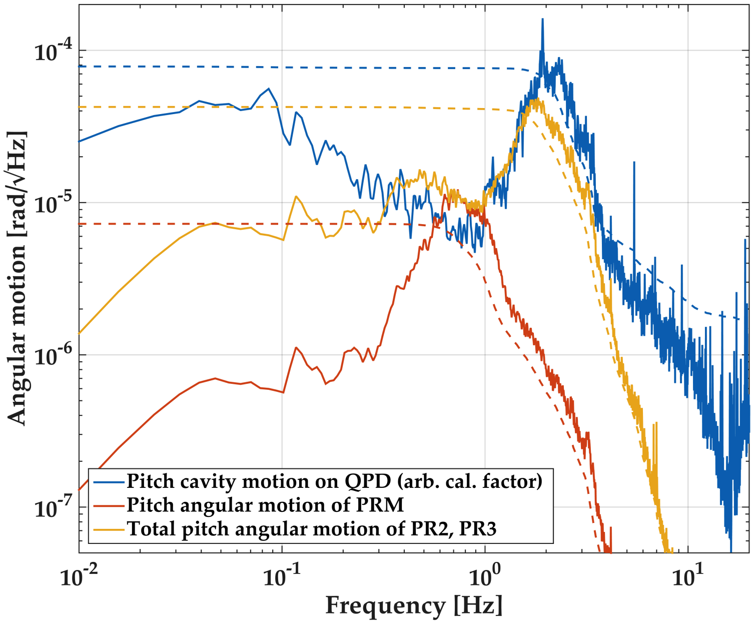

Figure 4.10: Pitch motion of power recycling cavity mirrors due to seismic motion compared to cavity axis motion. Measured horizontal seismic motion is propagated through the isolation stack (same for all optics) and then through the SOS or folding mirror suspensions' coupling from longitudinal motion at the suspension point to angular motion. Red trace shows the angular motion of the power recycling mirror due to seismic motion. Gold trace shows the estimate total angular motion of the folding mirrors PR2 and PR3. The contribution from a single folding mirror is calculated, and then multiplied by two (since the beam will reflect off of each folding mirror twice per round trip), and then multiplied by $\sqrt{2}$ to estimate the incoherent addition of the motion of each mirror. Blue trace shows the measured pitch motion of the power recycling cavity axis, multiplied by an arbitrary factor such that the shape can be roughly compared with the mirror motions. Dashed traces show RMS of each spectra.

This large amount of seismic motion is particularly problematic for the recycling cavities, since they are nearly unstable. The stability criterion for a 2-mirror cavity is

$$0 \leq g_1 g_2 \leq 1, \tag{4.9}$$

where each g-factor is given by

$$g = 1 - \frac{L}{RoC}, \tag{4.10}$$

with $L$ the length of the cavity and $RoC$ the radius of curvature of the mirror. The more general form of the stability criterion for a multiple mirror cavity (such as our recycling cavities) has been derived by K. Arai [64] as

$$0 \leq \frac{A + D + 2}{4} \leq 1, \tag{4.11}$$

where $A$ and $D$ are elements from the ray transfer analysis matrix ("ABCD matrix") of the round trip path through the cavity. Note that for a 2-mirror cavity, this simplifies to Equation 4.9.

With the mirror parameters listed in Table 4.3, the power recycling cavity is expected to have a stability number of about 0.89. The values for the radii of curvature in Table **??** for the 40 m's recycling cavity folding mirrors are different than those for the original installation of the mirrors. The polish for these mirrors was not well-specified, only stating that the mirrors should be "flat". After great difficulty trying to control the recycling cavities, we measured the radii of curvature of the folding mirrors and found them all to be close to -600 m. Having two convex mirrors made the stability number from Equation 4.11 for both recycling cavities very nearly unity.

Rather than order new mirrors (with a lead time greater than 6 months) that have the dichroic coatings to maximally transmit green light while being highly reflective for infrared, the folding mirrors have all been installed backward. For each optic, the main beam passes through the anti-reflective (AR) coating and through the glass substrate before hitting the highly reflective surface. The transmission through the AR coating and the substrate add to the loss of the mirrors and also alters the effective radius of curvature of the optic. For our folding mirrors, the new effective radii of curvature is approximately 400 m for each, as shown in Table 4.3. The cavity stability achieved was deemed an acceptable trade-off to the increase in loss of the cavity.

One major issue with having marginally stable cavities is that the frequency spacing between the fundamental $\mathrm{TEM}_{00}$ mode and higher-order spatial modes is quite small. The spacing between the $\mathrm{TEM}_{00}$ mode and a higher order $\mathrm{TEM}_{nm}$ mode is given by

$$\Delta f_{(00),(nm)} = (n+m)\frac{\arccos\left(\sqrt{\frac{A+D+2}{4}}\right)}{\pi} f_{\mathrm{FSR}}, \tag{4.12}$$

where $A$ and $D$ are again the elements from the ray transfer matrix and $f_{\mathrm{FSR}}$ is given in Equation 3.7. The closer the stability number is to unity, the easier it will be for higher order modes to become co-resonant in the cavity, which makes the cavity challenging to control. Angular motion, for example seismically-induced and coupled through the folding mirrors, can make it even more likely that we will see higher order mode resonances.

We have a quadrant photodiode located at the pickoff of the power recycling cavity ("POP" in Figure 4.1), and in the absence of the arm cavities, we can use this to feedback to the PRM, and suppress the angular motion of the cavity axis. However, when the full interferometer is locked, both the RF sideband and the main carrier light will be resonant in the power recycling cavity. The carrier light's apparent angular motion will be dominated by the cavity axes of the long arm cavities, so we would like to use only the angular information of the RF sideband. We do not have photodiodes at the POP port that are only sensitive to the RF light, so we cannot use the POP QPD in feedback when the full interferometer is locked. Section 5.5 will discuss an alternative, specifically the use of seismic noise cancellation to make the cavity more stable.

## 4.4 Locking: acquiring control

Once we have all the pieces of an interferometer, they must be made to work together. In particular, we need a system that will allow us to hold all of the mirrors such that the laser beam resonates throughout the interferometer. All of the techniques described here were applied to the power recycled Fabry-Pérot Michelson configuration of the 40 m Lab, with the signal recycling mirror misaligned. Future work includes commissioning of the dual recycled configuration of the interferometer.

For any given Fabry-Pérot cavity, we use Pound-Drever-Hall reflection locking, as described in Appendix B, as well as [27, 28]. In order to do this, we use an electro-optic modulator (EOM) to imprint phase modulations on the main laser light. An EOM has a birefringent crystal through which the light will propagate. We apply a voltage along one of these axes to modulate the phase of the light in that polarization component. Ideally, the light incident on the EOM crystal will be perfectly S-polarized, and matched to the vertical axis of the EOM. If the light is slightly mis-polarized, the component in the vertical direction will be modulated, but the component in the horizontal direction will be unaffected. This results in a rotation of the polarization of the light exiting the EOM. As soon as this passes through a polarized optical element (for example a polarizing beam splitter that only transmits a single polarization), the rotation of the light's polarization will cause some amount of the light to be rejected, resulting in residual amplitude modulation (RAM) of the RF sideband. This can cause offsets in our length degree of freedom error signals, and cause our cavities to be detuned. See [65] for details on how RAM can affect interferometers.

For the situation of many coupled cavities (i.e. 2 or more Fabry-Pérot cavities with shared mirrors), we must find a way to derive an error signal for each of the individual length degrees of freedom. R. Ward [39] found that it was very challenging to do this with only a single laser and 2 RF sidebands. For Advanced LIGO, an auxiliary arm length stabilization (ALS) system will be used. This system was first prototyped at the 40 m Lab [46, 47]. The ALS system (described in more detail in Section 4.4.1) is designed to allow the long arm cavities to be controlled at a fixed length, but not necessarily in the presence of a valid PDH signal from the main laser.

While the arms are held off resonance, we use the Arai 3f locking scheme [31] to acquire and maintain control of the short corner degrees of freedom (MICH and PRCL for the power recycled interferometer, MICH, PRCL, and SRCL for the dual recycled interferometer). The key feature of the 3f error signals is that they are less sensitive than the usual first order error signals to the change in effective reflectivity of an arm cavity when we eventually bring them to resonance.

The reflectivity of the compound mirror created by the arm cavities is given by

$$r_{\mathrm{cav}} = -r_i + \frac{t_i^2 r_e e^{-i\phi}}{1 - r_i r_e e^{-i\phi}}, \tag{4.13}$$

where $r$ and $t$ are the amplitude reflectivity and transmission of the input and end cavity mirrors as in Equation 3.14 and Equation 3.15, and $\phi$ is the phase of the light after travelling the length of the cavity, as defined in Equation 3.17. When the arm cavity is fully resonant, $\phi = 2\pi$, but if the arm cavity is anti-resonant, $\phi = \pi$. Since this changes the resonance condition of the carrier light in the power recycling cavity (assuming we have the PRC anti-resonant for the carrier when the arms are held off resonance), it will change all of the 1f PDH signals at the reflection port of the interferometer, making it impossible to maintain control of any cavities using the 1f signals during this transition.

The 3f signal is comprised of all components that beat and create a signal at a frequency 3 times the modulation frequency. For example, as can be seen in Figure 4.11, $(f_1) * (-2f_1)$ and $(3f_1) * (f_0)$ will both contribute at different levels. Since these are intermodulations of harmonics of the initial RF modulations, the amplitude of these 3f signals will be small, and so the signal to noise ratio will not be excellent. We will use these signals only until the arm cavities are brought to resonance, at which point we can transition to the larger SNR 1f PDH signals.

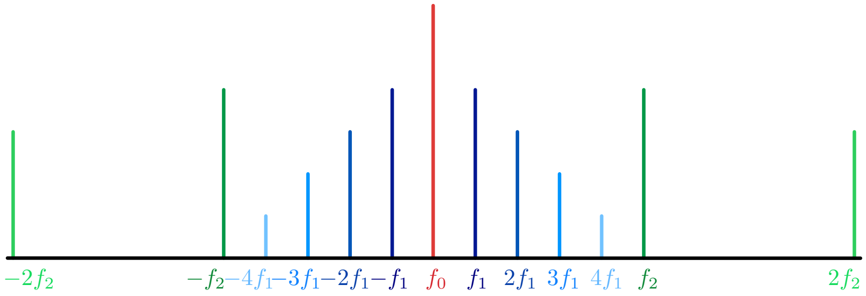

Figure 4.11: Cartoon of RF sidebands for 3f signal generation. Red is the main carrier light. Dark blue is the initial modulation at 1f (11 MHz), and lighter blues are harmonics of that modulation. Dark green is the initial modulation at 2f (55 MHz), and light green is the second harmonic.

We have found that the 3f signal is sensitive to the relative phase and amplitude between the RF modulations, because of potential cancellation between various products. For example, some signal is created at the $4f_1$ frequency on the input beam, as a result of the $f_2$ signal being imprinted on the light that already has the $f_1$ signal and harmonics thereof. This spurious 44 MHz sideband will create extra signal at $3f_1$ by beating against the main $f_1$ sideband. For simplicity, since we are only working on the PRFPMI configuration, we only use the $3f_2$ signals for locking at this time. Investigations into this phenomenon will be continued at the 40 m Lab during commissioning of the DRFPMI configuration.

Once the vertex degrees of freedom are controlled, we use the ALS system to bring the arms into resonance for the main laser. The power recycling cavity, coupled with the arm cavities, makes the effective linewidth of the CARM degree of freedom much smaller than the single-arm case. For the 40 m Lab, the linewidth is approximately 20 pm, while the RMS motion of the arm cavity is 23.5 pm in the very best case, from the noise of the ALS system [46].

It is important to note here that although the residual motion of the cavity is larger than the linewidth, the PDH signal is still a useful error signal for the arm cavities. Eventually, of course, we will require that the cavity remain in the linear region of the error signal, but during the transition from ALS control to PDH control (or at any time while acquiring control of a cavity), we only need the error signal to be of the correct sign. For example, if the PDH signal is positive when the cavity is too long, at any time the error signal is positive the feedback loop will work to shorten the cavity, and *vice versa*. This has the effect of helping to slow down the velocity of the mirrors, and force them toward the desired zero-crossing point.

We utilize this fact by slowly blending the low frequency components of the PDH signal for the CARM and DARM degrees of freedom in with the ALS error signal (described in more detail in Section 4.4.2). Once the residual motion of the arm cavities has decreased to within a few cavity linewidths, we discontinue use of the ALS error signals entirely. This method allows us to skip several intermediate candidate error signals for the long arm cavities, but it does require that we transition both CARM and DARM degrees of freedom roughly simultaneously.

After the arm cavities are held on-resonance by the low noise PDH signals, we transition the vertex degrees of freedom to their final 1f reflection signals, which are far less noisy than the temporary 3f error signals. For the PRCL degree of freedom, it is critical that the CARM residual noise be very small, since the error signals will show up in the same phase at the reflection port. We must fully suppress the CARM motion before we will be able to sense the PRCL error signal adequately.

### 4.4.1 Auxiliary green locking

At the 40 m Lab, we utilize the arm length stabilization (ALS) "green locking" system to control the long arm cavities independently of the main laser. We use information about the difference between the main laser's frequency and the frequency that would resonate in the cavities as an error signal to control the cavity length. Here, the method of deriving that error signal is described.

As discussed in K. Izumi's thesis [47], we use a delay line frequency discriminator to read out the beat frequency between the main laser and the auxiliary laser. Since the auxiliary laser's frequency is locked to the arm cavity (following the motion of the cavity), this beatnote provides information about the main laser's frequency relative to the length of the arm cavities.

The delay line frequency discriminator ("beatbox") electronics are described in D1102241 [66], with some modifications [67, 68].

The output of the broadband photodiode [69] gives us a signal proportional to

$$\sin(\omega_{\mathrm{RF}} t), \tag{4.14}$$

where $\omega_{\mathrm{RF}}$ is the beat frequency between the two lasers. A portion of this signal is picked off and sent through a 30 m delay line cable, such that at the end of the cable the signal is proportional to

$$\sin\left(\omega_{\mathrm{RF}}t + \frac{2\pi l}{v}\omega_{\mathrm{RF}}\right), \tag{4.15}$$

where $l$ is the length of the delay line and $v$ is the speed of propagation in the cable, which is roughly $\frac{2}{3}$ the speed of light. The rest of the signal is split in two, with one portion (the Q-phase) delayed 90° out of phase from the other (the I-phase).

The delay line signal is mixed with each of the RF signals, such that the beatbox has

$$\begin{aligned} I &\propto \sin\left(\omega_{\mathrm{RF}}t\right) \times \sin\left(\omega_{\mathrm{RF}}t + \frac{2\pi l}{v}\omega_{\mathrm{RF}}\right) \\ Q &\propto \sin\left(\omega_{\mathrm{RF}}t + \frac{\pi}{2}\right) \times \sin\left(\omega_{\mathrm{RF}}t + \frac{2\pi l}{v}\omega_{\mathrm{RF}}\right). \end{aligned} \tag{4.16}$$

After applying a lowpass, we are left with

$$\begin{aligned} I &\propto \cos\left(-\frac{2\pi l}{v}\omega_{\mathrm{RF}}\right) \\ Q &\propto \cos\left(\frac{\pi}{2} - \frac{2\pi l}{v}\omega_{\mathrm{RF}}\right). \end{aligned} \tag{4.17}$$

These are the signals that are sent to the digital system. Unfortunately, neither is linear in $\omega$ over a very large range. However, we can combine the $I$ and $Q$ phase signals to infer the frequency of the beatnote over a wide range. We rotate the $I$ and $Q$ signals using a standard rotation matrix,

$$\begin{pmatrix} I' \\ Q' \end{pmatrix} = \begin{pmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{pmatrix} \begin{pmatrix} I \\ Q \end{pmatrix}. \tag{4.18}$$

This gives us

$$I' \propto \sin\left(\frac{\pi}{2} - \left(\frac{2\pi l}{v}\omega_{\mathrm{RF}} - \phi\right)\right), \tag{4.19}$$

and similar for $Q'$. If we chose $\phi$ at each moment in time such that $I' = 0$ by applying a simple feedback loop, then the argument of the sin is always zero. This means that

$$\phi = \frac{2\pi l}{v}\omega_{\mathrm{RF}} - \frac{\pi}{2} \tag{4.20}$$

within the bandwidth of the loop (which we chose to be about 2 kHz). The closed loop response of this feedback loop gives us some frequency dependence to the sensor in our ALS locking loop, but the 2 kHz bandwidth is much higher than any of our other loops, so it provides a nearly-flat response.

This "phase tracking" technique allows us to precisely know the relative frequency between the main laser and the auxiliary laser, even when the arm cavity is freely swinging over many MHz. The value of $\phi$ is used as the error signal that can be fed back to the arm cavity mirrors to stabilize the length of the cavity (since the main laser's frequency is stable, stabilizing the arm so that the relative frequency between the main and auxiliary lasers is stationary holds the arm length fixed).

### 4.4.2 Frequency dependent error signal blending

In order to transition the long arm cavities from ALS control to PDH control, we chose to blend the error signals, as shown in Figure 4.12. Error signal blending is more convenient than control signal blending, since the system's stability is guaranteed, as long as the two error signals are calibrated to have the same sensitivity and see the same plant.

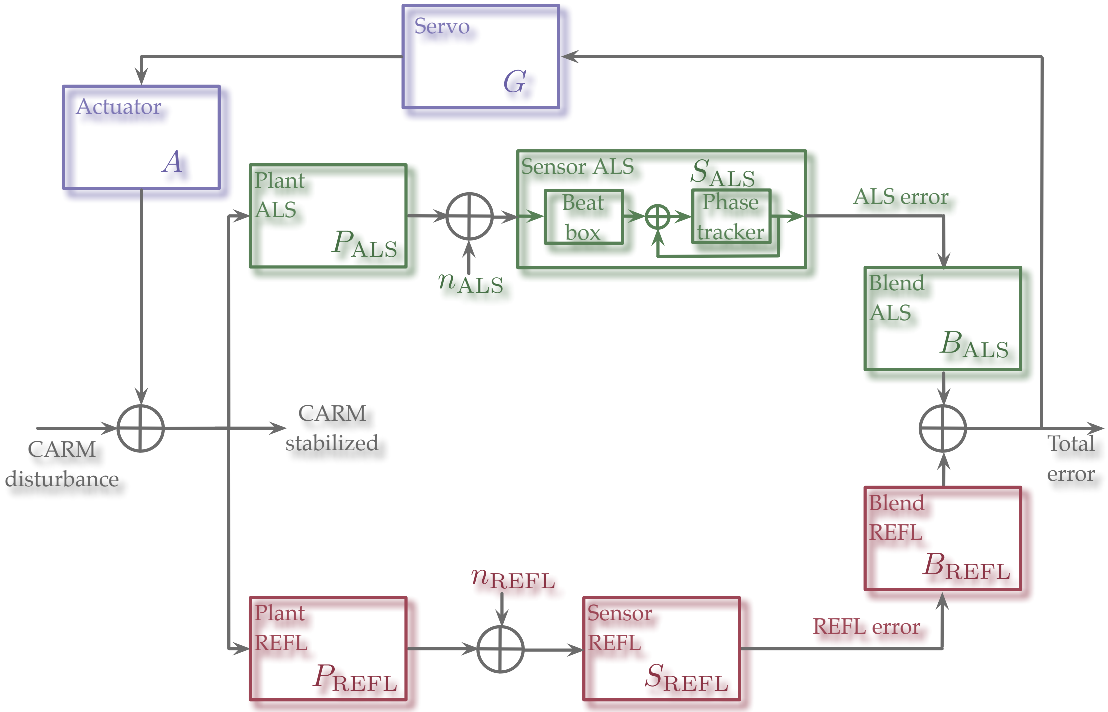

Figure 4.12: Block diagram of error point blending.

In our system, the ALS system sees a plant that includes a pole due to the optical response of the green light in the cavity, at 79.3 kHz. The PDH error signal, because of the coupled cavity phenomenon discussed in Section 4.4, will instead see a pole below 200 Hz. We account for the difference in the two plants by placing a zero at the coupled cavity pole frequency in the path of the PDH error signal at $B_{\mathrm{REFL}}$ just before the summing node. Figure 4.13 shows the two individual loop gains, as well as the combined total loop gain, partway through the transition.

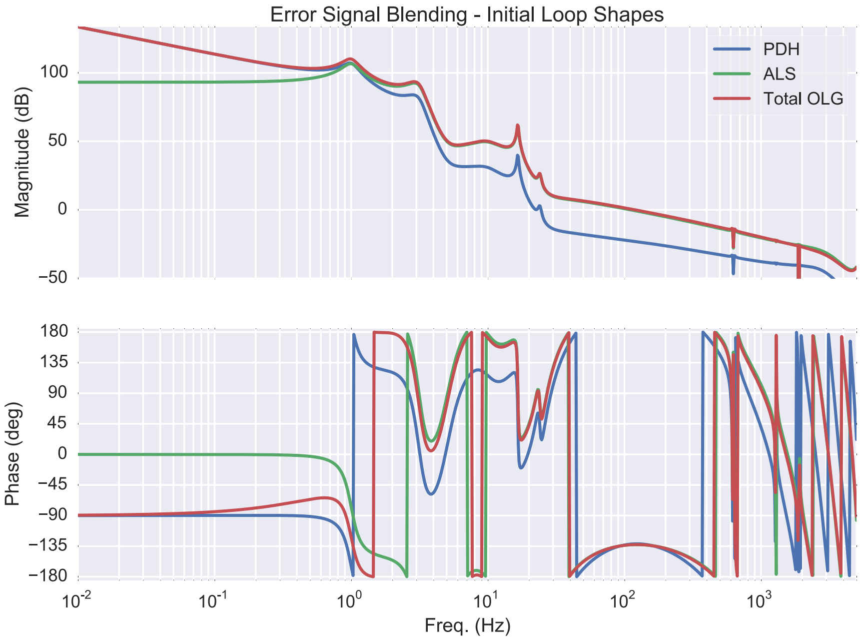

Figure 4.13: Individual loop gains, as well as the combined total loop gain, partway through the transition.

To begin the transition between the two loops, once we have a valid PDH error signal we start to blend in the low frequency component of the signal by placing a pure integrator (pole at 0 Hz) in the PDH blend block. As the gain of the PDH loop is increased, the crossover point between the two error signals increases in frequency. This crossover is unconditionally stable as we increase the crossover frequency. Once the gain of the PDH loop is sufficient, we de-boost the low frequency components of the ALS signal, and finally turn $G_{\mathrm{ALS}}$ to zero.

This scheme allows us to take advantage of the lowest noise components of each error signal throughout the transition, while maintaining the benefits of error point blending.

Figure 4.14 shows an example of a smooth transition of the CARM and DARM error signals using this technique. Each degree of freedom has a loop topology as shown in Figure 4.12. Notably, the residual RMS of the blended error signal decreases, and the power transmitted through the cavity becomes stable.

### 4.4.3 Feedforward decoupling of loops

An alternative scheme for transitioning from ALS control of the arm cavities to PDH control uses a feedforward path to effectively decouple the 2 loops. This “fools” the ALS loop into thinking it

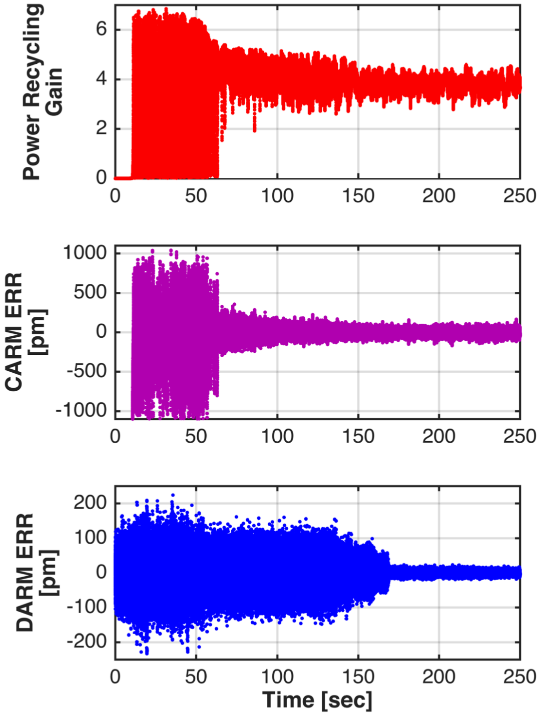

Figure 4.14: Time series of transition from ALS to IR PDH control. The error signals are taken from the "total error" point for each system, after the blend. The arm cavities begin with an offset of approximately 3 nm, so there is no resonance of the infrared laser beam in the arm cavities (the power recycling gain is representative of this value). At ~10 s, the CARM offset is removed. At ~60 s, the integrators for both CARM and DARM are engaged, and the CARM degree of freedom is almost fully transitioned to the PDH error signal, although the transition is not yet completed. At ~170 s the DARM degree of freedom is transitioned completely to its PDH signal and the CARM transition is finished.

maintains full control, while allowing the PDH loop to engage.

Figure 4.15 shows a block diagram of the proposed control system. Green blocks represent the ALS loop and red blocks represent the PDH loop, also referred to here as the "Refl" loop. Initially, $G_{\mathrm{REFL}} = 0$, so we only have a simple single loop system. Once the arm cavities are close to resonance and we have a valid PDH error signal, we begin increasing the gain of the secondary loop. Without any decoupling, the two loops can act against one another, and push the system into instability.

The purple decoupler block feeds the control signal of the Refl loop through an appropriate transfer function to the error point of the ALS loop. The transfer function of the decoupler is chosen so that any signal that arrives at the ALS loop's error point due to the Refl loop is entirely cancelled out. If we calculate using the nodal matrix technique described in Appendix A the transfer function

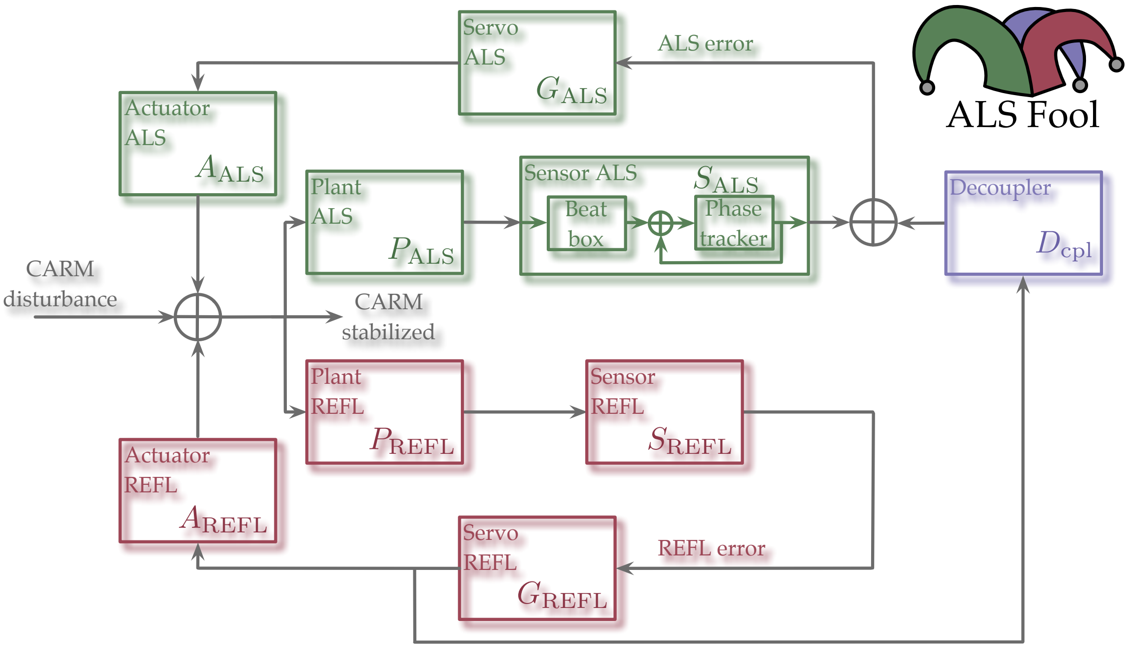

Figure 4.15: Block diagram of feedforward decoupling of loops.

of the system shown in Figure 4.15 from the output of $G_{\text{REFL}}$ to the input of $G_{\text{ALS}}$, we find

$$\left(G^{out}_{\text{REFL}} \rightarrow G^{in}_{\text{ALS}}\right) = \frac{D_{\text{cpl}} + A_{\text{REFL}}P_{\text{ALS}}S_{\text{ALS}}}{1 - A_{\text{ALS}}G_{\text{ALS}}P_{\text{ALS}}S_{\text{ALS}} - A_{\text{REFL}}G_{\text{REFL}}P_{\text{REFL}}S_{\text{REFL}} - D_{\text{cpl}}A_{\text{ALS}}G_{\text{ALS}}G_{\text{REFL}}P_{\text{REFL}}S_{\text{REFL}}}. \tag{4.21}$$

Setting this to zero implies that $D_{\text{cpl}}$ should be set to

$$D_{\text{cpl}} = -A_{\text{REFL}}P_{\text{ALS}}S_{\text{ALS}}. \tag{4.22}$$

If we use this value of $D_{\text{cpl}}$ to examine the stability of the full system's transfer function, we find that the closed loop is

$$\text{CLG} = \frac{1}{1 - H_{\text{ALS}} - H_{\text{REFL}} + H_{\text{ALS}}H_{\text{REFL}}} \tag{4.23}$$

if

$$H_{\text{ALS}} = A_{\text{ALS}}G_{\text{ALS}}P_{\text{ALS}}S_{\text{ALS}} \qquad \text{and} \qquad H_{\text{REFL}} = A_{\text{REFL}}G_{\text{REFL}}P_{\text{REFL}}S_{\text{REFL}}. \tag{4.24}$$

The open loop gain of the full system is then

$$\text{OLG} = H_{\text{ALS}} - H_{\text{REFL}} + H_{\text{ALS}}H_{\text{REFL}}. \tag{4.25}$$

For a single arm Fabry-Pérot cavity, the bode plot of the open loop gain in Equation 4.25 is shown in Figure 4.16.

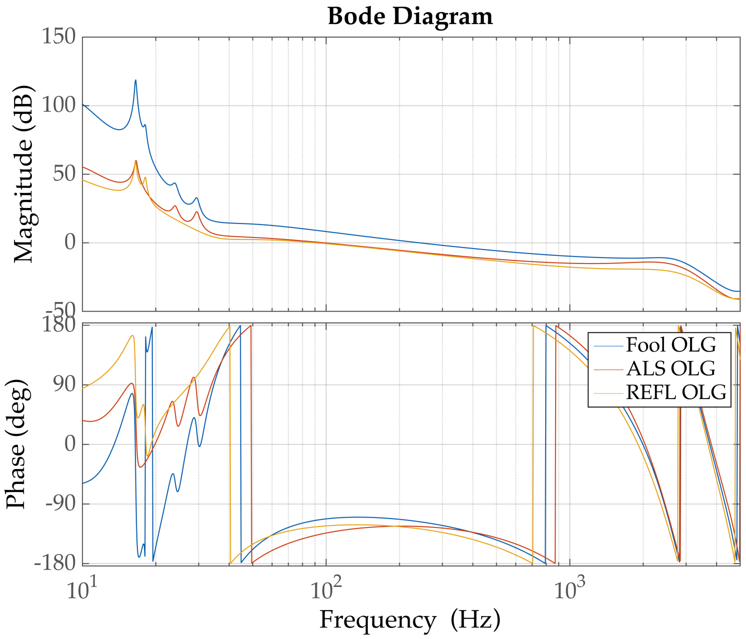

Figure 4.16: Open loop gain of feedforward decoupled loops. Shown are the individual loops, as well as the combined (“fool”) loop from Equation 4.25.

While this technique has not yet been implemented for the full PRFPMI at the 40 m Lab, we have tested it on a single arm cavity. Figure 4.17 shows the measured transfer function between the PDH control signal and the ALS error point. We see in the blue trace that the two paths to the summing node (through $A_{\mathrm{REFL}}P_{\mathrm{ALS}}S_{\mathrm{ALS}}$ or through $D_{\mathrm{cpl}}$) are well matched since the magnitude of the ratio is very close to 0 dB. The green trace shows that we get more than 30 dB isolation almost everywhere below 100 Hz [70].

We find that it is possible to increase the gain of the Refl loop, and the system stays stable. Conveniently, if one of the mirrors is given an impulse large enough that the PDH signal is no longer valid, the ALS loop maintains control of the cavity. The cavity comes back to resonance and is once again controlled by the PDH loop within about 40 ms, as seen in Figure 4.18 [71].

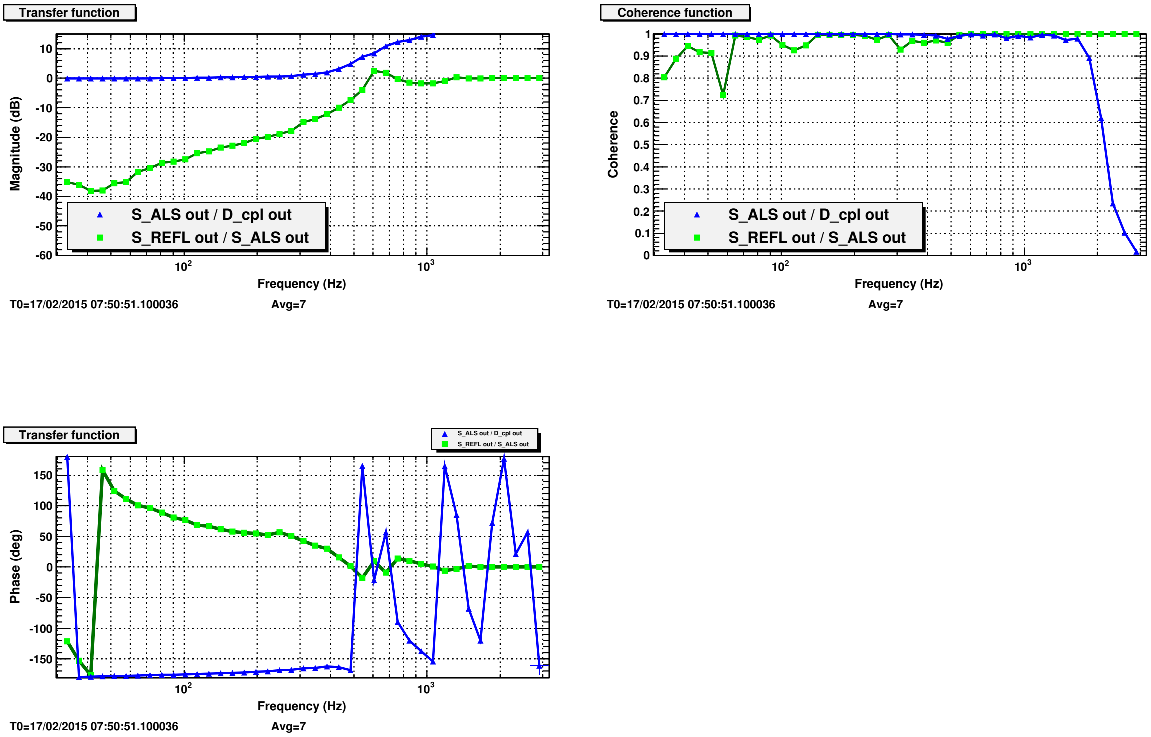

Figure 4.17: Measured loop decoupling. We see in the blue trace that the 2 paths to the summing node (through $A_{\text{REFL}}P_{\text{ALS}}S_{\text{ALS}}$ or through $D_{\text{cpl}}$) are well matched since the magnitude of the ratio is very close to 0 dB. The green trace shows that we get more than 30 dB isolation almost everywhere below 100 Hz.

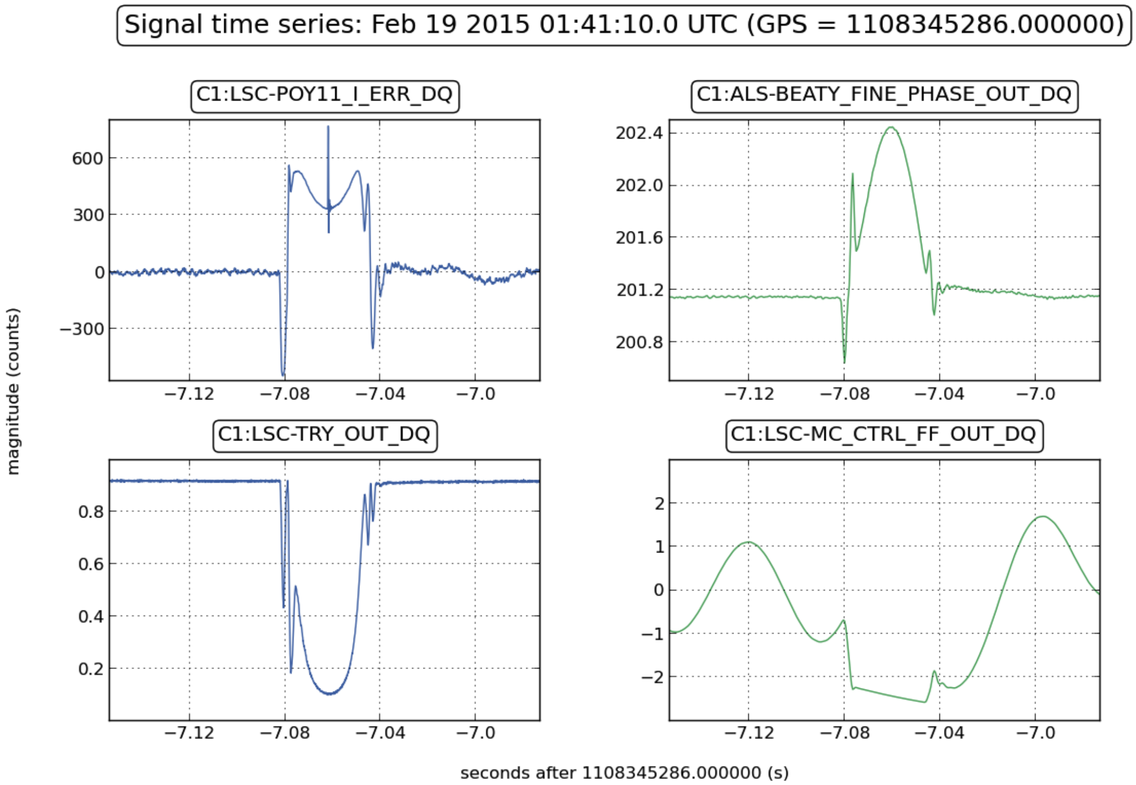

Figure 4.18: Impulse response of decoupled loop system, for single Fabry-Pérot cavity. ETMY is given an impulse that would normally unlock the PDH-only system; however, the ALS loop maintains control, and the cavity is relocked within about 40 ms [71].

## 4.5 Characterization of the 40m interferometer

Here we present various characterization measurements of the 40 m interferometer. Figure 4.19 shows a comparison of the measured open loop gain of the DARM loop, as compared to the modeled expectation. Figure 4.20 shows a measurement of the CARM open loop gain, near the ~15 kHz unity gain frequency. Figure 4.21 show calibrated error ("in") and control ("out") spectra of these two loops. Table 4.3 and Table 4.4 compare the optical and interferometer cavity designs for Advanced LIGO, and the 40 m interferometer.

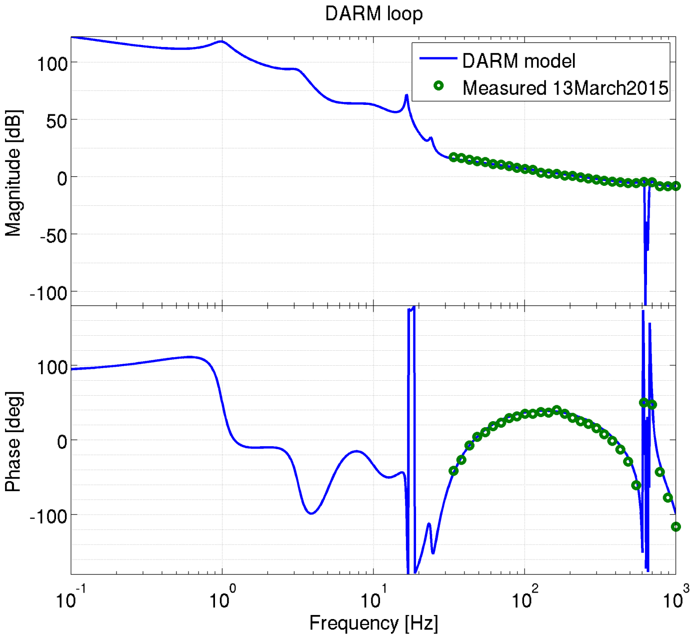

Figure 4.19: Measured vs. modelled DARM loop for the 40 m's PRFPMI configuration.

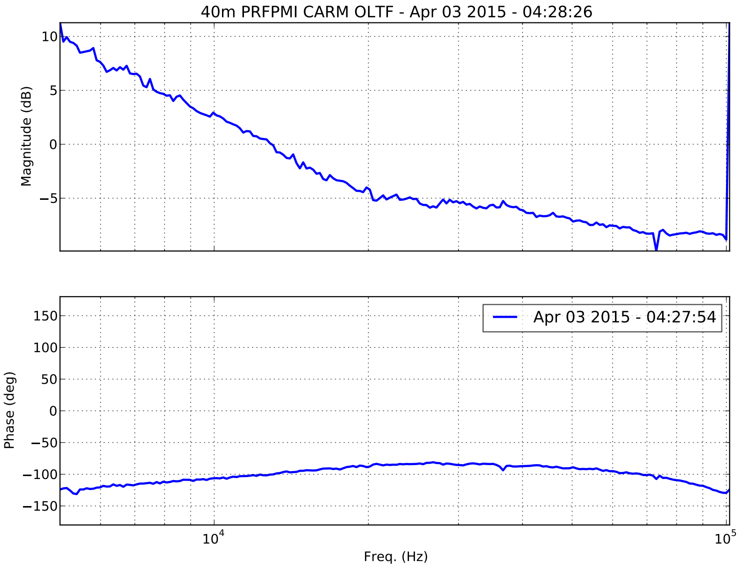

Figure 4.20: Measured CARM loop, near unity gain frequency, for the 40 m's PRFPMI configuration.

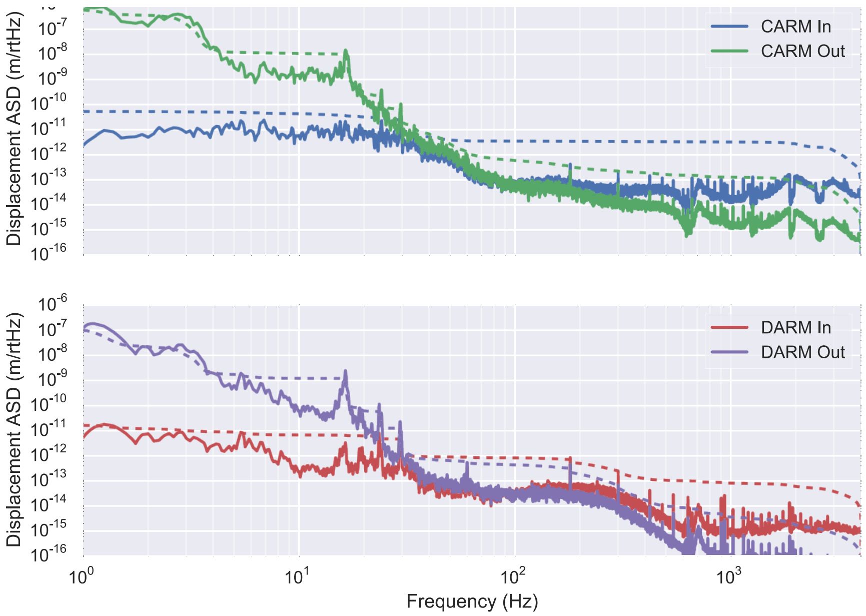

Figure 4.21: CARM and DARM error and control spectra for the 40 m's PRFPMI configuration, without compensation for filters in use.

| | aLIGO design | 40m design | 40m measured | | 40m detuned RSE |
|---|---|---|---|---|---|
| BS | | | | | |
| RoC | > 300 km [72] | flat, > ± 5,625 m [73] | 10,320 m [74] | | 10,320 m [74] |
| HR T (1064 nm) | 50 ± 0.5 % [75] | 50 ± 1 % [76] | | | 50 ± 1 % [76] |
| AR R (1064 nm) | < 50 ppm [75] | 600 ± 100 ppm [76] | | | 600 ± 100 ppm [76] |
| ITM | | | | | |
| RoC | 1934, -5, +15 m [77] | > 5,625 m [78] | X = 83.1e3 m [79] | | >4000 m [39] |
| | | | Y = -883.4e3 m [79] | | |
| HR T (1064 nm) | 1.4 ± 0.007 % [80] | 1.4 ± 0.05 % [81] | 1.384 % [79] | | 0.5% [39] |
| AR R (1064 nm) | < 50 ppm [80] | 500 ± 100 ppm [81] | 417 ppm [79] | | 600 ± 100 ppm [82] |
| HR T (532 nm) | < 1 % [80] | 1.0 - 2.5 % [81] | 1.094 % [79] | | n/a |
| AR R (532 nm) | not specified | < 1000 ppm [81] | 197 ppm [79] | | n/a |
| ETM | | | | | |
| RoC | 2245, -5, +15 m [83] | 57.37 ± 0.6 m [84] | X = 59.48 m [79] | | X = 56.3 m [39] |
| | | | Y = 60.26 m [79] | | Y= 56.8 m [39] |
| HR T (1064 nm) | 5 ± 1 ppm [85] | 15 ± 10 ppm [86] | 13.7 ppm [79] | | 10 ppm [39] |
| AR R (1064 nm) | < 500 ppm [85] | < 300 ppm [86] | 93.3 ppm [79] | | < 300 ppm [87] |
| HR T (532 nm) | 3 - 15 % [85] | 5.0 ± 3 % [86] | 4.579 % [79] | | n/a |
| AR R (532 nm) | 0.1 - 2 % [85] | < 1000 ppm [86] | 950 ppm [79] | | n/a |
| PRM | | | | | |
| RoC | -11.00 ± 0.11 m [88] | 115.5 ± 5 m [89] | 122.1 m [79] | | 348 m [39] |
| HR T (1064 nm) | 3.0 ± 0.1 % [90] | 5.75 ± 0.25 % [91] | 5.637 % [79] | | 7% [39] |
| AR R (1064 nm) | < 50 ppm [90] | < 300 ppm [91] | 160 ppm [79] | | < 300 ppm [92] |
| PR2 | | | | | |
| RoC | -4.56 ± 0.02 m [93] | flat, not specified | 413 m [94] | | n/a |
| PR3 | | | | | |
| RoC | 36.00 ± 0.17 m [95] | flat, not specified | 409 m [94] | | n/a |
| SRM | | | | | |
| RoC | -5.69 ± 0.06 m [96] | 142 ± 5 m [97] | 148.1 m [79] | | 365 m [39] |
| HR T (1064 nm) | 20 ± 1 % [98] | 10.0 ± 0.5 % [99] | 9.903 % [79] | | 7% [39] |
| AR R (1064 nm) | < 50 ppm [98] | < 300 ppm [99] | 67.6 ppm [79] | | < 300 ppm [100] |
| SR2 | | | | | |
| RoC | -6.43 ± 0.03 m [101] | flat, not specified | 413 m [94] | | n/a |
| SR3 | | | | | |
| RoC | 36.00 ± 0.17 m [102] | flat, not specified | 391 m [94] | | n/a |

Table 4.3: Parameters for 40m optics versus Advanced LIGO. “RoC” is radius of curvature of highly reflective side of optic. “HR T” is power transmission of highly reflective side of optic. “AR R” is power reflectivity of anti-reflection coating. The 40 m beam splitter (along with the input mode cleaner optics) was not replaced during the dichroic upgrade, and so some measured values are unavailable. PR2, PR3, SR2, and SR3 are all highly reflective for infrared light. PR2 and SR2 are also transmissive for green. Note that the values for PR2, PR3, SR2, and SR3 radii of curvature for the 40 m are all effective values, as seen by light transmitted through the substrate of the optic [103]. Some values in the “detuned RSE” column are design values, while others are measured. ETM RoCs there are average measured values. They are reported for comparison of the dichroic 40m with the former incarnation.

| | aLIGO design | 40m design | 40m measured | | 40m detuned RSE |
|---|---|---|---|---|---|
| RF modulation | | | | | |
| Primary frequency | 9.099474 MHz [104] | 11 MHz | 11.066 MHz | | 33.2 MHz [39] |
| Modulation depth | | | 0.190 ± 0.003 [105] | | |
| Secondary frequency | 45.497355 MHz [104] | 55 MHz | 5*primary | | 166.0 MHz [39] |
| Modulation depth | | | 0.2564 ± 0.0003 [105] | | |
| Arm cavities | | | | | |
| Length | 3994.5 m [104] | 37.795 m | X = 37.79 ± 0.05 m [106]<br>Y = 37.81 ± 0.01 m [106] | | X = 38.4583 ± 0.0001 m [39]<br>Y = 38.6462 ± 0.0003 m [39] |
| Free spectral range | 37.52 kHz [104] | | | | |
| Finesse (1064 nm) | 450 [104] | 450 | | | X = 1206 ± 10 [39]<br>Y = 1220 ± 3 [39] |
| Finesse (532 nm) | | 100 | | | n/a |
| Cavity pole (single arm, 1064 nm) | | | | | X = 1616 ± 14 Hz [39]<br>Y = 1590 ± 4 Hz [39] |
| g-factor | 0.8303 [104] | | | | X = 0.317 [39]<br>Y = 0.322 [39] |
| Power recycling cavity | | | | | |
| Length | 57.656 m [104] | 6.753 m [107] | 6.759 m [107] | | 2.257 m [39] |
| Free spectral range | 2.6 MHz [104] | | | | |
| Finesse | 114 [104] | | | | |
| Cavity pole | | | | | |
| g-factor | 0.8214 [104] | | 0.89 [94] | | |
| Recycling gain | | | 4 | | |
| Signal recycling cavity | | | | | |
| Length | 56.008 m [104] | 5.399 m [107] | 5.399 m [108] | | 2.151 m [39] |
| Free spectral range | 2.67 MHz [104] | | | | |
| Finesse | | | | | |
| Cavity pole | | | | | |
| g-factor | 0.8699 [104] | | 0.92 [94] | | |
| Michelson | | | | | |
| Schnupp asymmetry | 5.04 cm [104] | | 3.9 ± 0.2 cm [109] | | 45.1 cm [39] |

Table 4.4: Parameters for 40m interferometer versus Advanced LIGO. Note that some values, such as the primary RF modulation frequency, change with time, and so approximate values are quoted.

# Chapter 5

# Global seismic noise cancellation

The next generation of interferometers for gravitational-wave detection, including the Laser Interferometer Gravitational Wave Observatory (LIGO), will have unprecedented sensitivity to astrophysical events [5]. At low frequencies (~10 Hz) it is likely that the displacement noise of the suspended mirrors will be limited at the $10^{-22}$ m/$\sqrt{\text{Hz}}$ level by fluctuations in the Newtonian gravitational forces [110, 111, 112]. The sources of the fluctuations are density perturbations in the environment (e.g., seismic and acoustic) and mechanical vibrations of the nearby experimental apparatus. While care will be taken to mitigate the sources of all of these fluctuations, further reductions of this Newtonian noise may be made by carefully measuring the source terms and subtracting them from the data stream (offline) or in realtime on hardware by applying cancellation forces to the mirror.

In addition, direct seismic vibrations cause significant non-stationarity of the LIGO detectors. In particular, through nonlinear processes, seismic noise below 10 Hz has been shown to limit the performance of gravitational-wave detectors. Of the external disturbances that the control systems must overcome, one of the most problematic is the persistent and large 0.1 - 0.3 Hz double-frequency 'microseismic' peak, generated by wave activity in the oceans [113]. At the microseismic peak the amplitude of the ground motion is on the order of several $\mu$m/$\sqrt{\text{Hz}}$, but the level of seismicity present at the observatories varies widely (by a factor of ~10) [114]. The relative displacement of the test masses in this band is of the same order as the horizontal displacements caused locally by the microseism. These levels of seismic noise present a host of problems, including:

- Some amount of the translational ground motion is converted to angular motion of the mirrors due to cross-coupling in the mirror suspensions. These misalignments create fluctuations in the power circulating throughout the cavities, driving feedback loops towards instability or modulating the couplings between length degrees of freedom (e.g., MICH to DARM coupling). Several of the angular control loops have bandwidths of only a fraction of a Hz. Even in the event of single arm locks the full interferometer may not hold resonance due to angular motion in the power recycling cavity.

- Large control forces at low frequency could, for Initial and Enhanced LIGO, generate excess noise in the signal band through non-linearities in the mirror actuation. Examples of such non-linearities are saturation of the actuator electronics and Barkhausen noise in the ferromagnets used to actuate the mirrors. Non-Gaussian transients from these mechanisms pollute the gravitational wave data stream as background events, which is why Advanced LIGO no longer uses electromagnetic actuators on the final pendulum stage.

- Some amount of light is scattered from each mirror surface, and subsequently re-enters the readout path after reflecting off of vibrating surfaces. When the relative motion of the mirror and the other surfaces in the scattering path is larger than the laser wavelength ($\lambda \sim 1\ \mu$m), the phase noise introduced by this scattered light is experienced not only at the frequency of the motion but as broad-band noise up to a cutoff frequency determined by the relative velocity [115].

In order to mitigate these couplings, we remove seismic noise in real-time by applying forces to the active isolation tables that support the suspended optics. To determine how to transform the seismic information from an external witness sensor into an appropriate cancellation force, we utilize optimal Wiener filters [116]. Seismic isolation is regularly in use at the LIGO sites, but usually only in a local sense – cancellation forces are calculated using only seismometer information from a nearby sensor. Some work has been done in the past to use feedforward to reduce the overall motion between sets of mirrors, although that work had a high frequency technical limit of about 0.3 Hz [113]. Here we strive to expand upon that work to a broader range of frequencies and apply seismic noise cancellation to the interferometer in a global sense, since it is the distances between the mirrors that we require to be undisturbed, rather than individual mirror positions.

To demonstrate the efficacy of this noise cancellation technique for future use, we demonstrate below that offline subtraction of seismic noise can be done using static Wiener filtering [116] based on an array of seismic sensors. We demonstrate feed-forward vibration isolation on a suspended Fabry-Pérot interferometer using Wiener filtering and a variant of the common Least Mean Square (LMS) adaptive filter algorithm. We compare the experimental results with theoretical estimates of the cancellation efficiency. Using data from the LIGO 5th Science Run, we also estimate the impact of this technique on full scale gravitational wave interferometers. In the future, we expect to use this technique to also remove acoustic, magnetic, and Newtonian gravitational noise perturbations from the LIGO interferometers. This noise cancellation technique is simple enough to implement in standard laboratory environments and can be used to improve SNR for a variety of high precision experiments.

While the Wiener filtering occurs entirely in the time-domain, we examine plots in the frequency-domain of the filtered and unfiltered signals to determine the level of subtraction achieved. In Section 5.2 we describe the Wiener simulations on a cavity in our lab. In Section 5.3, we make similar estimates for one of the 4 km LIGO interferometers. In Section 5.4 and Section 5.5 we demonstrate the performance of real-time seismic noise cancellation systems on the Enhanced LIGO interferometers, and at the 40 m Lab. Section 5.6 discusses the extension of these techniques to adaptive noise cancellation systems.

## 5.1 Wiener filters

To find a linear filter that will improve a chosen signal, we must first define what it means to 'improve' the signal. We define an error signal

$$e(n) = d(n) - y(n), \tag{5.1}$$

where $n$ is the time index, $d(n)$ is the noisy "target" signal and $y(n)$ is the approximation of $d$ from the independent witness sensor. This is given by

$$\vec{y} = \vec{w}^{\mathrm{T}}\vec{x}, \tag{5.2}$$

where $x$ is the measurement of the external disturbance from the witness sensor, and $w$ is the finite impulse response (FIR) filter that we will solve for. The figure of merit ($\xi$) that we use for calculating the Wiener filter coefficients in this case is the expectation value of the square of the error signal,

$$\xi \equiv E[e^2(n)] = E[d^2(n)] - 2\vec{w}^{\mathrm{T}}p + \vec{w}^{\mathrm{T}}R\vec{w}. \tag{5.3}$$

Here, $E[*]$ indicates the expectation value of $*$, $\vec{p}$ is the cross-correlation vector between the witness and target signals, and $R$ is the autocorrelation matrix for the witness channels. When we find the extrema of Equation 5.3 by setting

$$\frac{d\xi}{dw_i} = 0, \tag{5.4}$$

we find

$$R\vec{w}_{optimum} = \vec{p}. \tag{5.5}$$

Equation 5.5 finds the time domain filter coefficients which minimize the RMS of the error $\vec{e}$ by optimizing the estimate of the transfer function between the witness sensors and the target signal. The error signal is now an estimate of the signal in $d$, without any noise. Since the matrix $R$ is of the block Toeplitz form, we take advantage of the Levinson-Durbin [117] method of solving problems of the form $\vec{b} = M\vec{a}$, where $M$ is a Toeplitz matrix. The Levinson method is considered weakly stable, as it is susceptible to numerical round-off errors when the matrix is close to degenerate (in our case, if two or more witness sensors carry nearly identical information about the noise source). However, for well conditioned matrices it is much faster than brute force inversion of the matrix [118].

## 5.2 Seismic noise cancellation applied to triangular ring cavity

At our 40 m prototype interferometer [39] lab at Caltech, both static Wiener filtering and adaptive filtering algorithms have been applied to the feedback signal of a suspended Fabry-Pérot triangular ring cavity (the mode cleaner described in Section 4.1). We have used two Güralp CMG-40T 3-axis seismometers [119] and several Wilcoxon 731A single-axis accelerometers [120] as our independent witness channels ($\vec{x}$ in Section 5.1), and the low-frequency feedback signal for the cavity length as the target channel ($\vec{d}$ in Equation 5.3) to reduce.

Figure 5.1 shows the locations of the witness sensors relative to the cavity mirrors. The mirrors of the cavity are suspended as pendulums with a resonance of ~1 Hz to mechanically filter high frequency noise, with the suspensions sitting on vibration isolation stacks to further isolate the optics from ground motion. The 'stacks' are a set of three legs supporting the optical table on which the mirror sits, with each of the legs consisting of alternating layers of stainless steel masses and elastomer springs [49, 121].

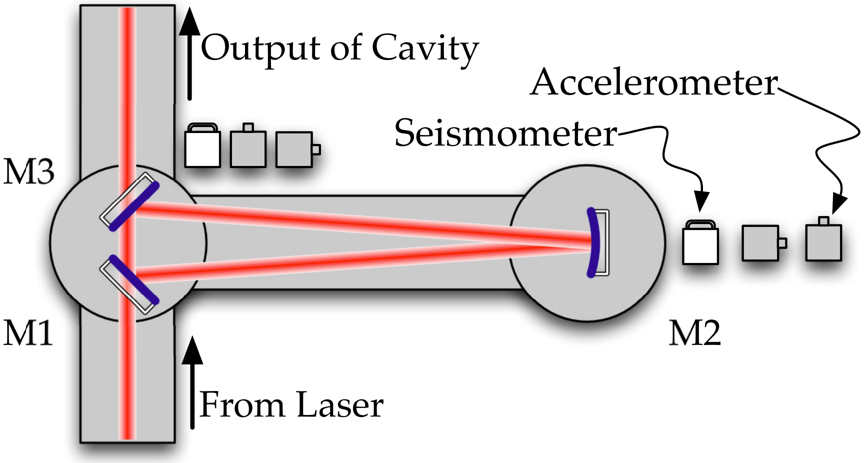

Figure 5.1: Locations of seismometers and accelerometers in relation to the cavity mirrors. Round trip length of the triangular cavity is 27 m.

We used Matlab [122] to import the data for the length feedback signal for our cavity, and to construct and apply the Wiener filters. Since the feedback control bandwidth is ≃50 Hz, the feedback signal can be used as an accurate measure of the seismic disturbance at low frequencies.

In Figure 5.2 we show results of a day-long simulation study. This study was done to determine the length of time we can use a set of static filters before updating. We use 1 hour of data to train and calculate a single Wiener filter, and then apply that filter to 10 minute segments of data for one day, using a 31 second long, 2000 tap filter with a sample rate of 64 Hz. In Figure 5.2b, we select a few typical traces to illustrate the capabilities of the filter, while in Figure 5.2a we show the full results as a spectrogram, whitened by normalizing to the spectra during the time the filter was being trained. We see large amounts of noise reduction both at the broad stack peak at ~3 Hz and around the 16 Hz vertical mode of the mirror pendula.

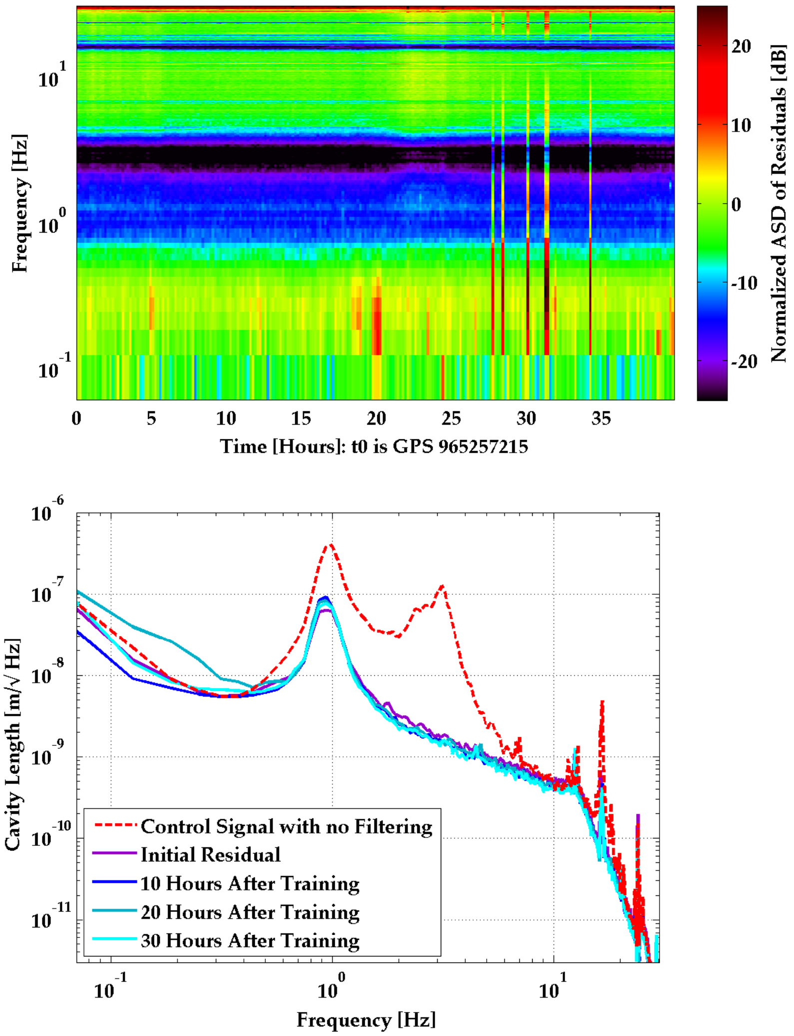

Figure 5.2: Result of offline seismic Wiener filtering on suspended triangular cavity. (a) Spectrogram showing the efficacy of a Wiener filter applied offline over a several hour period. Noticeably different traces between ~28 hours and ~34 hours are the result of non-stationary anthropogenic noise, not a decay of the filter's efficacy. (b) Amplitude spectral density of the control signal. Dotted red is without subtraction, purple is initial residual, and progressively lighter blues are 10 hours, 20 hours, and 30 hours after filter was trained.

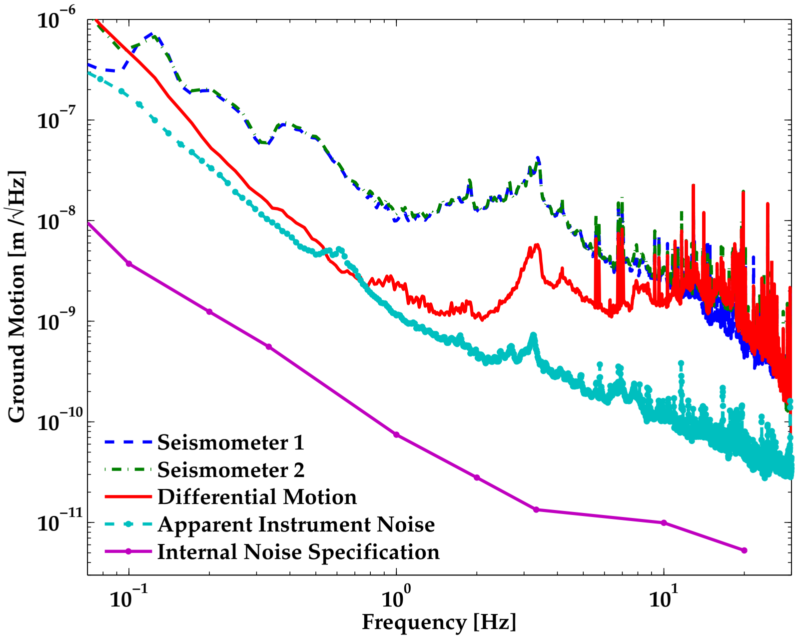

Figure 5.3: Shown are the spectra of the individual seismometers (blue dashed and green dash-dot), the manufacturer's spec for the seismometers' internal noise (purple solid-circle), and the differential ground motion along the 13.5 m length of the cavity (red solid). We also show the differential noise of the seismometers with the seismometers collocated in a stiff seismic vault (teal dash-circle); in principle, this is a measurement of the actual seismometer noise floor. It is unknown what uncorrelated noise is present in our sensors, making the teal trace so much larger than the specification.

We also include the noise contributions of our seismometers in Figure 5.3 to demonstrate how close we are able to get to the fundamental limit of Wiener filtering. Since the Wiener filter accepts, as inputs, the signals from the witness sensors (which have true ground motion information plus self-noise of the instruments and noise in the readout electronics), all of these noise contributions are filtered and added back into our data stream, limiting our ability to suppress ground motion below these levels. In Figure 5.3, we show that the differential ground motion over the length of the cavity is not much larger than the instrument noise of the seismometers. In other words, the ground noise over length of the cavity is strongly correlated below $\sim 1$ Hz and so the differential motion is much smaller than the motion of any individual sensor. Currently our measurement of the differential ground motion is limited by the apparent instrument noise of the seismometers, represented by the teal trace in Figure 5.3. The apparent instrument noise is significantly higher than the specification, which indicates that there is some unknown noise which is uncorrelated between two seismometers, even when they are placed very close together. We will use lower noise sensors and readout electronics and better thermal/acoustic isolation of the seismometers in order to get better performance on such short baselines.

The limit to the performance of the feed-forward subtraction seems to be a combination of low frequency noise in the seismometers and the feedthrough of noise from the auxiliary controls systems of the cavity (e.g. angular controls, pendulum damping servos, etc.).

## 5.3 Estimated impact for Enhanced LIGO

One of the LIGO sites in Livingston, Louisiana has had a hydraulic external pre-isolator (HEPI) actuation system installed since 2004 (the other LIGO site in Hanford, Washington received a HEPI system as part of the Advanced LIGO upgrade) [123]. This HEPI system is designed to actuate on the seismic isolation stacks that support the suspended LIGO optics to actively reduce seismic noise. Initial implementation of the HEPI actuators only included local seismic isolation between 0.1-5 Hz to reduce anthropogenic noise, tidal effects, and the microseism [42].

To estimate how the global Wiener filtering technique should scale up to a full size interferometer, we analyzed data from the 5th LIGO Science Run [5]. While this analysis was done as offline post-processing, results from later tests executed on the LIGO interferometers using the HEPI actuators during the 6th LIGO Science Run are discussed in Section 5.4.

Instead of a single cavity, in this case we explored the subtraction of seismic noise from the differential arm length feedback signal (which is an accurate measure of the low frequency ground noise). The sensors are placed close to the ends of the interferometer arms and at the beamsplitter as shown in Figure 5.4.

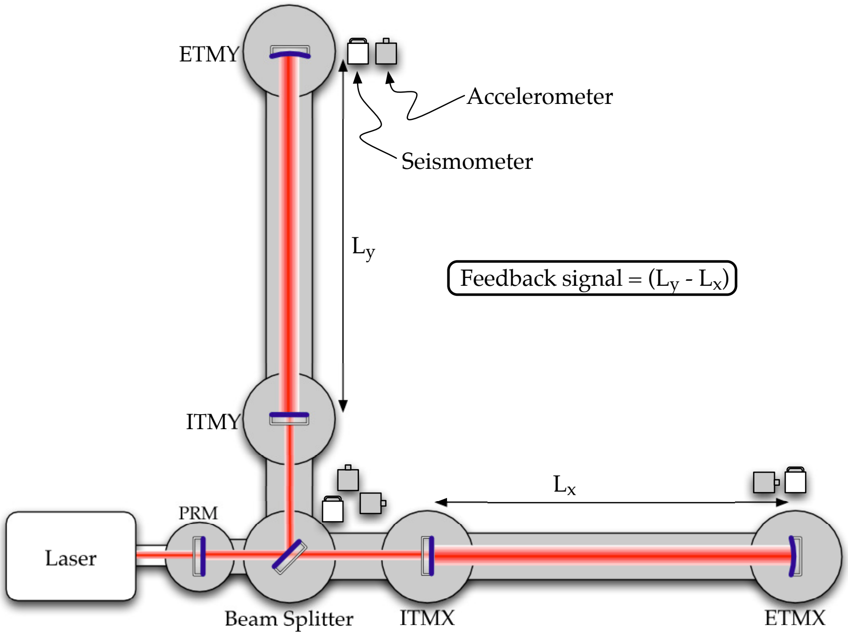

Figure 5.4: Schematic layout of seismometers and accelerometers relative to interferometer mirrors.

Figure 5.5 shows the resulting subtraction efficacy for a static filter. The variation in the 0.1-0.3 Hz band comes from variation in the ambient level of the double frequency microseismic peak [124]. The structure in the 1-15 Hz band is the usual increase in anthropogenic noise during the workday.

Even though some excess noise is added in the dips around 3-5 Hz and 7 Hz, the filter reduces the main contributors to the RMS of the control signal, and the reduction is remarkably stable over the 30 day timespan. This static filter does inject an unacceptable amount of noise above 20 Hz, which we will eliminate in the future by using more aggressive pre-weighting to disallow such noise amplification before calculating the Wiener filter.

Figure 5.6 shows the subtraction if we use an acausal filter, retraining it every 10 minutes, for the same 30 day data set. This filter performs much better than the static version. While we cannot apply an acausal filter in real-time, we can utilize causal adaptive filters such as those discussed in Section 5.6 to achieve nearly the same effect as long as the seismic environment does not change appreciably on time scales less than 10 minutes.

Residuals for both Figure 5.5 and Figure 5.6 were calculated using 46 second long Wiener filters of 3000 taps at a sample rate of 64 Hz.

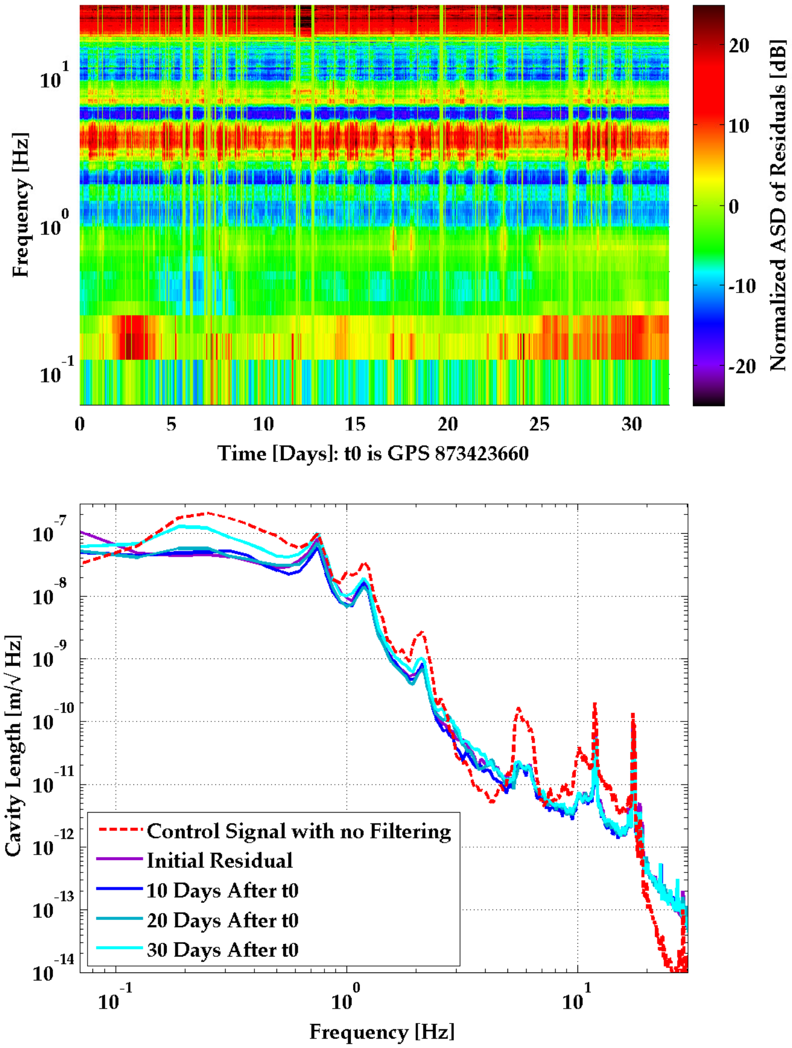

Figure 5.5: Result of offline simulated seismic Wiener filtering on the 4km LIGO Hanford interferometer. (a) Traces are amplitude spectra normalized to the unfiltered control signal (red trace in b), which is at a time during the filter's training. Filter was trained on 6 hours of data, then applied in 10 minute segments. Vertical stripes indicate times when the interferometer was not operational. Seismic subtraction is fairly constant on a one month time scale, although it is not particularly effective for times when seismic noise is significantly different from the training time. (b) Selected individual spectra from (a) above. Dotted red trace is before subtraction, purple trace is initial residual, and progressively lighter blues are 10, 20, and 30 days after the filter was trained.

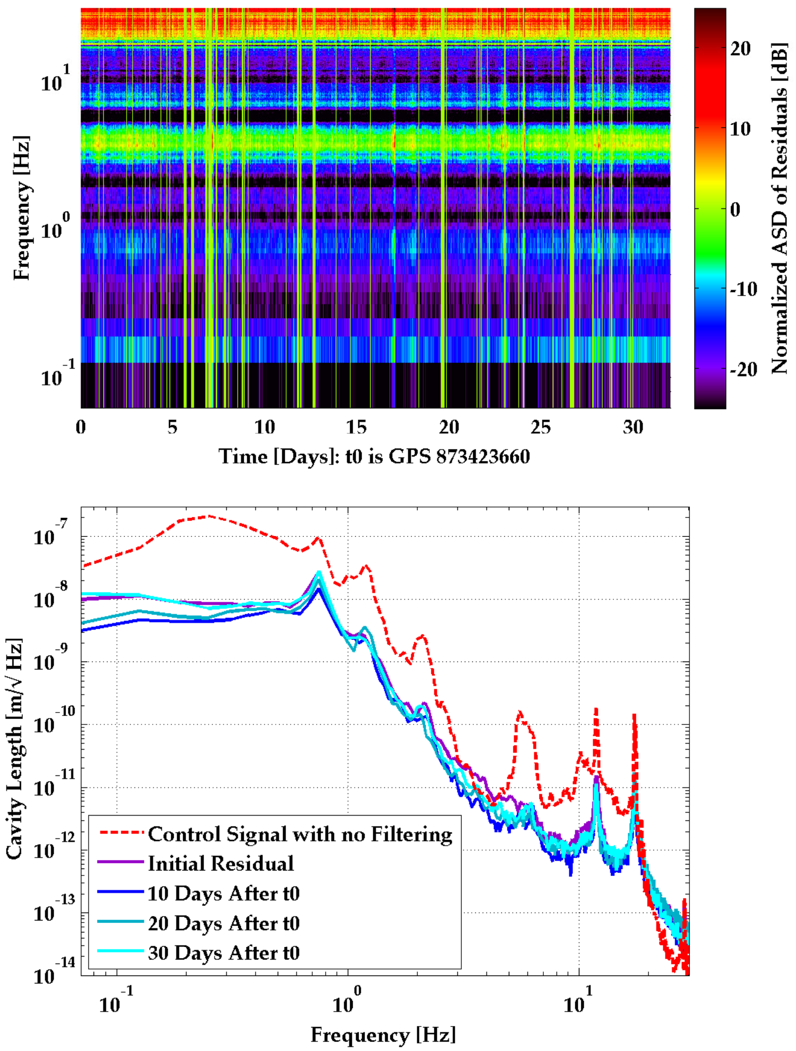

Figure 5.6: Result of offline simulated seismic Wiener filtering on the 4km LIGO Hanford interferometer, using an acausal filter on the same 30 day data set. (a) Traces are amplitude spectra normalized to the unfiltered control signal (red trace in b). A filter is trained on, and then applied to 10 minute segments of data. Seismic noise is more effectively suppressed using this constantly updated filter, implying that the transfer function is changing on a relatively short time scale, and that it is advantageous to update the filter more often than once per month. (b) Selected individual spectra from (a) above. Dotted red trace is before subtraction, purple trace is initial residual, and progressively lighter blues are 10, 20, and 30 days after beginning.

## 5.4 Static implementation for Enhanced LIGO

In the winter of 2009-2010, this static global seismic noise cancellation was implemented at the LIGO sites during the 6th Science Run. The FIR Wiener filters were estimated in MATLAB [122] using the Levinson-Durbin algorithm [117, 118], and then fit to a set of IIR coefficients to reduce the computational time in the real-time control system. The conversion from FIR to IIR was done with the Vectfit [125] software package. The intermediate FIR filters were composed of 1000's of taps and processed data with a sample rate of 64 Hz, allowing for subtraction down to a few 10's of mHz. Approximately 1 hour of data was used to train each filter. Each test mass is in the vicinity of at least one seismometer (see Figure 5.7), which measures motion in three perpendicular directions, X, Y, and Z, where Z is vertical and the two horizontal directions are roughly aligned to the arms of the interferometer. The length control is coupled most strongly to motion along the axes of the respective cavities, but cross-couplings to the other directions are non-negligible. Regardless of the direction of the witness signal generating a feed-forward correction, the actuation is applied in the direction in which the laser beam is propagating for that particular chamber.

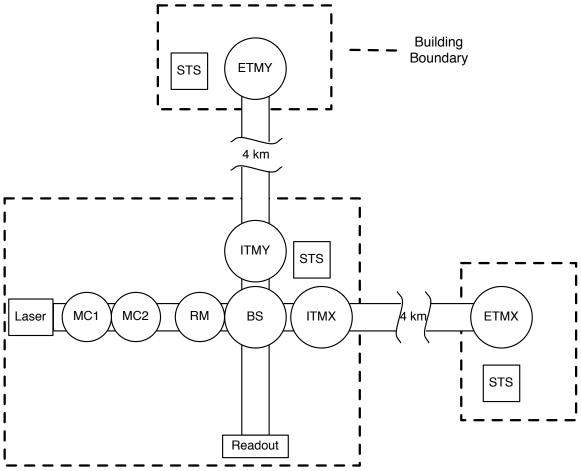

Figure 5.7: Location of seismometers and vacuum chambers. ETMX/Y and ITMX/Y are the locations of the end and input test masses, respectively. RM is the power recycling mirror, BS the beamsplitter, MC1 and MC2 the mode cleaner chambers, and STS refers to a Streckeisen STS-2 3-axis seismometer. At LHO the end stations were equipped with multiple single-axis Geotech GS-13 single-axis seismometers instead of STS-2s.

Offloading control signals to actuators located in the external seismic isolation systems mitigates several of the problems mentioned in the introduction to Chapter 5. In order to properly subtract the filtered witness signals, the transfer function from our point of actuation to the cavity control signal must be measured and divided out. A diagram showing the relevant pieces of the mechanical structure can be seen in Figure 5.8. The number of mechanical components separating the mirrors from the ground creates a complicated transfer function with many resonant features.

At LLO, the active seismic isolation system is HEPI (hydraulic external pre-isolator) [42, 126]. A

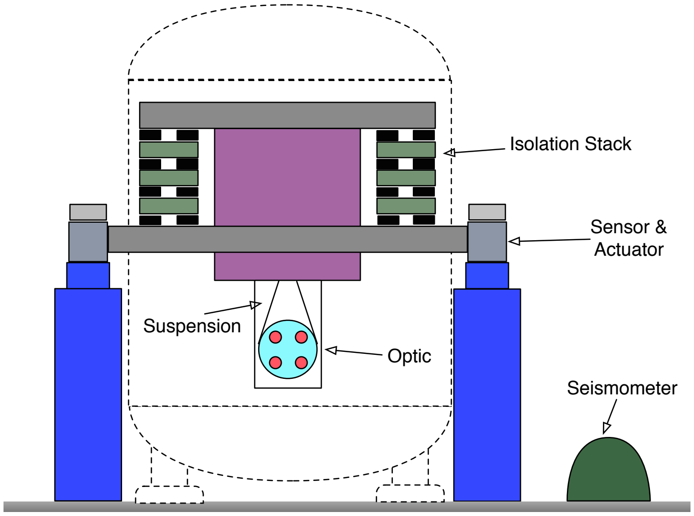

Figure 5.8: Cartoon of Initial / Enhanced LIGO seismic isolation system, with single-loop suspension and passive table isolation stack. The interferometer optic is suspended from the isolation stack. Signals from the ground seismometer are applied just below the stack via piezo (LHO) or hydraulic (LLO) actuators as indicated in the figure. The circles on the optic represent the four electromagnetic actuators on the back of the mirror. The dashed line represents the vacuum chamber.

similar system installed at LHO at the time of this work, PEPI, used piezoelectric actuation. HEPI is an Advanced LIGO (aLIGO) system which is now in place at both LLO and LHO. Figure 5.9 shows a schematic of the HEPI system [42]. The Wiener filter feed-forward path was added in parallel to the existing sensor correction path, which contained hand-tuned filters designed to match HEPI's position sensors to local seismometers on the ground. Typically this sensor correction reduced the differential motion sensed by the suspended cavities to $\frac{1}{10}$ of the ground motion in the microseism band. All results shown are improvements on top of this existing isolation. The HEPI actuators [127] provide the ability to move its payload by ± 700 $\mu$m; the maximum range of PEPI was ± 90 $\mu$m. We excited these actuators while the detector was locked and monitored the mirror motion to measure the mechanical response. We again used Vectfit to fit these measurements. An example of these transfer function measurements can be seen in Figure 5.10. An example Wiener filter can be seen in Figure 5.11. Both the Wiener filter and mechanical response are only fit accurately up to a few Hz. Low-pass filtering was applied to the feed-forward signal to prevent noise injection at higher frequencies.

In the winter of 2009-2010, we began commissioning of feedforward paths for the other interferometric control signals, starting with the power recycling cavity control at LLO. This work was performed during a time of particularly high microseismic activity, which impaired the Livingston detector's ability to remain in lock. As can be seen in Figure 5.12 (a), a significant reduction in control signal was achieved, with the RMS being reduced by more than a factor of two. As a result there was also a reduction in the power fluctuations in the recycling cavity, which tended to drive

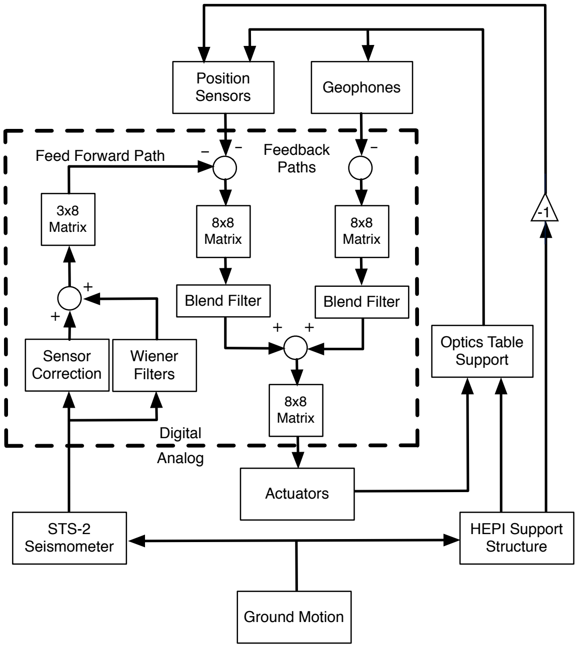

Figure 5.9: The external isolation systems measure the motions of their payloads using position sensors and geophones, providing feedback via hydraulic or piezoelectric actuators. Ground motion is also measured with a nearby seismometer, which is digitally processed and subtracted from the position sensor signal, a process called sensor correction. For chambers containing test mass optics this signal is then blended with the geophone signal into a so-called 'super-sensor' which has good combined sensitivity over a wider frequency band than either sensor individually. These 'super-sensor' signals are fed to the actuators. The presence of conversion matrices indicate a transformation between the different coordinate bases used by particular sensors/actuators. The use of four horizontal and four vertical position sensors/geophones allows for the sensing and control of eight degrees of freedom: three translational, three rotational, and two over-constrained. This schematic is adapted from Figure 3.8 of Wen [42].

the interferometer control loops towards instability, and greatly improved the instrument's duty cycle. This cavity is nearly geometrically degenerate [128] (g-factor ~1) and therefore is especially sensitive to misalignments.

Sending feed-forward signals developed to minimize one cavity control signal may inject noise into other length controls, since the cavities share mirrors. To avoid this problem, feed-forward paths were implemented in a serial fashion, such that extra motion injected into other degrees of freedom could be corrected by the Wiener filters constructed for subsequent feed-forward paths. For example, Figure 5.12(b) shows the effect of filters designed to minimize the short Michelson control signal, calculated on top of the existing power recycling feed-forward path. Above ~0.2 Hz the Michelson signal is reduced, with the improved isolation of the power recycling cavity mostly preserved, albeit with some noise injection in the PRC control spectrum around 1 and 2 Hz, as well as a sharp feature in between those frequencies (see the thick red trace in Figure 5.12(b)).

Global feed-forward for the differential arm length control was implemented at both obser-

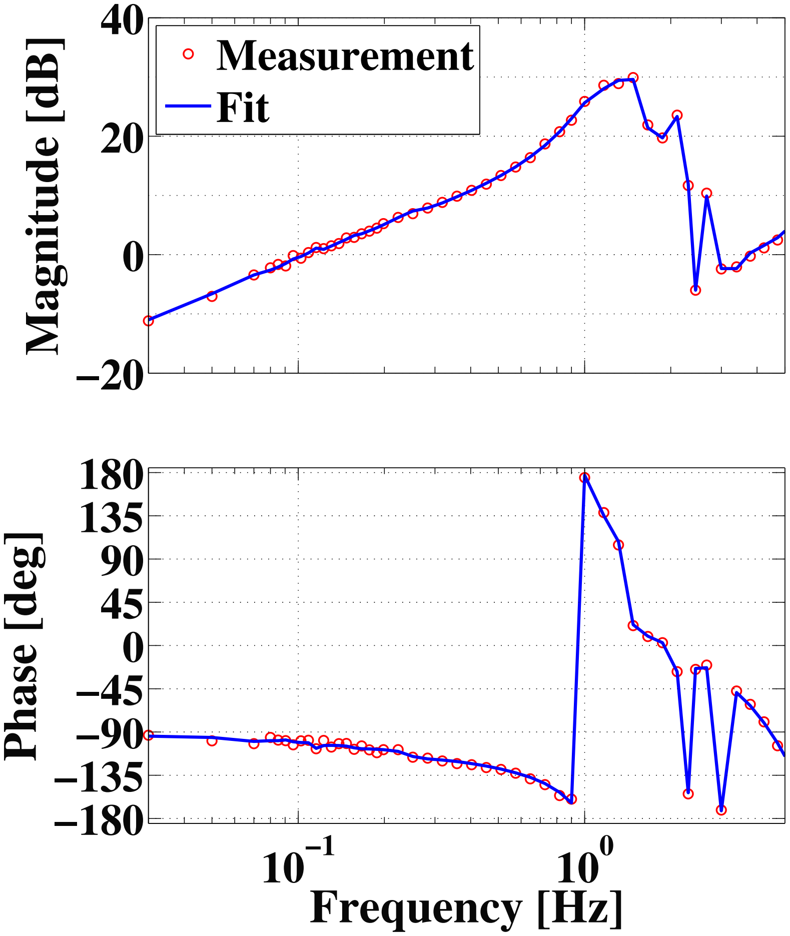

Figure 5.10: Example plant transfer function, fit using Vectfit. The average statistical error of points shown was ~2.5% in magnitude and ~ 1.5° in phase.

vatories with similar results. Since the DARM degree of freedom is sensitive to the motions of mirrors in both the end buildings, as well as the corner station, seismometers in all three locations were incorporated into the calculation of the optimal feed-forward filter. Unlike the cavities in the corner station, the arm cavity's baseline is significant compared to the seismic wavelength at these frequencies, which are around 7 km at the Livingston observatory for example[113]. As such the cavity does not experience much common mode rejection of the seismic motion between the two buildings where its mirrors are located, and may even experience differential motion greater than either building's local seismic noise. The sensor correction scheme successfully reduced the differential cavity motion by using these seismometers as the low-frequency inertial reference, and the Wiener filter feedforward scheme improved this isolation by incorporating the cavity feedback signal's information about low-frequency differential motion. Figure 5.13 and Figure 5.14 show that the overall RMS of the control signal is reduced by a factor of ~2.5 for each interferometer by implementing the new scheme. For LLO this improvement in performance was balanced by

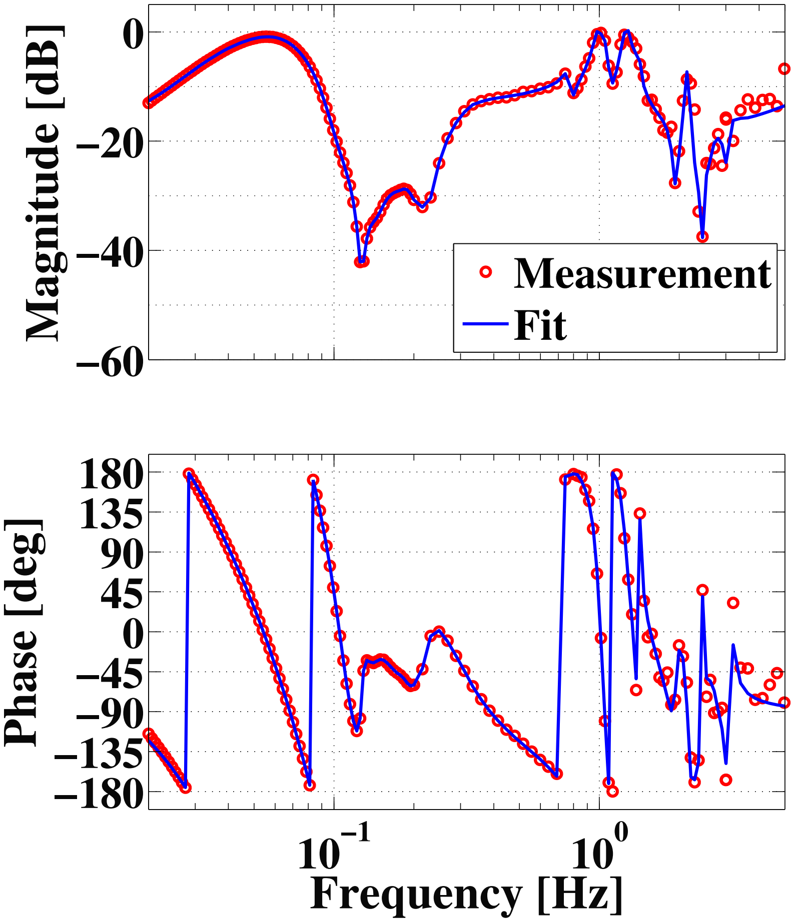

Figure 5.11: Example Wiener filter, corrected for the mechanical plant response, also fit using Vectfit.

a slight noise increase above 1 Hz. Some noise was also injected by Wiener filter feed-forward at LHO, below 0.1 Hz, and above 3 Hz.

The couplings of ground motion to the feedback signal can potentially be time dependent. Since the approach used here is not adaptive, this could lead to a degradation of the filter's subtraction efficacy over time. Figure 5.15 shows the performance of DARM feed-forward at LLO when first implemented and eight months later. While the overall reduction in RMS motion was originally a factor of ~2.5 the same feedforward filters provided ~20% less isolation after eight months, reducing the RMS motion by a factor of two. There is no obstacle to retraining new Wiener filters to potentially recover subtraction performance; however, if there is a change in the mechanical plant remeasuring the transfer function is a time consuming process, due to the high accuracy required.

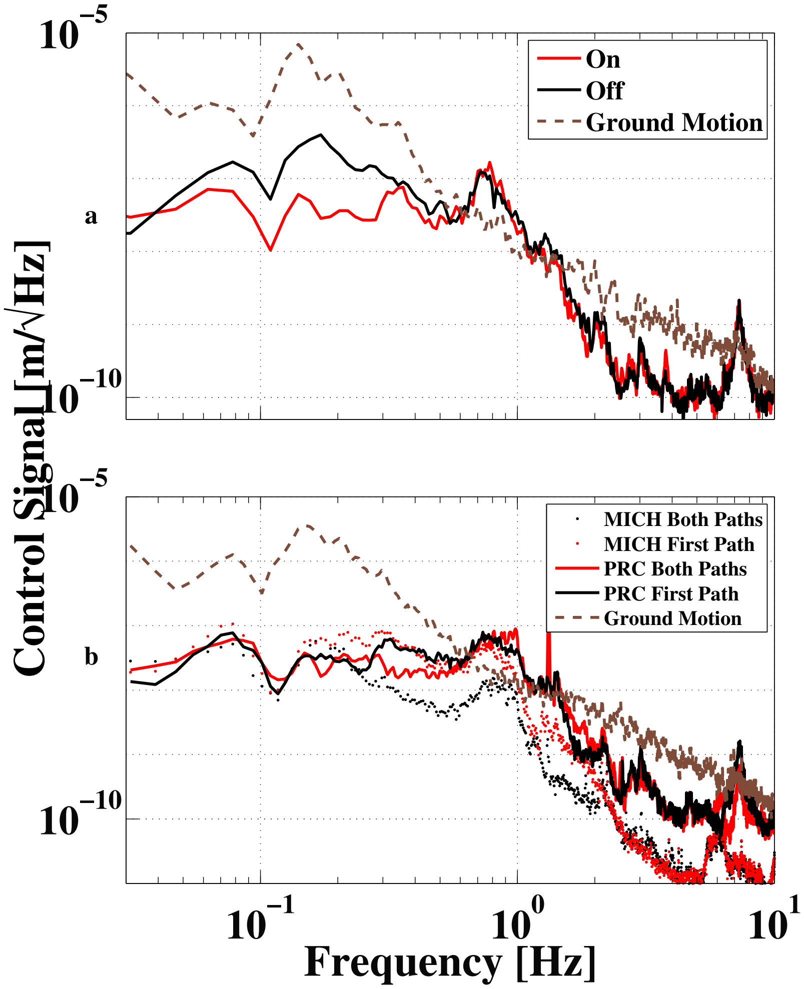

Figure 5.12: The upper plot shows the reduction in power recycling cavity control signal achieved using Wiener filter feedforward, improving the isolation of the relative mirror motion to $\simeq \frac{1}{100}$ of the local ground motion. The lower plot shows the implementation of Wiener filter feedforward on the short Michelson cavity and power recycling cavity simultaneously.

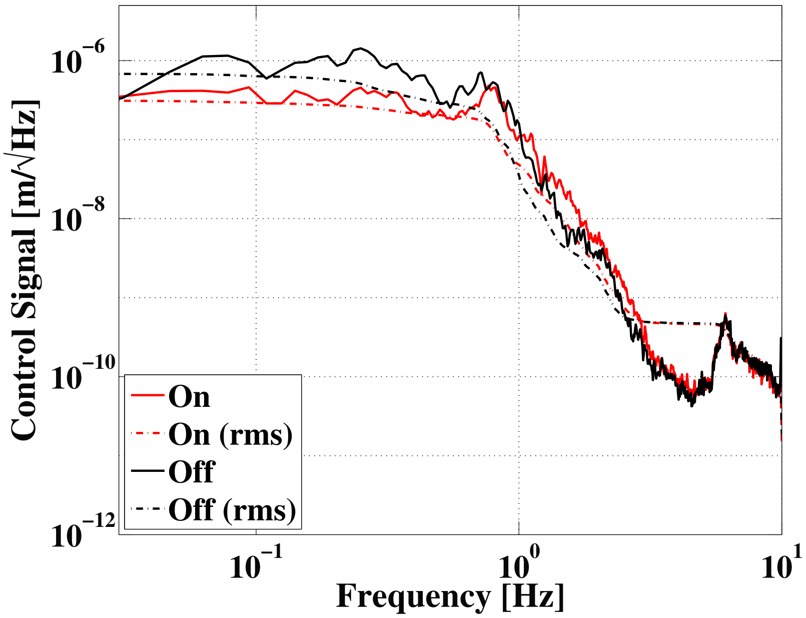

Figure 5.13: Reduction in differential arm length control from Wiener filter feedforward signals at LLO. The overall RMS of the feedback control was reduced by a factor of almost three, with some noise injection above 1 Hz.

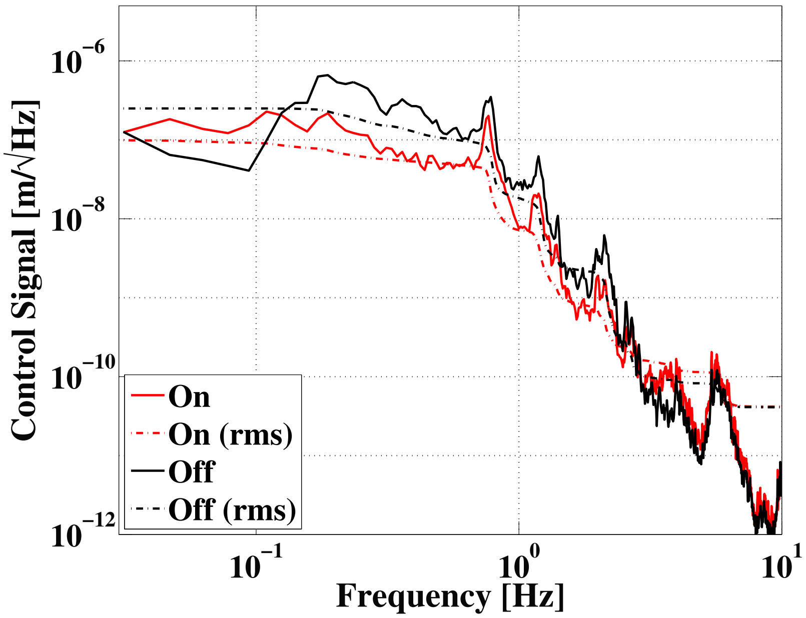

Figure 5.14: Reduction in differential arm length control from Wiener filter feedforwad signals at LHO. The RMS is reduced in a simliar fashion to LLO, with some excess noise at a few Hz and below 100 mHz.

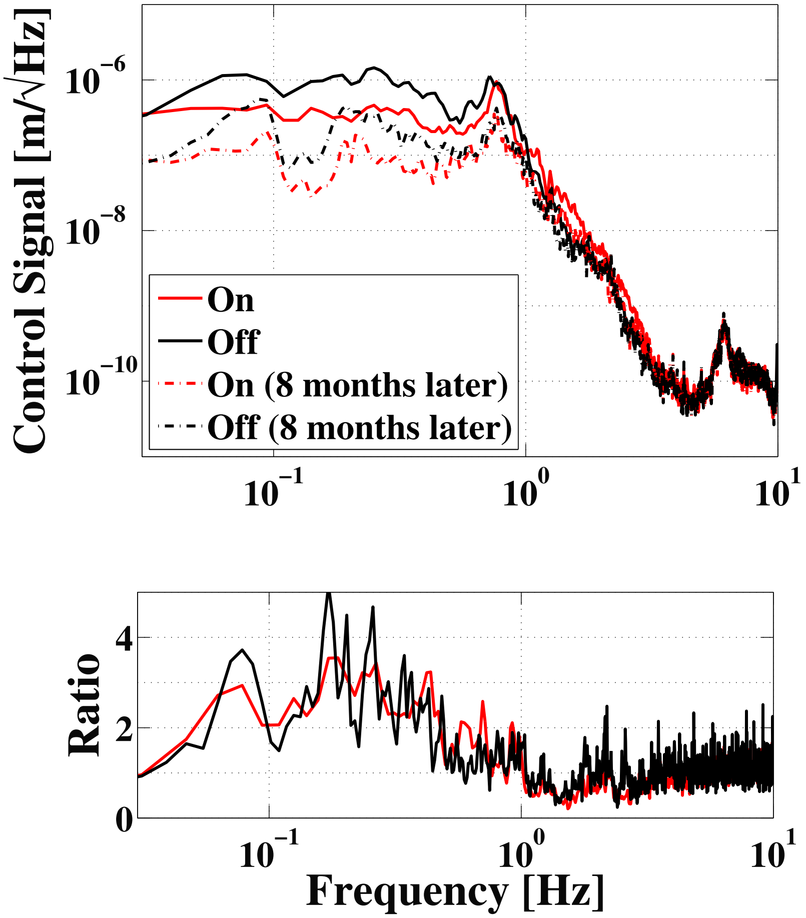

Figure 5.15: Performance of DARM Wiener filter feedforward at the time of filter training (upper traces) and 8 months later (lower traces). The two sets of spectra are from times with significantly different seismic input. The lower plot show the ratio of feed-forward off/on for the two time periods.

## 5.5 Cancellation of seismically-induced angular motion in power recycling cavity

To help alleviate the problems of angular motion in the power recycling cavity discussed in Section 4.3, feedforward was applied in real-time at the 40 m Lab. We measure the angular motion of the cavity axis with a quadrant photodiode (QPD) located at the POP port (see Figure 4.1). Note that this is done in the PRMI-only configuration (end mirrors misaligned) so that only the RF sideband is resonant in the cavity, and the QPD signal is not contaminated with information about the long arm cavity axis motion. We utilize a Trillium T-240 3-axis seismometer [129], mounted on a granite block on the ground next to the power recycling cavity, as the witness sensor for the ground motion.

As with Section 5.4, we must pre-filter the witness sensor data by the actuator transfer function. In this case, we measure this transfer function by actuating on the power recycling mirror and reading the QPD signal. This transfer function and the fit to the transfer function are shown in Figure 5.16. Since the accuracy of the measurement and the fit can both limit the feedforward subtraction (see Section A.2 for details), the coherence for each data point in Figure 5.16 is above 0.95, and for most of the points is above 0.99.

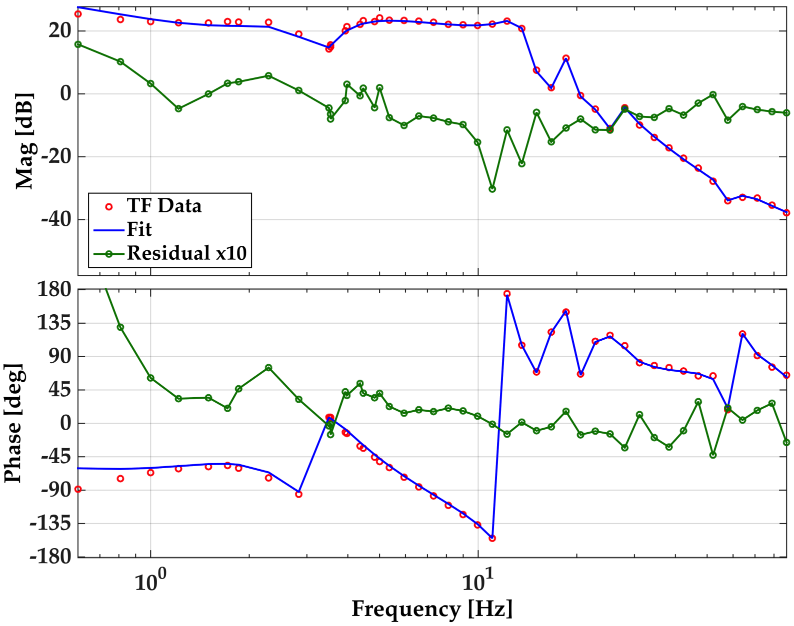

Figure 5.16: PRMI pitch actuator transfer function, actuating on PRM and reading out at the POP QPD. Red dots are measured data points with coherence of at least 0.95, the blue trace is the fit to the transfer function, and the green trace indicates the residual mismatch between the measurement and the fit.

We then calculate the Wiener filters for each seismometer axis. Figure 5.17 shows an example

Wiener filter that was calculated, for the seismometer axis in the same direction as the cavity axis in the power recycling cavity. The red dots show the ideal calculated filter, and blue trace shows the fit to that filter that will be used in the real time system for the actual feedforward. The green trace indicates the residual mismatch (multiplied by a factor of ten for clarity) between the calculated filter and the fit. The coherence between the seismometer witness signal and the QPD target signal was above 0.5 for the range $0.7\,\mathrm{Hz} < f < 4\,\mathrm{Hz}$, and above 0.9 for the sub-range $1.2\,\mathrm{Hz} < f < 3.5\,\mathrm{Hz}$. The Wiener filter's fit is weighted by the coherence, so that frequencies of high coherence are exceptionally well fit, at the expense of frequencies of lower coherence (where the signal from this witness channel is not important). In addition to these filters, each witness channel was filtered with a 10 Hz elliptic lowpass to avoid injection of sensor noise into the system. It was found that a 0.01 Hz highpass filter was also required, although that may be due to a lack of accurate actuator measurements at lower frequencies. Future iterations of this feedforward will investigate this feature.

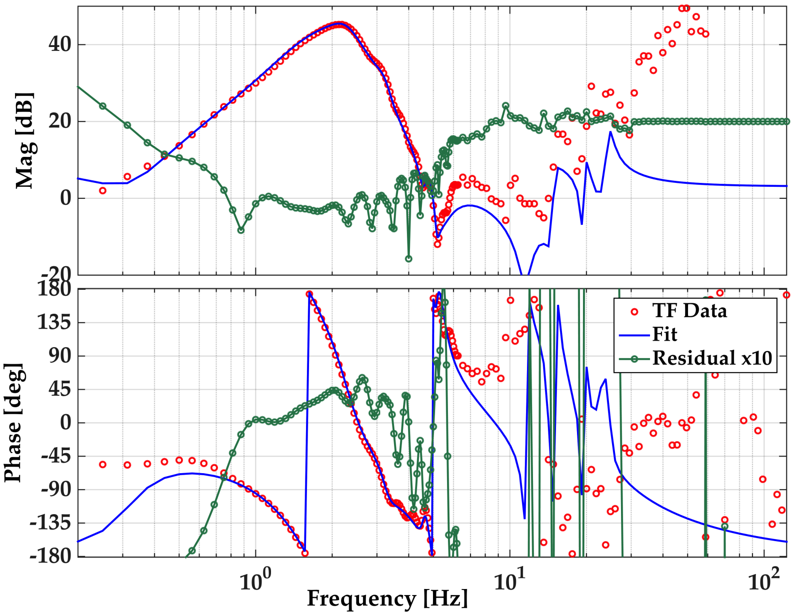

Figure 5.17: Example Wiener filter for PRC angular feedforward [130]. The red dots show the ideal calculated filter, and blue trace shows the fit to that filter that will be used in the real time system for the actual feedforward. The green trace indicates the residual mismatch (multiplied by a factor of ten for clarity) between the calculated filter and the fit.

Figure 5.18 shows the results for the pitch degree of freedom of this feedforward applied in real-time. Both pitch and yaw feedforward were applied simultaneously. For the pitch degree of freedom, only the two horizontal axes of the seismometer are used, but for the yaw degree of freedom all three axes are used. Red traces are without the feedforward, black traces are with the feedforward on, where both sets of measurements were taken within a few minutes of each other. RMS values are also shown.

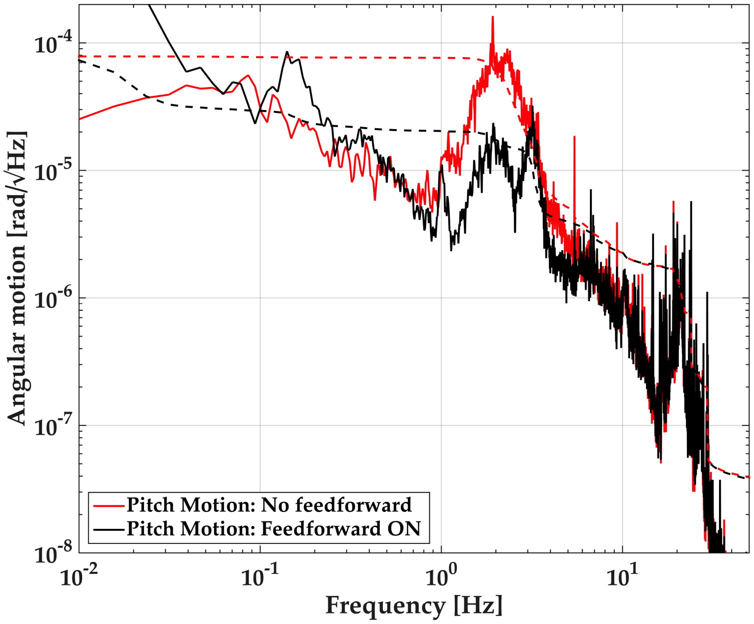

Figure 5.18: Suppression of angular motion in PRMI [131]. Red traces are without the feedforward and black traces are with the feedforward on. Both sets of data were taken within a few minutes of each other. RMS values are also shown.

Figure 5.19 shows our measured subtraction factor (blue trace) compared to the predicted noise suppression (orange trace). The predicted trace assumes noise-free actuators, as well as perfect fitting of the calculated Wiener filters. Here, a number greater than one indicates noise suppression while a number less than one indicates noise injection. We note that extra angular noise is injected below approximately 0.05 Hz. This is likely related to the need for the 0.01 Hz highpass filter, and will be investigated further in future measurements.

Interestingly, the residual intensity noise (RIN) of the cavity shown in Figure 5.20 decreases over a much broader band than the direct angular motion. It is expected that the power in the cavity should be more stable when the cavity axis motion is reduced. The fact that the RIN improvement is over a broader band implies that there are non-linear couplings present, such as scattered light effects.

While these results are shown in the PRMI-only configuration, they are most helpful when locking the full PRFPMI interferometer. Since the 3f locking signals are so sensitive, it is very challenging to maintain lock of the vertex degrees of freedom long enough to fully transition CARM and DARM. However, with this angular noise suppression, we are able to maintain lock of the vertex degrees of freedom quite easily for several minutes at a time, which enabled the CARM

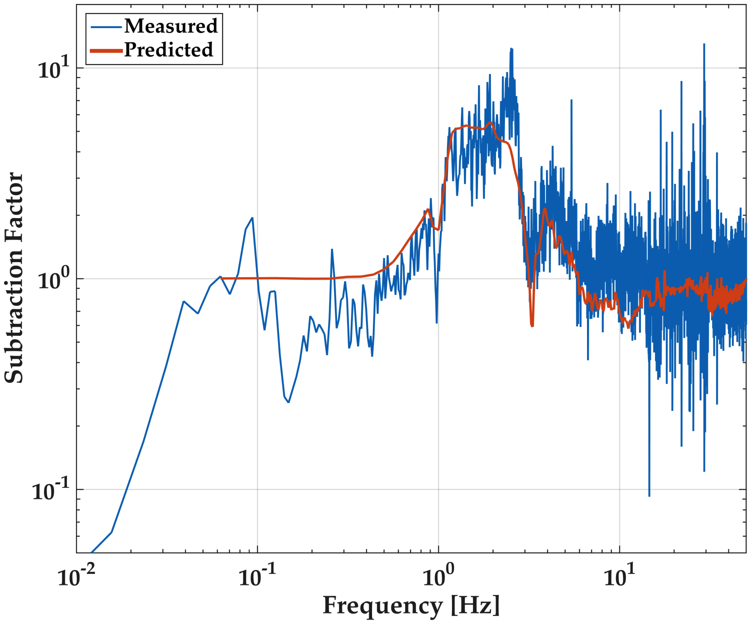

Figure 5.19: Actual versus predicted noise suppression for pitch degree of freedom. Measured trace is the ratio of the traces from Figure 5.18. A number greater than one indicates noise suppression while a number less than one indicates noise injection.

and DARM transitions and full lock of the 40 m interferometer.

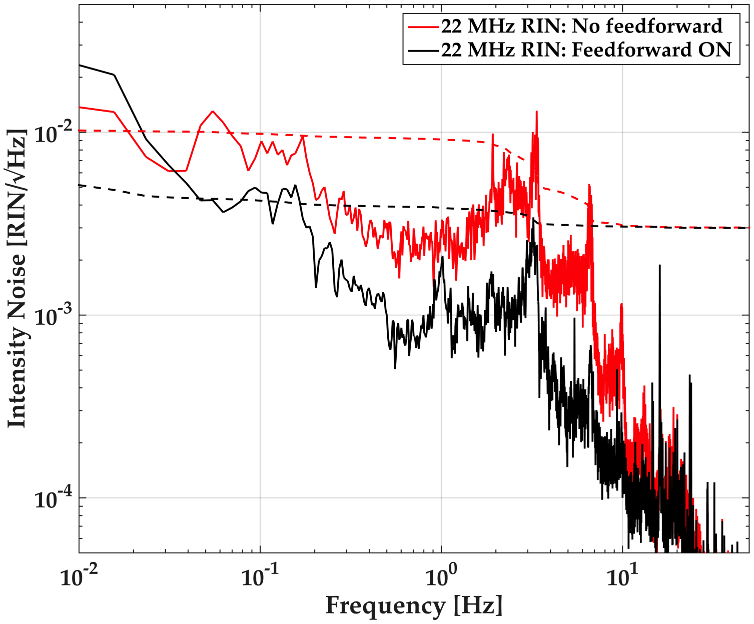

Figure 5.20: RIN suppression due to improved angular noise in PRC. Red (no feedforward) and black (feedforward on) traces were measured at the same time as the traces in Figure 5.18. RMS values are shown with dashed traces.

## 5.6 Least mean squared adaptive noise cancellation

In case the transfer functions between the sensors and the target are changing with time, it would be useful to use a filter whose coefficients change with time. Such an adaptive filter could also take into account changes in the 'actuator'. The most simple and common implementation of an adaptive filter is the Least Mean Squares (LMS) algorithm [132, 133]. Figure 5.21 shows a block diagram of an LMS adaptive filter.

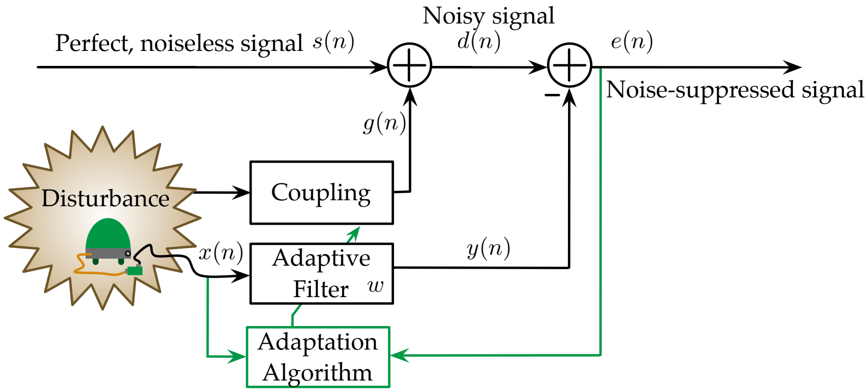

Figure 5.21: Block diagram of adaptive LMS filter.

The LMS algorithm is a canonical gradient descent algorithm. Much like Section 5.1, we define our error signal as $e(n) = d(n) - y(n)$, where $d(n)$ is the signal of interest at each time $n$, contaminated by some noise and $y(n) = wx(n)$ is the output of our adaptive filter if $x(n)$ is a witness to the noisy disturbance and w represents our filter.

To derive the update law that will determine the filter coefficients at each step, we first find the derivative of the expectation value of the square of our error function, $E[e^2(n)]$, with respect to the filter coefficients. This gives us

$$\begin{aligned} \nabla_w E[e^2(n)] &= \nabla_w E[e(n)e^*(n)] \\ &= 2E[(\nabla_w e(n))\, e^*(n)] \end{aligned} \tag{5.6}$$

with

$$\begin{aligned} \nabla_w e(n) &= \nabla_w \left(d(n) - wx(n)\right) \\ &= -x(n). \end{aligned} \tag{5.7}$$

So,

$$\nabla_w E[e^2(n)] = -2E[x(n)e^*(n)]. \tag{5.8}$$

We increment the filter coefficients $w$ in the opposite direction of the gradient,

$$\begin{aligned} w(n+1) &= w(n) - \frac{\mu}{2}\nabla_w E[e^2(n)] \\ &= w(n) + \mu E[x(n)e^*(n)], \end{aligned} \tag{5.9}$$

where $\mu$ is a step-size parameter chosen to determine how large the change to $w$ will be at each step. Since we want to follow potential changes in the system, we will not assume that we know $E[x(n)e^*(n)]$ *a priori*. We estimate this by

$$E[x(n)e^*(n)] \simeq \frac{1}{N}\sum_{i=0}^{N-1} x(n-i)e^*(n-i) \tag{5.10}$$

using a time series of length $N$. Since this will be implemented in real-time with causal filters, out time series for these estimates will be of length $N = 1$. For such instantaneous estimates of the expectation value,

$$E[x(n)e^*(n)] \simeq x(n)e^*(n). \tag{5.11}$$

With this estimate, Equation 5.9 becomes

$$w(n+1) = w(n) + \mu x(n)e^*(n). \tag{5.12}$$

In the long-time limit, Equation 5.12 converges to the optimal Wiener filter in the mean [134].

Since we update the filter coefficients at each time step, the filter designed using Equation 5.12 will oscillate around the true optimum.

To investigate the stability of the filters and extra noise due to the non-perfect convergence to the ideal Wiener solution as a function of $\mu$, we must look at the eigenvalues of the autocorrelation matrix $R$, defined in Section 5.1. This can be done, for example, by looking at past data, which is sufficient for this purpose. If $\lambda_i$ is the $i$th eigenvalue of $R$, then it can be shown [135] that for stability, we require

$$0 < \mu \leq \frac{1}{3\sum_i \lambda_i}. \tag{5.13}$$

Also, the mean square error will have excess noise of a factor of

$$E[e^2]_{\mathrm{LMS}} \simeq \left(\mu \sum_i \lambda_i\right) E[e^2]_{\mathrm{Wiener}}, \tag{5.14}$$

more than the ideal Wiener case.

The algorithm implemented at the 40 m Lab includes a set of modifications to the most simple LMS case. We use a "leaky normalized filtered-x LMS algorithm" [136],

$$w(n+1) = w(n)(1-\tau) + \mu(n)x(n)e(n). \tag{5.15}$$

We use the *leaky* modification of including a factor of $(1-\tau)$ to allow the response of the filter to decay over time. This prevents transients in $\vec{x}$ from affecting the filter permanently. $\tau$ must be chosen to be large enough to provide stability against transients, but small enough that it does not compromise the filter's ability to converge. This is often chosen empirically. We use the *normalized* modification of the algorithm such that the size of $\mu$ relative to $\vec{x}$ does not change. While we examined the approximate stability in Equation 5.13 by looking at past data (under the assumption that the past data is representative of future data), we must actually satisfy the stability criterion for every point in time. For this, we have made $\mu$ time dependent in Equation 5.15, where

$$\mu(n) = \frac{\mu}{\vec{x}(n)^T \vec{x}(n)} \tag{5.16}$$

and it is $\mu(n)$ that must satisfy the stability criterion.

While both the leaky and normalized modifications are quite common, we also include the *filtered-x* variant of the algorithm. This variant acknowledges that there exist phase delays in the path of the target signal which cannot be approximated by the LMS method alone [137]. To account for these phase delays, we filter the incoming witness signals with filters approximating those in the target signal path. Once we have matched the delays in the two different paths, we implement the regular LMS optimization to find the coefficients we will use in our FIR filter.

We apply this technique (which we call Online Adaptive Filtering - the "OAF" system) to the same triangular cavity as in Section 5.2. Once again, we use the cavity length feedback signal as our targeted signal to minimize, and a similar layout of independent witness sensors as shown in Figure 5.1. Unlike Section 5.2 and Section 5.3 which were simulations using previously collected data, here we are actuating on the cavity in realtime. Figure 5.22 shows results using 250 second long, 4000 tap filters with $\mu = 0.2$ for the mode cleaner and approximately $\mu = 0.04$ for the arm cavities, and $\tau = 0$ at a 16 Hz sample rate. The on/off traces in the adaptive case are similar to estimates made in the static Wiener filtering case (Figure 5.2).

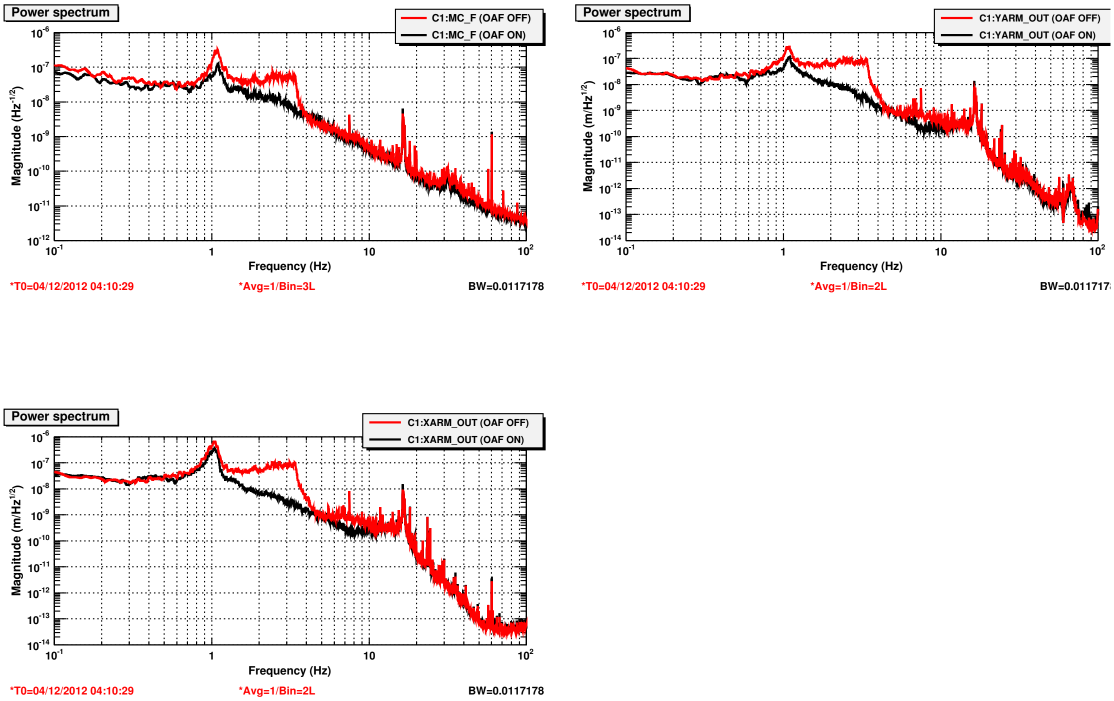

Figure 5.22: Online Adaptive Filter performance: the spectral density of the cavity length fluctuations are shown with the feed-forward on (lower black trace) and off (upper red trace) [138]. OAF was allowed to adapt simultaneously for all three cavities (two individually-locked arms and the mode cleaner).

# Chapter 6

# Newtonian gravitational noise cancellation

Fluctuations in the local Newtonian gravitational field present a limit to high precision measurements, including searches below a few tens of Hz for gravitational waves using laser interferometers. Early estimates of Newtonian noise by Weiss [139], Saulson [140], Hughes and Thorne [141], and Beccaria et al. [112] have made increasingly better estimates of the seismic environment and thereby, the gravitational noise. In this work, we update the estimations of Newtonian noise from the two LIGO sites. We also present a model of the perturbing gravitational field and evaluate schemes to mitigate the effect by estimating and subtracting it from the interferometer data stream, expanding on work pioneered by G. Cella [142]. Information about the Newtonian noise is obtained from simulated seismic data. The method is tested on causal as well as acausal implementations of noise subtraction. In both cases it is demonstrated that broadband mitigation factors close to ten can be achieved removing Newtonian noise as a dominant noise contribution. The resulting improvement in the detector sensitivity will substantially enhance the detection rate of gravitational radiation from cosmological sources.

The multi-stage vibration isolation systems [126, 143] developed for gravitational wave detectors should, in principle, be capable of reducing the direct influence of the ambient seismic noise to below the quantum and thermodynamic limits of the interferometers. Unfortunately, there is no known way to shield the detectors' test masses from fluctuating gravitational forces. As shown in Figure 6.1, our calculations estimate that the fluctuations in the local Newtonian gravitational field will be the dominant source of the mirror's positional fluctuations below 20 Hz. This noise source has been referred to as gravity gradient noise or Newtonian noise in previous literature.

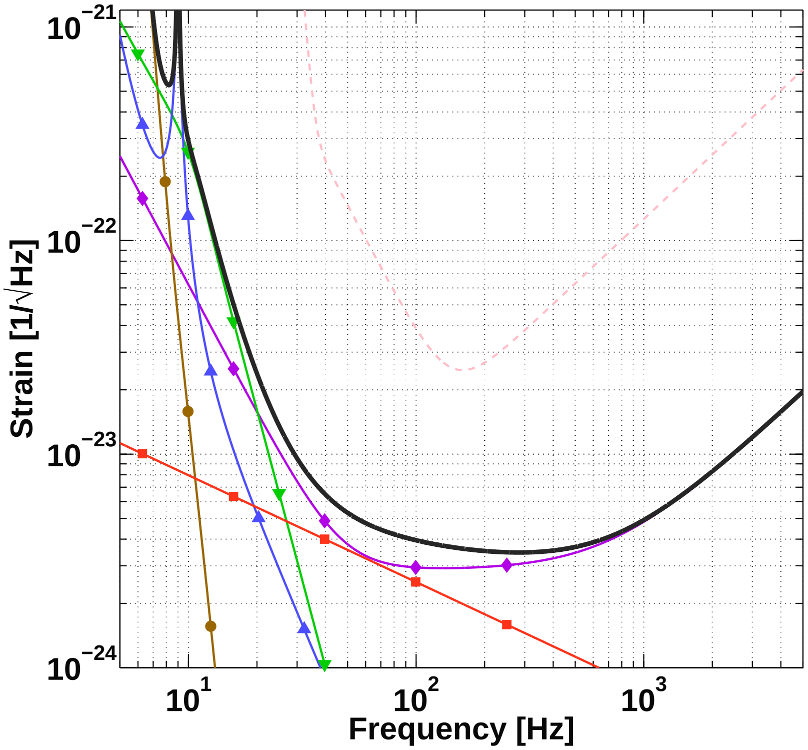

Figure 6.1: Strain noise spectral density of Advanced LIGO (black). The sensitivity of the initial LIGO (pink - dashed) is shown for comparison. The Newtonian gravity noise (green) is dominating the Advanced LIGO noise below 20 Hz. Other traces shown are other, non-gravitational, limits to the sensitivity: direct seismic vibrations (brown), quantum radiation pressure and shot noise (purple), mirror thermal noise (red), and mirror suspension thermal noise (blue).

## 6.1 Estimating Newtonian gravitational impact on Advanced LIGO

In 2011, several measurements were carried out at the LIGO sites to define a Newtonian noise budget [8]. Accelerometers were used to monitor vibrations on water pipes, near exhaust fans, on top of the buildings and on the walls. Sound spectra were measured with microphones inside and outside of the LIGO buildings. The resulting Newtonian noise estimates for each of these sources are summarized in Figure 6.2. In addition, the plot contains a representative noise model for potential upgrades to the advanced detectors such as Advanced LIGO and Advanced VIRGO, which we refer to as 3rd generation ground-based detectors [144]. Future detectors built at new sites, such as the proposed Einstein Telescope, we call 4th generation detectors [145]. More specifically, the strain noise model (excluding Newtonian noise) that we use to simulate interferometer noise

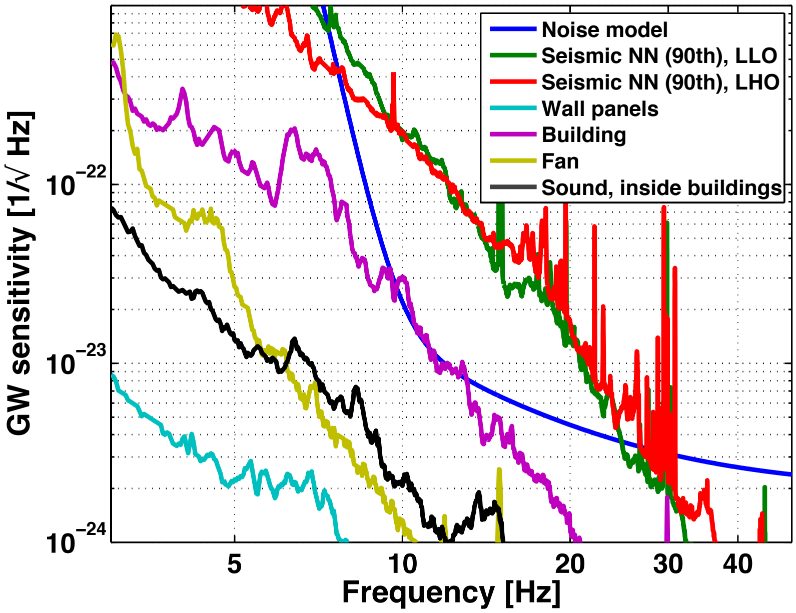

Figure 6.2: Seismic 90th percentile Newtonian noise estimates for the LIGO Livingston (LLO) and Hanford (LHO) sites (green and red lines), third generation strain noise model (blue line), and additional Newtonian noise estimates from vibrations of walls (cyan line), building tilts (magenta line), exhaust fans (beige line), and sound waves inside buildings (black line). Seismic contributions are the only Newtonian noise source significant for third generation detectors and earlier. Building tilt will be important for detectors beyond the third generation, but is not a dominating noise source at this time.

for a 3rd generation detector is:

$$\sqrt{S_h(f)} = \left(\frac{10^{-23}}{(f/10\,\mathrm{Hz})^{14}} + \frac{10^{-23}}{(f/10\,\mathrm{Hz})^{2}} + \frac{2\times 10^{-24}}{(f/40\,\mathrm{Hz})^{1/50}}\right)\frac{1}{\sqrt{Hz}}. \tag{6.1}$$

It is based on current best estimates of technology advance to mitigate instrumental noise such as thermal noise, seismic noise, and quantum shot noise. The Newtonian noise is simulated as test-mass displacement noise. To convert from the displacement noise of a single mass into strain noise, we multiply by two to account for the incoherent sum of four masses and then divide by the interferometer length, $L = 4\,\mathrm{km}$, to get strain.

All measured curves in Figure 6.2 are derived from 90th percentiles of spectral histograms similar to the one shown in Figure 6.7 for the seismic measurement at the LIGO Livingston site. Note that the seismic curves for both LIGO sites presented in Figure 6.2 are more recent, using a more accurate, non-averaged, analysis of the seismic percentiles, as compared to earlier measurements [8].

According to these estimates, seismic Newtonian noise is the only significant Newtonian noise contribution for 3rd generation and earlier detectors, so other contributions to Newtonian noise such as building vibrations or air pressure fluctuations are not considered in Section 6.2 and beyond.

Figure 6.2 shows our estimate of the strain noise due to several candidate Newtonian noise sources. We note that seismic Newtonian noise is likely to be the dominant source that we will see in aLIGO, with the tilt of the end station buildings the next most important source. We expect that other sources such as vibrations of the wall panels and the air handler fans will be less important. In addition to the sources considered for this plot, estimates of Newtonian noise from the chambers should be included as presented in LIGO-T070192 [146].

In the general case, if the displacement $\left(\vec{\xi}(\vec{r}, t)\right)$ of a mass is weak, then the test mass acceleration due to Newtonian noise can be estimate by the integral over the entire medium,

$$\delta\vec{a}_{\mathrm{NN}}(\vec{r}_0, t) = G \int \mathrm{dV}\, \frac{\rho_0(\tilde{\mathrm{r}})}{|\tilde{\mathrm{r}} - \tilde{\mathrm{r}}_0|^3} \left(\tilde{\xi}(\tilde{\mathrm{r}}, \mathrm{t}) - 3(\tilde{\mathrm{e}}_{\mathrm{r}} \cdot \tilde{\xi}(\tilde{\mathrm{r}}, \mathrm{t}))\tilde{\mathrm{e}}_{\mathrm{r}}\right), \tag{6.2}$$

where $\rho_0$ is the density of the medium, $G$ is the gravitational constant, $\vec{r}_0$ is the position of the mirror, $\vec{r}$ points to locations in the medium, and $\vec{e}_r$ is the unit vector pointing from $\vec{r}_0$ to $\vec{r}$ [147].

The following subsections will describe the measurements that we have done to obtain more information about each of these sources, as well as how we arrive at the estimate of the Newtonian noise contribution for each.

### 6.1.1 Vibration of outer walls

The LIGO buildings are constructed such that they have double walls, the inner of which is a structural wall, while the outer wall is made of thin sheet metal, with an approximately 1 meter gap between the two walls. The outer wall panels are attached in small sections — about 2 m wide — to a metal beam framework. To calculate the Newtonian noise, we assume that the vibrations of wall panels can be described as an incoherent sum over contributions from individual panel sections with coherence length equal to 2 m along both directions of the wall, each panel section vibrating in normal direction to the wall like a drumhead.

We calculate the contribution for a single wall panel that is very near the beam axis, and assume that if we include all of the sections on the front and back walls they will add to give some amount of cancellation, in addition to not vibrating in the same direction as the beam axis, so they will add to about the order of one single section. When considering both the front and back walls, we need to multiply by $\sqrt{2}$ if they add incoherently.

We choose to keep the model of the panel Newtonian noise simple since, as shown in Figure 6.2, the wall panel vibrations are not close to being a limiting source. We are confident that the estimated strain noise shown in Figure 6.2 is an overestimate of the wall vibration Newtonian noise, since it is calculated assuming that the entire wall panel vibrates with the maximum amplitude; however, the edges of each panel are bolted to the wall structure, which has displacement that is weaker by an order of magnitude than the maximum displacement of the center of a panel.

Since we have supposedly an overestimate of the noise, and this estimate is so far from being a limiting source, we will likely not look into more detailed measurements and simulations of wall panel vibrations in near-future investigations.

To calculate the strain noise due to the wall panels, we start as usual with the dipole approximation of Newtonian noise test mass acceleration:

$$\vec{a}_{NN} = G \int dV \rho \frac{1}{r^3} \Big(\vec{\xi} - 3\big(\vec{e}_r \cdot \vec{\xi}\big)\vec{e}_r\Big). \tag{6.3}$$

We assume that, for the panels close to the beam axis, $\vec{e}_r \parallel \vec{\xi}$, so $\big(\vec{e}_r \cdot \vec{\xi}\big)\vec{e}_r \sim \vec{\xi}$, and $\Big(\vec{\xi} - 3\big(\vec{e}_r \cdot \vec{\xi}\big)\vec{e}_r\Big) \sim -2\vec{\xi}$. We change from acceleration to displacement by dividing the Newtonian noise amplitude spectral density by $(2\pi f)^2$. We then multiply by $\sqrt{2}/L$ to get strain assuming that the wall panels are much closer to the test masses at the end stations, so we only include noise from the end stations. We also approximate the integral over volume as giving us the mass of a single wall panel being uniformly displaced. So we have

$$h_{NN} = G\,(\text{Mass of panel}) \frac{4\xi}{D^3} \frac{1}{(2\pi f)^2} \frac{1}{L}, \tag{6.4}$$

where $D$ is the distance between the wall(s) and the test mass at the end stations along the direction of the arm, $\xi$ is the measured wall displacement, and $L$ is the length of the LIGO arms.

We assume an individual wall panel has an area of $4\,\text{m}^2$, a thickness of 3 mm, and a density of $7000\,\text{kg/m}^3$. These values are only estimates, but they are sufficiently accurate to allow us to rule out the Newtonian noise from vibrating panels as being important. Figure 6.3 shows representative spectra of the wall panel vibrations, measured with a Wilcoxon 731-207 accelerometer [148]. The data from Figure 6.3, combined with Equation 6.4 results in the cyan curve of Figure 6.2.

### 6.1.2 Building tilt

From all possible types of motion of the end-station buildings, the tilt or rocking mode along the direction of the arms would produce the strongest Newtonian noise since the walls move in phase and Newtonian noise adds up from the two walls "in front of" and "behind" the test masses. So the conservative Newtonian noise estimate will take the total mass of the two walls and assume coherent displacement. Here, the Newtonian noise contribution of the tilt of the buildings is calculated in much the same way as the wall panels.

The coherent motion yields a factor of 2 in Newtonian noise from back and front wall. We also assume that the displacement of the walls increases linearly from 0 displacement at the ground to maximum displacement at the roof. This is based on the accelerometer spectra measured at the base and roof of the buildings, which show that displacement at the roof is much stronger. So to obtain an effective uniform displacement of the walls, we divide the spectra measured at the roof by 2. Finally, we make the additional assumption that this type of Newtonian noise is significantly

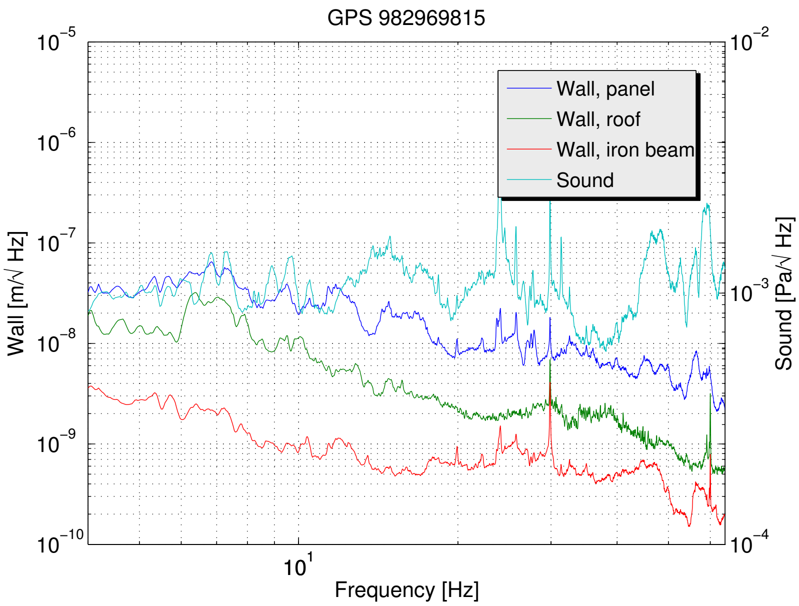

Figure 6.3: Various measurements of wall displacement at Livingston Y end station. 'Wall, panel' displacement is an accelerometer directly mounted to the outer wall sheet metal, near the ground. 'Wall, roof' is an accelerometer mounted to the outer wall (not panel), a few inches below the roof of the building. 'Wall, iron beam' is an accelerometer mounted to the iron beam structure of the outer wall to which the sheet metal panels are mounted.

weaker at the corner station since the buildings are larger and walls in the direction of the arms are further away from the test masses.

Then strain due to tilt of the buildings is

$$h_{NN} = \sqrt{2}G\,(\text{Mass of wall})\frac{2\xi}{D^3}\frac{1}{(2\pi f)^2}\frac{1}{L}. \tag{6.5}$$

Figure 6.3 shows spectra of accelerometer data, where the stronger motion measured by the roof accelerometer indicates that tilt could indeed be the dominant mode of the building. This noise spectrum, combined with Equation 6.5, results in the magenta curve of Figure 6.2. We use 25 tons (2.5e4 kg) as the mass of a single wall of the end stations, and we use 5 meters as the distance between the walls and the end test masses. This estimated total mass of each wall is a potentially bad guess, and since Newtonian noise from buildings is sufficiently close to aLIGO sensitivity, a more detailed calculation should be carried out based on a better estimate of the total mass of the wall.

### 6.1.3 Air pressure fluctuations

Perhaps one of the previously most concerning Newtonian noise candidates is air pressure fluctuations (since it is very difficult to subtract them if significant). According to our measurements, it appears that sound and air pressure fluctuations are not going to be a limiting noise source for aLIGO.

We start with the same Newtonian noise acceleration as calculated for seismic compressional body waves, since the physics is the same:

$$a_{NN} = \frac{8\pi}{3} G\rho\xi. \tag{6.6}$$

We then convert to strain amplitude using the same $1/(2\pi f)^2$ and $2/L$ terms considering incoherent contributions from both ends of each arm. Since sound data are pressures and not displacements, we need to convert:

$$\delta\rho = -\rho_0\left(\vec{\nabla} \cdot \vec{\xi}\right) \tag{6.7}$$

$$\delta\rho = -\rho_0\left(\vec{k} \cdot \vec{\xi}\right) \tag{6.8}$$

$$\delta\rho = -\rho_0\left(k\xi\right). \tag{6.9}$$

So

$$\xi = \frac{\delta\rho}{\rho_0}\frac{1}{k} = \frac{\delta\rho}{\rho_0}\frac{c}{2\pi f}, \tag{6.10}$$

where the density change is linked to the measured pressure change via

$$\frac{\delta\rho}{\rho_0} = \frac{1}{\gamma}\frac{\delta p}{p_0}, \tag{6.11}$$

where $\gamma = 1.4$ is the adiabatic constant of air. We finally obtain the strain noise

$$h_{NN} = \frac{8\pi}{3} G\rho_0 \frac{\delta\rho}{\rho_0}\frac{c}{2\pi f}\frac{1}{(2\pi f)^2}\frac{2}{L}. \tag{6.12}$$

We use 100 kPa as the average air pressure $p_0$, and air mass density $\rho_0 = 1.3\,\mathrm{kg/m^3}$. The speed of sound is $c = 330\,\mathrm{m/s}$.

Since the sound inside and outside of the buildings have comparable spectral densities, we will integrate Newtonian noise all the way to the chambers, in contrast to the calculation in Teviet Creighton's paper [149] where the noise is only integrated outside of the buildings, leading to larger suppression (although suppression is not significant at 10 Hz). Integrating the Newtonian noise from sound waves over large volumes is certainly simplistic since sound waves do not propagate freely due to building walls, trees, chambers, etc., but this is okay since it leads to an overestimation of the noise.

Figure 6.4 shows representative spectra from inside the X endstation building at the Hanford site.

Note that the LIGO microphones (in this plot represented by ”Microphone, BSC9”) are sufficiently sensitive down to ∼ 3 Hz; however, below 3-4 Hz, the response of the LIGO microphones decreases. This data, combined with Equation 6.5, results in the black curve of Figure 6.2.

The strain curve shown in Figure 6.2 is calculated using microphone data, and so cannot be trusted below 3-4 Hz. At very low frequencies we will need to utilize infrasound sensors; however, these are unnecessary for the ∼10 Hz regime where the Newtonian noise contributions are likely to limit future LIGO detectors.

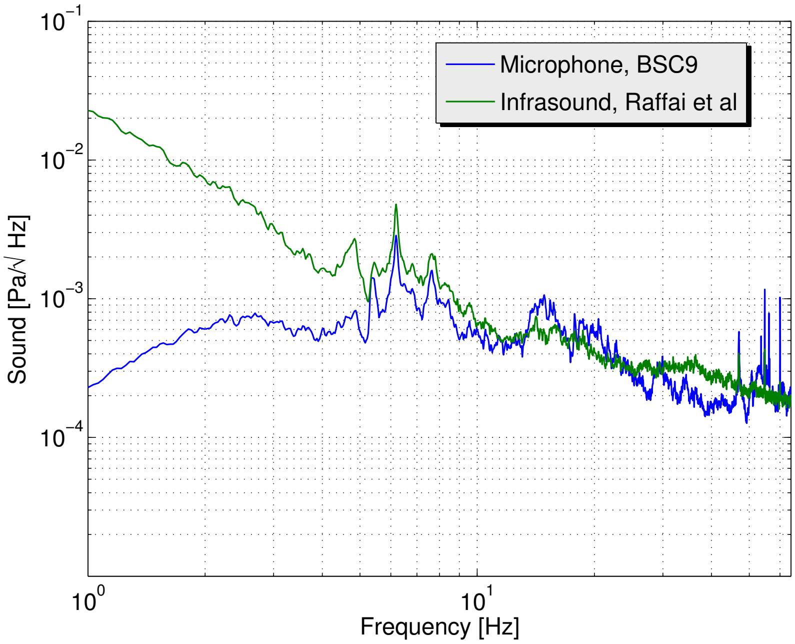

Figure 6.4: Sound spectra taken at the X end station at Hanford.

### 6.1.4 Air handler fans

The air handler fans are large fans used in the air conditioning of the experimental halls at the sites. The fans are very large, and have strong vibrations, and could be a potential Newtonian noise source. The main problem with the noise calculation for the fans is that in principle we need to consider macroscopic motion of fan parts instead of microscopic vibrations. The proper way to calculate this is to expand the Newtonian noise into contributions from the different mass moments of the fan that oscillate at multiples of the rotation frequency $\Omega$. The lowest oscillating orders are the dipole and quadrupole moments. The dipole moment gives rise to Newtonian noise that decreases with distance $r$ between test mass and source as $1/r^3$. One can show that the noise from the quadrupole moment falls as $1/r^5$, and so on for the higher moments. Now, the simple model that we will use in the following is that ideally, all low order mass moments of the fan should vanish by design (assuming that there is a high symmetry of the rotating parts). Then the dominant contribution to Newtonian noise would come from the residual oscillating mass dipole moment, which certainly exists since the fan is vibrating. In fact, we will assume that the vibrations measured at the fan are due to the residual mass dipole moment alone, which is a reasonable assumption since a changing dipole moment requires that a second body attached to the fan (the ground) compensates the associated oscillating momentum. So we will use the dipole formula (Equation 6.3) to calculate Newtonian noise from fans using the vibration spectra. The link between residual dipole moment and vibrations certainly needs to be investigated in more detail, but it shall serve as starting point for a simplified model.

To calculate the fans' Newtonian noise contribution we use

$$h_{NN} = G\,(\text{Mass of Fan})\frac{\xi}{D^3}\frac{2}{L}\frac{1}{(2\pi f)^2}. \tag{6.13}$$

The distance $D$ is 12 m between the fans at the end stations and the end test masses. The distance between the corner station fans and the input test masses is about 27 m. We assume that the vibrating mass is 1000 kg.

We show in Figure 6.5 spectra of an accelerometer epoxied near the base of one of the fans at the Livingston site. The fan's vibration is so strong that we were not able to directly measure it on the fan. So if our noise model is correct, then the Newtonian noise from fans would be much stronger than indicated by the gold trace in Figure 6.2. Further measurements need to be taken on these fans. Note that in the vibration spectral histogram there is a 'quiet time' and a 'loud time'. Periodically, the Livingston fans increase their vibration significantly. The Hanford site does not have this bimodal vibration, and more closely resembles the 'quiet' times.

### 6.1.5 Seismic noise

We currently anticipate that surface seismic noise will be the dominant Newtonian noise source in Advanced LIGO. We measured seismic noise both inside and outside of buildings at both sites using Streckeisen STS-2 [150] and Güralp CMG-40T [119] seismometers.

To estimate the strain noise due to the seismic motion, we use the simple model of evanescent waves propagating on the surface of a homogeneous medium. The Newtonian noise acceleration of the test mass in horizontal direction can be written as

$$a_{\mathrm{NN}}(f) = (\text{Numerical Factor})2\pi \mathrm{i} G\rho_0\xi 2\pi f e^{-hk}, \tag{6.14}$$

where $\xi$ is the vertical surface displacement at the test mass, $k$ denotes the (horizontal) wavenumber of the field at frequency $f$, $G$ is Newton's constant, $\rho_0$ is the average or unperturbed density of the

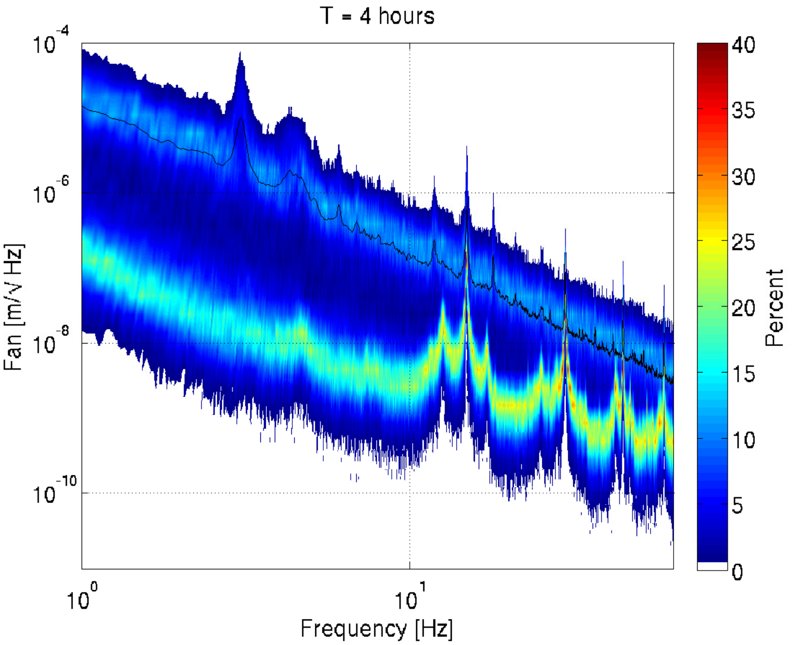

Figure 6.5: Spectral histogram showing displacement due to fan motion, measured near the base of the fan. Measuring the fan directly was not possible, as the accelerometers were completely saturated.

ground, and $h$ is the height of the test mass above ground. The value of the numerical factor depends on the mode content of the seismic field. Density changes due to vertical displacement of the surface produced by Rayleigh waves is partially cancelled by density changes inside the ground due to compressional components of the field. In this case, the numerical factor is about 0.8 depending on the Poisson ratio of the ground.

To convert to strain noise, we divide by $(2\pi f)^2$ and multiply by $2/L$ where $L = 4\,\mathrm{km}$ is the length of the arms. The latter factor is based on the assumption that all four test masses feel incoherent Newtonian noise with the same spectral density. The numerical estimate of the seismic Newtonian noise uses $\rho_0 = 2500\,\mathrm{kg/m^3}$ and the 90th percentile of the two histograms shown in Figure 6.6. This data, combined with Equation 6.14, results in the red and green traces in Figure 6.2.

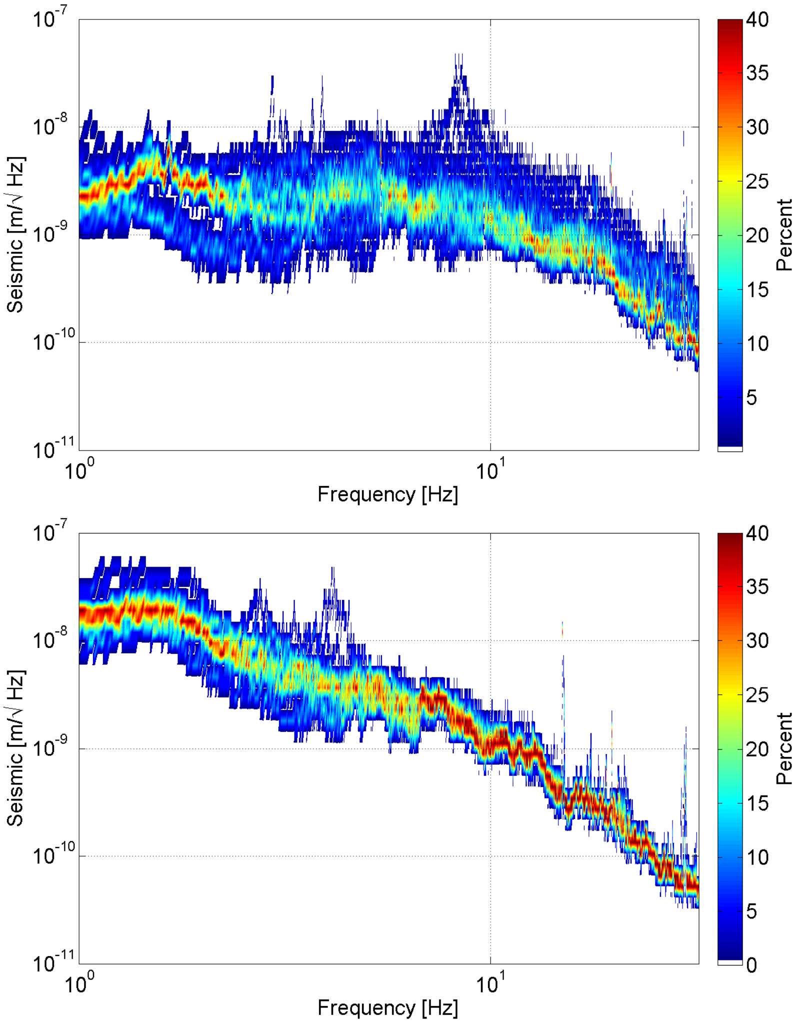

Figure 6.6: Histogram of spectra of vertical ground displacement measured inside the LVEAs at Hanford (top plot) and Livingston (bottom plot).

## 6.2 Optimal seismic arrays

### 6.2.1 Simulation of seismic Newtonian noise

Since Newtonian noise cannot, at this time, be directly measured, we must base our estimates of subtraction capabilities on simulated noise. We attempt to obtain a sufficiently accurate estimate of the Newtonian noise based on information about its source, which in this case is the seismic field. Other contributions to Newtonian noise such as building vibrations or air pressure fluctuations are not considered here because, as discussed in Section 6.1, they are not expected to limit GW interferometer sensitivities in the next ~10 years.

In this Section, we give a description of the time-series generator for the seismic fields, the associated Newtonian noise, and other instrumental noise of the interferometer and seismometers.

We also discuss the suitability of our simulation as an estimate of Newtonian noise at the LIGO sites. The problem is set up as a full time-domain simulation of seismic fields and instrumental noise. Instrumental noise such as seismometer noise or test-mass displacement not generated by Newtonian noise is treated as stationary. In contrast, we do not assume stationarity of the seismic field. Here the attempt is to simulate a seismic field that is comparable to previously measured seismic noise, and to make the content of the field as complex as possible in order to test Newtonian noise subtraction schemes on challenging scenarios. Still, due to computational limitations, simplifications are necessary. In the general case, if ground displacement $\vec{\xi}(\vec{r},t)$ is weak, then the test-mass acceleration due to Newtonian noise can be estimated by Equation 6.2. This equation is valid for arbitrary seismic fields and represents the noise imprinted on the test mass due to Newtonian noise. In our simulation, we only consider seismic fields composed of surface waves (Love or Rayleigh waves, described in Figure 6.21). This simplification is enforced by computational limitations since generating Newtonian noise time series from simulated 3D seismic fields would require months-long simulation runs. We expect this assumption to be reasonable, since surface waves are expected to have much larger amplitudes than body waves [141], and so surface waves give the dominant contribution to Newtonian noise at the surface; however, seismic array measurements currently in progress at the LIGO sites will confirm this and several other of our assumptions for our particular sites.

Freely propagating surface waves like Rayleigh waves and their overtones produce Newtonian noise in such a way that there is always an effective 2D representation of the problem (which is not generally true for all supported wave fields, such as scattered waves). This implies that the numerical simulation can be set up as a 2D simulation over the surface, which is why Equation 6.15 and Equation 6.16 only describe vertical displacement. This approach was chosen to reduce computational costs, and does not over simplify the subtraction problem as long as scattering of seismic waves is weak.

The simulated seismic field is composed of two wave types: wavelets and symmetric surface waves. Wavelets represent seismic waves from far-field sources, while symmetric surface waves represent disturbances due to local sources. The vertical displacement due to a wavelet is described by

$$\xi(\vec{r},t) = \xi_0 \exp(-\tau^2/(2\Delta T)^2)\cos(2\pi f\tau + \phi_0) \tag{6.15}$$

with $\tau = t - \vec{k}\cdot(\vec{r}-\vec{r}_0)/(2\pi f)$. Twenty wavelets are injected for each second of time series randomly distributed over the entire simulation time so that wavelet numbers can be different each simulated second. Frequencies $f$ are drawn from a uniform logarithmic distribution between 8 Hz and 30 Hz, which includes the full range of frequencies for which Newtonian noise is expected to be dominant. Wavelet durations $\Delta T$ are uniformly distributed between $10/f$ to $20/f$, to represent that wavelets can vary in duration depending on the type and source of the disturbance. The distribution of wave vectors $\vec{k}$ is isotropic, to represent far-field sources from all directions. The average speed of sound for seismic waves in the ground is approximately 200 m/s [151] so we allow seismic speeds in our simulation to vary uniformly from 195 m/s to 205 m/s. This variance in speed is a brute-force method to simulate wave scattering, but it is very likely an overestimation of the effect. The initial location $\vec{r}_0$ of the wave maximum lies in the direction of the back-azimuth of the incident wave such that the wave is guaranteed to reach the location of the test mass within the simulated duration of the time series. The initial phase $\phi_0$ of any single seismic wave is not a critical parameter. rWhat is important is that not all seismic fields in our simulation have the same phase, so $\phi_0$ is drawn from a uniform distribution between 0 and $2\pi$.

The second type of wave in the simulation is the symmetric surface wave, described by Bessel functions, with vertical displacement

$$\xi(\vec{r}, t) = \xi_0 J_0(k_0 R) \cos(2\pi f t + \phi_0) \tag{6.16}$$

with $R = |\vec{r} - \vec{r}_0|$. Equation 6.16 represents fields from sources located at $\vec{r}_0$ with distance $r_0$ drawn uniformly between 10 m and 20 m. Sources more distant than 20 m appear at the test mass as distant sources, represented by wavelets as in Equation 6.15. All other parameters are obtained in the same way as for the wavelet, where as before the variation in seismic speed leads to a corresponding variation of the wave number $k_0$. We assume that local sources do not vary strongly over the relevant time scales (defined by the subtraction procedure; see following Sections), so that the local sources are considered stationary. A fixed number of 10 waves from local sources is used. Below 1 Hz, there are typically no more than 2 waves present at a time [152], so we expect that, while the number of waves present can increase with the frequency of the seismic waves, twenty distant and ten local sources is a conservative overestimate of the complexity of seismic fields we will see at the LIGO sites.

The full simulation covers a surface area of 100 m × 100 m with the test mass at its center, which is larger than the area from which interesting Newtonian noise contributions are expected [147]. The number of grid points along each dimension is $N = 201$ so that the grid-point spacing is 0.5 m. We choose a $201 \times 201$ point grid as a compromise between overall grid area, grid spacing, and computational time. The test-mass is suspended 1.5 m above ground, which is approximately the height of the LIGO test masses. As the effective 2D representation is based on the surface term of the gravity perturbations and not the full dipole form [147], we convert the integral in Equation 6.2 into a discrete sum over grid nodes. Using only the surface contribution to the integral, the test mass acceleration along the direction of the interferometer arm is

$$a_{\mathrm{NN}}^{\mathrm{arm}}(t) = G\rho_0 \mathrm{dS} \sum_{\mathrm{l}=1}^{\mathrm{N}\times\mathrm{N}} \frac{\xi_{\mathrm{l}}(\mathrm{t})}{\mathrm{r}_{\mathrm{l}}^2} \cos(\phi_{\mathrm{l}}), \tag{6.17}$$

where dS is the area of the square enclosed by four neighboring grid points, $\xi_l(t)$ is the vertical displacement of grid point $l$ at time $t$, $r_l$ is the distance between the grid point to the test mass, and $\phi_l$ is the angle between the vector pointing from the test mass to the grid point and the direction of the interferometer arm. The sum over grid points in Equation 6.17 is used to determine the time series of the Newtonian noise at the test mass. Time series for each seismometer in Section 6.2.3 and Section 6.2.4 are calculated separately using Equation 6.15 and Equation 6.16, so seismometer locations are not restricted to coinciding with grid points.

We utilize models of the instrumental noise of seismometers and the strain noise of an interferometer to more accurately determine the Newtonian noise subtraction efficacy, as described in Section 6.2.3 and Section 6.2.4. The instrumental noise of all seismometers is simulated with spectral densities that are white (frequency independent) in units of velocity and have a value of $10^{-10}\,\mathrm{m}/\sqrt{Hz}$ at 10 Hz. This is a conservative estimate for commercial geophones.

The seismic spectrum itself plays a minor role for the purpose of this paper, but nevertheless we defined distributions for $\xi_0$ in Equation 6.15 and Equation 6.16 in such a way that the spectral density approximates the median spectrum measured at the LIGO sites. The plot in Figure 6.7 shows the histogram of unaveraged 128 s spectra measured at the LIGO Livingston site over a time of one year during the last science run, and the black curve represents the average spectrum of the simulated seismic field. The average spectral density derived from the histogram is about a factor of two to three larger at frequencies between 10 Hz and 30 Hz than the model used in [141] with correspondingly larger Newtonian noise spectrum.

The strain noise model (excluding Newtonian noise) that we use to simulate interferometer noise for a 3rd generation detector is the same as Equation 6.1. This is a representative noise model for proposed upgrades to Advanced LIGO and Advanced VIRGO, which we refer to as 3rd generation ground-based detectors [144]. Future detectors built at new sites, such as the proposed Einstein Telescope, we call 4th generation detectors [145].

The sampling frequency for all time series is $f_{\mathrm{s}} = 100\,\mathrm{Hz}$ and the observation time is $T = 100\,\mathrm{s}$. We plan to test our subtraction techniques on longer duration simulated data in the future; however, computational time restraints have kept us to this moderate duration for the time being. All time series are high-passed with corner frequency 5 Hz directly after being generated to avoid numerical problems. As can be seen in Figure 6.8, interferometer noise dominates Newtonian noise below 8 Hz and above ~20 Hz, so that we can safely ignore frequencies outside this range when testing subtraction methods.

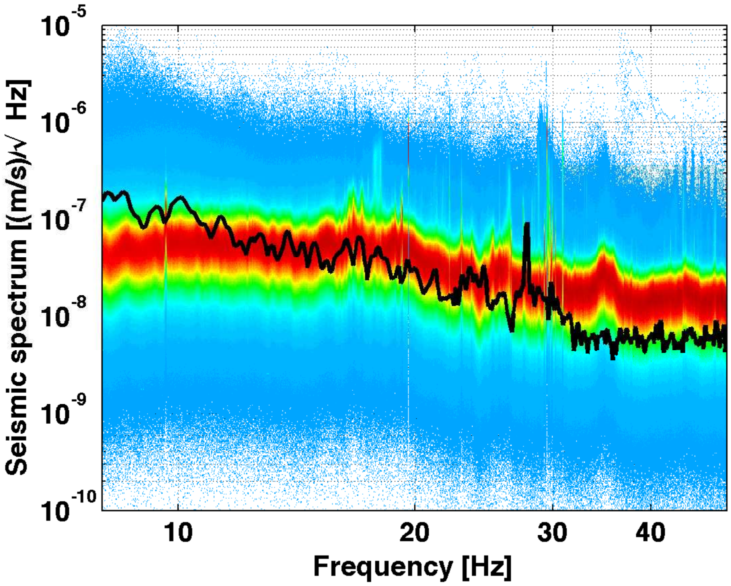

Figure 6.7: Histogram of one year of unaveraged 128 s seismic spectra measured during the 6th LIGO science run inside the corner station of the Livingston detector. The black curve is the spectral density of the simulated seismic field. The spectral histogram of the Hanford site is very similar for the frequencies plotted here.

### 6.2.2 Sensor array optimization

Optimization of seismic arrays with respect to Newtonian noise subtraction was discussed in [153]. The authors calculated subtraction residuals analytically by evaluating explicitly the correlation between seismometers and the test-mass as a function of seismometer locations. The average subtraction residual can be written as

$$R = 1 - \frac{\vec{C}_{\rm SN}^{\rm T} \cdot (C_{\rm SS})^{-1} \cdot \vec{C}_{\rm SN}}{C_{\rm NN}}. \tag{6.18}$$

Here, $\vec{C}_{\rm SN}$ is the cross-correlation vector between seismometers and the Newtonian noise acceleration of the test-mass, $C_{\rm SS}$ is the cross-correlation matrix between seismometers, and $C_{\rm NN}$ is the Newtonian noise variance. These quantities can also be interpreted as (cross-)correlation spectral densities. Given a fixed number of seismometers, the optimal array is found by changing seismometer locations and minimizing $\sqrt{R}$. The equation is idealized, as it does not depend on any details about the way subtraction is implemented, i.e., whether a finite impulse response (FIR) filter is used, or some non-causal post-subtraction filter (see following two Sections for details). For this reason it describes the performance of all subtraction methods that are based on linear filtering, and the optimal array found by minimizing $\sqrt{R}$ is universal for all linear noise filters. Since it is very likely that different noise cancellation techniques will be combined in practice, it seems that

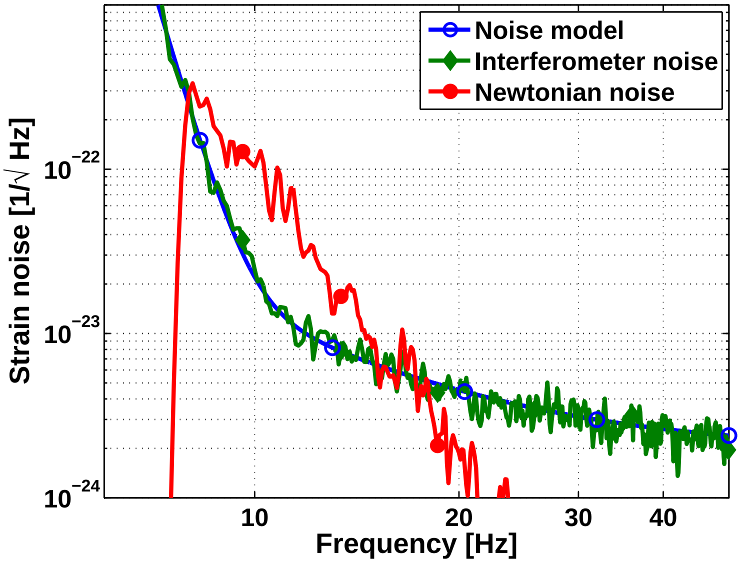

Figure 6.8: Spectrum of simulated Newtonian noise (red line), reference third generation sensitivity curve described in Equation 6.1 (blue line), and simulated interferometer noise based on this noise model (green line). Since other noise sources such as mirror suspension thermal noise and direct seismic vibrations will be the limiting noise sources for second and third generation detectors below ~ 8 Hz, we do not need to consider Newtonian noise at these low frequencies. We therefore do not include low frequency information in our Newtonian noise estimate, which creates a sharp cutoff when the simulated data is viewed in the frequency domain

optimization based on Equation 6.18 is the best one can do.

Correlation patterns of surface waves observed in nature are often well approximated by Bessel functions that characterize isotropic plane-wave surface fields [154, 155]. Adopting a more convenient normalization the corresponding seismic correlation $C_{\rm SS}$ between two points $\vec{r}_i$, $\vec{r}_j$ on the surface is given by

$$C_{\rm SS}(\vec{r}_i,\vec{r}_j) = J_0(2\pi|\vec{r}_i - \vec{r}_j|/\lambda) + \frac{1}{\rm SNR^2}\,\delta_{ij}, \tag{6.19}$$

where $\lambda$ is the length of the seismic wave, and SNR is the signal-to-noise ratio of the seismometers. To find out how well this theoretical model approximates the seismic correlation in the simulation, we calculated $C_{\rm SS}$ between seismometers of increasing distance using our simulated seismic fields. The result is shown in Figure 6.9, where we show that the correlation vs. distance of our simulated seismic fields match the theoretical correlation of seismic fields fairly well, albeit not precisely.

Other terms in Equation 6.18, using the same normalization as for Equation 6.19, are the Newtonian noise variance

$$C_{\rm NN} = 0.5 \tag{6.20}$$

and the correlation between seismic displacement and Newtonian noise acceleration of the test

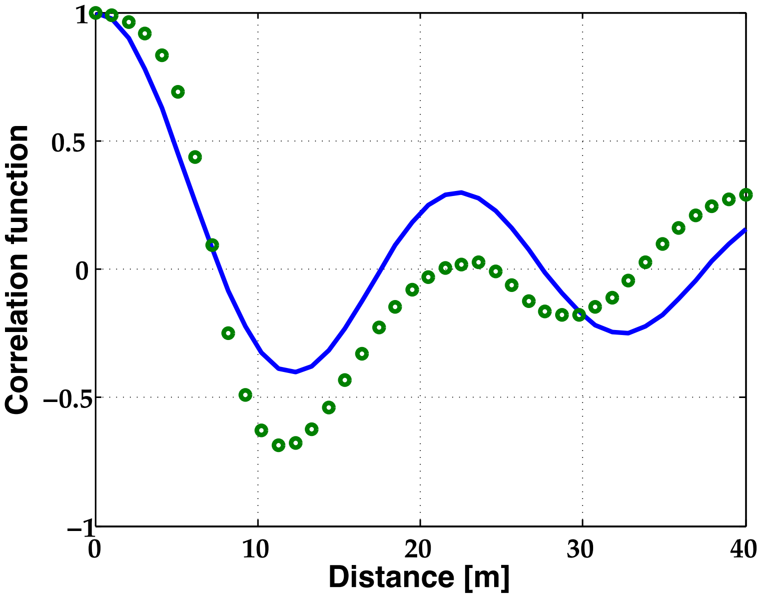

Figure 6.9: Comparison between the theoretical model for seismic correlation of isotropic plane-wave surface fields as described by Equation 6.19 (solid line), and the correlation calculated by the simulation of a field composed of wavelets and locally generated waves (dotted line). Contributions from local sources cause variations of the correlation curve at larger distances between simulation runs (see Section 6.2.1 for details). Therefore, theoretical and simulated correlations match at close distances and deviate strongly at larger distances. Since seismic fields in the context of Newtonian noise subtraction only matter very near the test mass, the match between simulated and theoretical correlations at small distances means that the optimal array determined analytically by minimizing $\sqrt{R}$ in Equation 6.18 should also perform well in simulation, and more importantly that the simulation should be representative of our real subtraction ability.

mass located at the origin

$$C_{\rm SN}(\vec{r}_i) = J_1(2\pi r_i/\lambda)\frac{x_i}{r_i}, \tag{6.21}$$

where $r_i = |\vec{r}_i|$, and $x_i$ is the projection of $\vec{r}_i$ onto the direction of the interferometer arm. Since Equation 6.18 is independent of seismic or Newtonian noise amplitudes, we can use any suitable normalization of the seismic field or the Newtonian noise.

Finding the optimal array is not a trivial task. The result of a stepwise optimization by placing one seismometer after another leads to array configurations very different from the optimum. For the model described by Equation 6.19 to Equation 6.21, the step-wise optimization yields a straight line of seismometers along the direction of the arm, approximately symmetric about the test mass, independent of the number of seismometers. Therefore, configurations close to the optimum can only be found by optimizing all seismometer locations simultaneously. A systematic numerical search for the optimum for more than a few seismometers is prohibitively computationally expensive, and approximate numerical optimization methods need to be applied. The array configuration that we call optimal in the following sections is shown in Figure 6.10. It

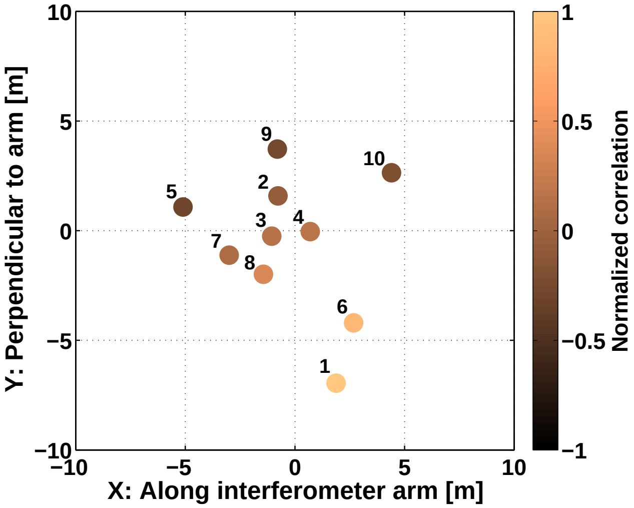

Figure 6.10: Locations of 10 sensors resulting from numerically minimizing the subtraction residual. The optimal array should be symmetric about the test mass located at (0,0), but the optimization routine was stopped once residuals were significantly better than will be required. The colors indicate the normalized seismic correlation between seismometer 1 and all other seismometers.

was found numerically by running a particle-swarm minimization code [156, 157] to optimize the location of 10 noiseless seismometers. It should be clear that the optimal array should have some kind of symmetry, so we know that this configuration is sub-optimal. While this configuration does not represent a global optimum, its subtraction performance should be sufficient for Advanced LIGO and 3$^{\mathrm{rd}}$ generation detectors. As many configurations yield similarly small subtraction residuals, we added further components to the cost function $\sqrt{R}$ to make sure that seismometers are not placed too close to each other. The array shown in Figure 6.10 is the result of minimizing this combined cost function.

As one can see from Equation 6.19 to Equation 6.21, the residual $R$ is a function of seismic wavelength, and therefore frequency, and broadband subtraction performance needs to be investigated. The subtraction residual of the array in Figure 6.10 was minimized at 10 Hz for a seismic wave speed of 200 m/s. In Figure 6.11 we show the subtraction residual as a function of frequency for various array configurations. We compare the array from Figure 6.10 with three different spiral configurations. A spiral seismic array can provide a compromise between high resolution (many sensors densely packed) very close to the test mass, and some sensors far away to capture information about longer wavelength seismic noise. One can see how the number of seismometers and the array size affect subtraction residuals. It is clear that a very small array does not perform well at low frequencies since it provides highly degenerate information at these frequencies whereas larger arrays sample a larger part of the seismic wave. A smaller number of seismometers simply leads to a broadband increase of subtraction residuals except for the smallest frequencies. We want to

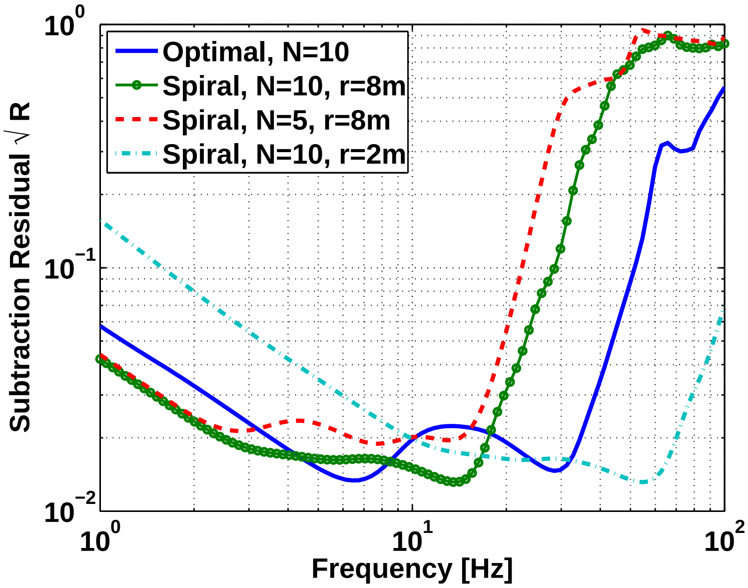

Figure 6.11: Subtraction residual as defined in Equation 6.18 vs. frequency for the array shown in Figure 6.10, and three different spiral configurations.The 'N=10, r=8 m' array is shown in Figure 6.13, and the 'N=10, r=2 m' array is the same, but with all seismometer coordinates scaled down by a factor of four. The 'N=5, r=8 m' array has two sensors at the same positions as numbers 1 and 10 in Figure 6.13, and three other sensors distributed along the two-turn spiral in between these two. It is assumed that the seismometers measure ground motion with SNR = 100 at all frequencies. The Rayleigh-wave speed is 200 m/s.

emphasize that these theoretical predictions only hold approximately for the numerical simulation presented in the following sections, since it does not account for details of the subtraction method as explained before.

Note that all arrays discussed here refer to sensors placed on the floor inside the LIGO buildings. In-chamber vibrations are already suppressed to the level of the noise floor of the best available sensors, so we are not able to measure any motion relevant to Newtonian noise inside the LIGO vacuum envelope. Thus, we use ground-mounted sensors to measure Newtonian noise outside the vacuum envelope.

### 6.2.3 Offline post-subtraction

For the purpose of this chapter, offline post-subtraction denotes the cancellation of noise in recorded data. The noise cancellation filter can therefore be causal or acausal. In this section, we will present a simple acausal implementation of the post-subtraction. The method chosen here is to cancel Newtonian noise on short segments of recorded interferometer time series. The basic idea is to optimally construct a vector of filter coefficients, one coefficient per seismometer axis, and then use these coefficients to form a linear superposition of the seismometer channels as the Newtonian noise estimate. For a general introduction to digital filtering techniques, please see, for example, [158]. In order to determine the effectiveness of this offline subtraction, we look at the residual interferometer sensitivity after removing the Newtonian noise:

$$D(t_m, t_{m+1}) = I(t_m, t_{m+1}) - \frac{\langle I^{\mathrm{c}}, \vec{S}^{\mathrm{c}} \rangle_m}{\langle \vec{S}^{\mathrm{c}}, \vec{S}^{\mathrm{c}} \rangle_m} \cdot \vec{S}^{\mathrm{c}}(t_m, t_{m+1}). \tag{6.22}$$

The residual $D$ in Equation 6.22 corresponds to the interferometer data $I$ minus the Newtonian noise estimate from the seismometer data. The time series used to calculate the Newtonian noise estimate are pre-conditioned with whitening and band-pass filters focusing on the 8 Hz to 30 Hz Newtonian noise band before the correlations are evaluated. We also found it necessary to apply an anti-aliasing window (we used the high-gain Nuttall window) for reasons that will be described below. All quantities subject to the preconditioning are marked with a "$c$". $\langle I^{\mathrm{c}}, \vec{S}^{\mathrm{c}} \rangle_m$ denotes the vector of cross correlations between the interferometer data $I^{\mathrm{c}}$ and all seismometers $\vec{S}^{\mathrm{c}}$ using data of segment $m$ acquired between $t_m$ and $t_{m+1}$. Note that $t_m$ and $t_{m+1}$ are start times of the $m^{\mathrm{th}}$ and $(m+1)^{\mathrm{th}}$ data segments, not time indices. Similarly, $\langle \vec{S}^{\mathrm{c}}, \vec{S}^{\mathrm{c}} \rangle_m$ is the cross-correlation matrix between all seismometers. This means that the filter used here will have one filter coefficient per seismometer for the entire time interval $t_m$ to $t_{m+1}$.

We must determine a reasonable time duration for each segment. Segments are too short if the spectral resolution is too small to disentangle seismic waves at different frequencies. Segments may be too long if the number of seismic waves in that time frame becomes large. A Wiener filter that sees many seismic waves may begin to average over the different waveforms and provide non-optimal noise suppression. Choosing the goldilocks segment duration is somewhat arbitrary; however, it is likely that the appropriate duration depends as much on the nature of the seismic field as on the frequency band targeted by the filter. With our simulation we found the best subtraction performance for 2 s long segments. This is an acausal technique, so testing can be done offline to determine the duration for which we see maximal Newtonian noise suppression on the real data.

Since filter coefficients are re-evaluated for each segment $m$, a simple subtraction of Newtonian noise estimates from consecutive segments can lead to discontinuities in the residual time series. For this reason the Nuttall anti-aliasing window is applied so that noise subtraction is suppressed at the beginning and end of a time segment. Consequently time segments are defined with overlap to provide continuous subtraction of Newtonian noise. Using the Nuttall window, we found excellent subtraction performance with 0.3 fractional segment overlap. Again, some investigation can be done to optimize this number for real data in the future.

Optimal array design has already been discussed in Section 6.2.2. We will compare the subtraction performance of the optimal array presented there with a circular, a spiral, and a linear array. All arrays contain ten seismometers, and are optimized in terms of the extent of the array relative to the location of the test mass. The linear array is simply a line of uniformly spaced seismometers along the direction of the arm extending 8 m away from the test mass in both directions. This linear array is slightly different from the result of the stepwise optimization discussed in Section 6.2.2, but the subtraction residuals are similar. The circular array consists of one seismometer under the test mass and nine seismometers in a circle of radius 5 m around the test mass. The configuration of the spiral array is shown in Figure 6.13 of the following section. The residuals of the noise subtraction (described in Equation 6.22) for each array are shown in Figure 6.12. The noise model represents the sensitivity curve of a potential upgrade of the advanced detectors not including the Newtonian noise, as described by Equation 6.1. Approximately, all arrays perform equally well in post-subtraction. The goal to reduce the Newtonian noise residuals to a level below the noise model is achieved over the entire Newtonian noise band except for the very smallest frequencies. In Section 6.2.4, we will investigate the possibility of combining the post-subtraction with an online feed-forward cancellation.

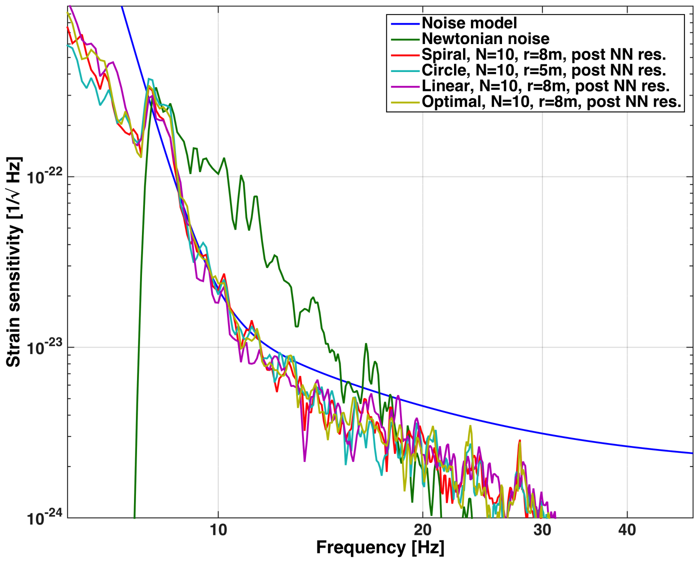

Figure 6.12: Offline Newtonian noise subtraction efficiency for third generation detectors. Spectrum of simulated Newtonian noise (green line), proposed third generation sensitivity curve (blue line), and Newtonian noise residuals of postsubtraction for a spiral array (red line), circular array (cyan line), linear array (magenta), and the optimal array (beige line). Filters derived from all four arrays reduce the simulated Newtonian noise to a level below other sources of interferometer noise as represented by the noise model.

### 6.2.4 Online feedforward subtraction

Online feedforward subtraction can be implemented in two ways. It is possible to continuously cancel Newtonian noise by exerting a cancellation force directly on the test masses. Alternatively, the cancellation can be done on interferometer data. If we had ideal, noise-free actuators, the residuals resulting from applying forces on the test masses and online feedforward cancellation applied to the data would be the same. Applying hardware cancellation forces could also be used to suppress the problem of any non-linear response of the detector to strong Newtonian forces, but is very technically challenging to implement [6]. Since we do not believe that near-future detectors will suffer from significant non-linear upconversion due to Newtonian noise, we only consider feedforward cancellation applied to the data.

The main difference between online feedforward and post-subtraction is that online subtraction can only be done with causal filters. Furthermore, the feedforward filter coefficients can only change slowly in time following slow changes of average correlations between seismometers and the test mass.

The feedforward subtraction scheme that we propose is based on a multi-input, single-output (MISO) finite-impulse response (FIR) filter that is continuously applied to the interferometer output to filter out the Newtonian noise as was already demonstrated successfully for seismic noise cancellation schemes [7]. The inputs consist of the seismometer channels, and the single filter output is the Newtonian noise estimate.

Average correlation between seismometers and interferometer data has a predictable form since average properties of the seismic field depend solely on the wave composition of the seismic field, which is characteristic for each site. This correlation pattern was investigated in Section 6.2.2, where we showed that the simplest theoretical model is a good representation even for the more complex wave composition that is used in our numerical simulation.

As we will show in the following, sufficient feedforward subtraction down to the level of other noise contributions can be achieved with a variety of array configurations including arrays that have seismometers with negligible correlation with the test mass Newtonian noise. The more important design factors are the number of seismometers and the size of the area covered by the array.

The only filter parameter that is predefined is the order of the FIR filter, i.e., the number of filter coefficients. The filter order essentially determines the time span of the filter. Therefore, similar to the post-subtraction scheme, we found that the order can be too high, in which case the seismic array cannot provide sufficient information to disentangle Newtonian noise contributions from individual seismic waves. The filter order is too low when an insufficient amount of data is used to accurately estimate the Newtonian noise from individual, resolved waves. We will later explain why the wave nature of the seismic field still matters in the context of feedforward cancellation. The FIR filter that yielded sufficient subtraction in all simulation runs has order $N = 50$ corresponding to a time span of 0.5 s. The MISO FIR filter coefficients were calculated from the 100 s long seismometer and test-mass time series generated as described in Section 6.2.1. All time series are pre-conditioned with band-pass and whitening filters. An example of a Bode plot of the filter for a spiral array is shown in Figure 6.13. The fact that for example seismometers 3 or 5 have relatively high filter magnitudes at some frequencies is interesting, since their correlation with the Newtonian noise is very small (as calculated by Equation 6.21). This situation can be described as a trade-off between gaining information about how Newtonian noise is generated close to each seismometer (the simple local model), and gaining information about how Newtonian noise integrates over the seismic field based on its wave nature.

The feedforward noise cancellation performance is shown in Figure 6.14. Since the FIR filter coefficients are the same for the entire time series, we included two Newtonian noise residuals, one for the Wiener filter that subtracts on the same time series used to calculate the filter coefficients, and a second one where the same filter is applied to subtract Newtonian noise from another time series. The two time series represent different sets of local sources and wavelets. The subtraction performance is very similar for the two cases, and therefore we can conclude that subtraction performance does not depend as much on the specific wave content of the seismic field as it depends on the average correlations between sensors and the Newtonian noise. While the Wiener filter applied to the data on which it was trained is an acausal use of the filter and could not be applied online, it is useful to see that the subtraction efficacy does not degrade for times that are not the training data for the filter. As with the post-subtraction, feedforward cancellation performed similarly for the circular, linear, and spiral arrays.

Finally, we investigate the possibility of combining the online feedforward cancellation with post-subtraction. Figure 6.15 shows the residual Newtonian noise spectra for the three subtraction methods. Overall, there is no clear advantage or disadvantage to combining the two methods. When both techniques are applied, Newtonian noise residuals are smaller at lower frequencies, but residuals are larger at higher frequencies. In conclusion, it was demonstrated that the standard static MISO FIR Wiener filter provides robust and sufficient subtraction results. Whereas a combination of feedforward and post-subtraction does not give further improvement in simulation, it could prove more effective in scenarios where strong occasional seismic disturbances leave significant residuals after feedforward cancellation.

In our simulation the feedforward filter used was implemented as a static Wiener filter; however, it is possible to let the filter coefficients adapt slowly to changes of the seismic field. This adaptive filter technology has many applications and is well established [133]. Also, once the array design has been chosen based on previous seismic measurements, cross-correlations observed with this array can help to find better array configurations. In other words, it will be possible to adapt to changing properties of the seismic field not only through adaptive filter technologies, but also through changes in the hardware configuration.

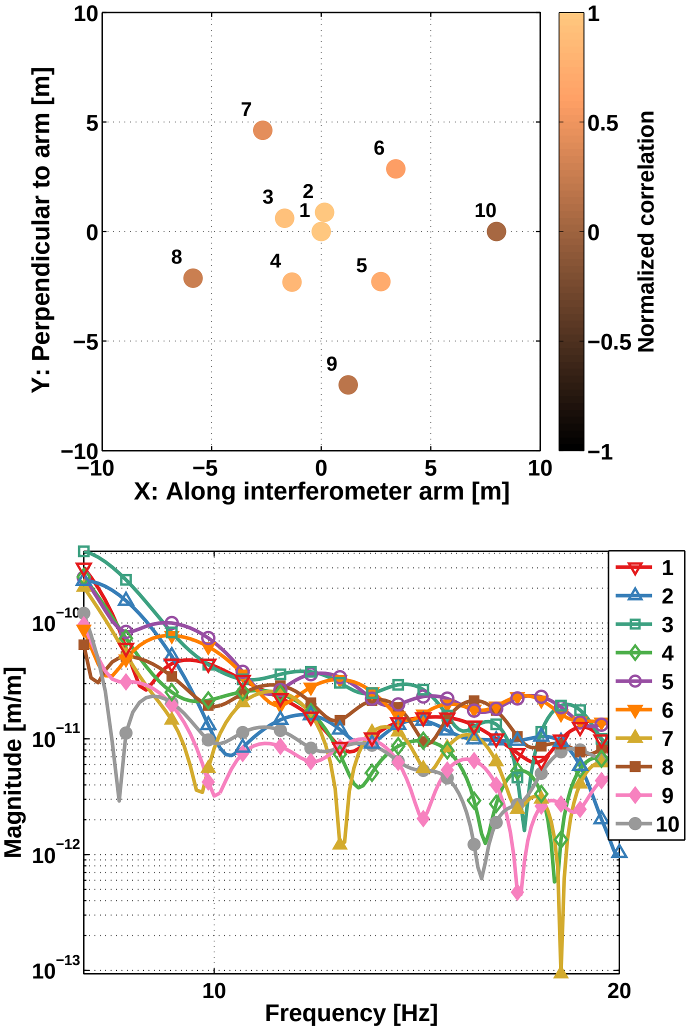

Figure 6.13: The upper plot shows the configuration of the spiral array. The colors correspond to the normalized seismic correlation between all seismometers and seismometer 1. The numbering of seismometers corresponds to the traces in the lower plot, which shows the magnitude of the FIR filter for each sensor in units of test mass Newtonian noise displacement over seismic displacement.

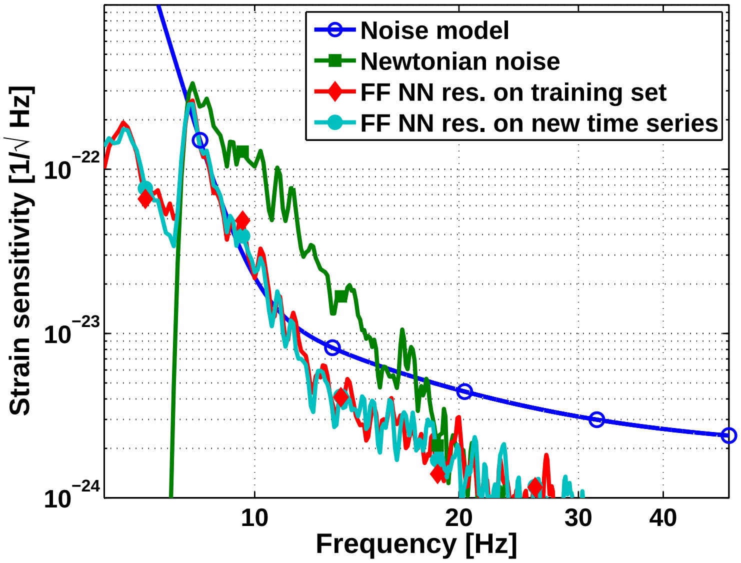

Figure 6.14: Spectrum of simulated Newtonian noise (green), proposed 3rd generation sensitivity curve (blue), and Newtonian noise residuals of feedforward subtraction on the training set (red), and on a second set of time series using the same filter (cyan) using the 10 sensor optimal array.

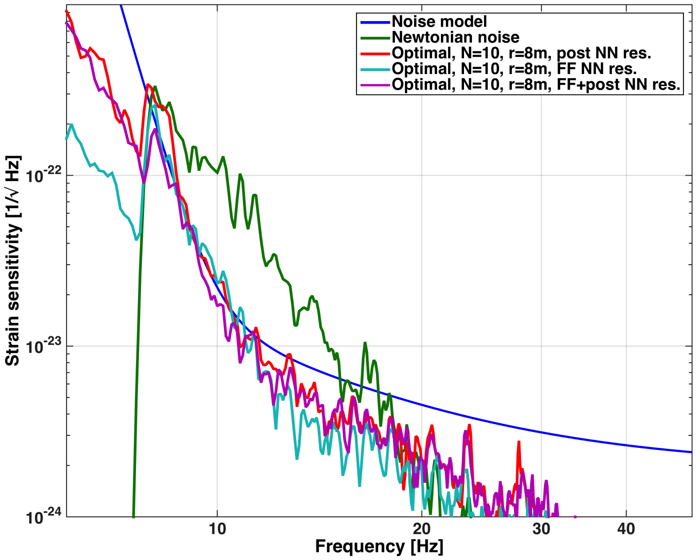

Figure 6.15: Feed forward Newtonian noise subtraction efficacy for 3rd generation LIGO detectors. Simulated Newtonian noise before subtraction (green), expected strain sensitivity (blue), Newtonian noise residuals after subtraction using post-processing (red), online feedforward (cyan), and both methods combined (magenta). Note that the combination of methods is close to the same level as either method individually. This indicates that we can safely apply feed forward subtraction in realtime, and clean up any leftover noise in post-processing if needed.

## 6.3 Comments on the possibility of suppressing the gravitational wave signal through subtraction

A common concern is that true gravitational wave signals could be subtracted out along with Newtonian noise or direct seismic feedforward. The primary reason this is not a concern is that seismic sensors are not directly sensitive to gravitational waves. If they were, sophisticated systems such as LIGO would be unnecessary. Here we briefly discuss several other smaller couplings between Newtonian noise subtraction signals and the gravitational wave channel, and why they are not a concern.

*1) Spurious electromagnetic coupling between test mass actuators and seismic sensors:* The LIGO detectors are controlled such that the test mass mirrors do not move in response to a gravitational wave. Rather, the feedback forces applied to the mirrors in the gravitational wave band contain the gravitational wave information, and combinations of the individual mirror feedback forces comprise the gravitational wave channel used for LIGO analysis. Spurious electromagnetic coupling between the actuators applying these feedback forces and the environmental monitoring sensors used for subtraction can lead to some of the gravitational wave signal being subtracted off unintentionally. This effect is very small, and can be further suppressed by measuring this coupling and correcting for it in the subtraction algorithm if necessary.

*2) Ground recoil due to active seismic isolation, measured by seismic sensors:* Feedback forces actuating on the active seismic isolation structures supporting each test mass contain the gravitational wave signal at a small level. The ground supporting the isolation system will recoil as a result of such a large mass moving. Seismic sensors in the vicinity will measure this recoil as seismic motion, and will attempt to subtract away the Newtonian noise due to this measured seismic motion. This is a $2^{\text{nd}}$ order coupling effect, and the correction to the gravitational wave signal is very small, so further calculations of this effect are outside the scope of this paper. As with the electromagnetic coupling, if necessary, this coupling can be measured and corrected for in the subtraction algorithm.

*3) Earth as a gravitational wave detector:* The Earth responds to gravitational waves with displacement amplitudes $h\frac{\beta}{\omega}$, where $h$ is the gravitational wave strain, $\beta$ is the speed of a shear wave in the ground, and $\omega$ is the frequency [159]. At all frequencies, this displacement is much smaller than is measurable in a local area by the best seismic sensors available, so this coupling is negligible. In order to extract gravitational wave information from the Earth, an array of seismometers covering the entire globe is required [160, 161].

*4) Very short training times for feed-forward filters allow random correlations between transient seismic and gravitational wave events:* A feedforward filter which has been trained on a very short data set can potentially remove signals from the original data stream. This is because such a filter is created using correlation information between seismic sensors and the gravitational wave channel.

During a short data set, there may be random correlations between transient seismic events and true transient gravitational wave events, which would create a filter capable of subtracting away the gravitational wave event. However, we only allow filters to be trained on data sets which are much longer than any burst or compact binary coalescence gravitational wave event that we expect, thus averaging over any seismic transients that could cause a problem for transient gravitational wave events in our data stream.

For all of these effects, it is possible that without correction, a very small amount of gravitational wave signal could be subtracted away from the gravitational wave channel. Once Advanced LIGO is constructed, tests can be done, such as injecting artificial signals into the detector, and measuring the amount by which they are suppressed by Newtonian noise subtraction.

We have shown that a relatively small number of medium sensitivity geophones or accelerometers can be used to estimate the Newtonian gravitational fluctuations with a reasonably high accuracy. Under our simplifying assumptions for the seismic fields and the structure of the ground, this allows us to use seismic data to subtract the gravitational noise due to seismic motion from the interferometer data stream well enough that Advanced LIGO, as well as the 3rd generation detectors, should not be limited by this terrestrial noise source.

We found that the array configuration has a minor impact on the subtraction residuals. The more important design parameters are the number of seismometers, the area covered by the seismic array, and proper preconditioning of the time series that are used for the Newtonian noise estimate.

Our numerical simulation needs to be developed further to test subtraction of other possible contributions to the seismic field that have mostly been considered insignificant for the Newtonian noise problem in advanced detectors in the past, as, for example, body waves and scattered waves. Testing cancellation of Newtonian noise by factors of 10 or more requires a more accurate simulation of seismic fields.

The offline, acausal subtraction scheme should naturally outperform the online, adaptive causal feedforward technique, but for the simple implementation of the post-subtraction used in this paper, the subtraction performances were comparable. To get latency for a cleaned-up data stream to be less than ~1 minute, we will do initial subtraction online and then make the final subtraction offline.

These Newtonian noise subtraction techniques will have a modest improvement on second generation detectors (Advanced LIGO, Advanced Virgo, KAGRA), but the true promise will come towards the end of the decade. At that time these techniques will be necessary to achieve the next order of magnitude improvement in astrophysical reach with third generation detectors.

## 6.4 Study of ground vibration content at the LIGO Hanford site

In 2012, from April through November, we deployed an array of 44 Wilcoxon Research 731-207 accelerometers [148]. The goal of this project was to measure the seismic and ground vibration composition *in situ*. Figure 6.16 shows the location of the accelerometers relative to the vacuum system and some potential noise sources, at the LIGO Hanford Y-end station. Figure 6.17 shows photos of some of the accelerometers in the array. Since these were scattered around the floor of a building, protective cones were used to identify their locations. Accelerometers were adhered directly to the concrete using double-sided tape. White cleanroom tape was used to secure the cones and cables to the ground.

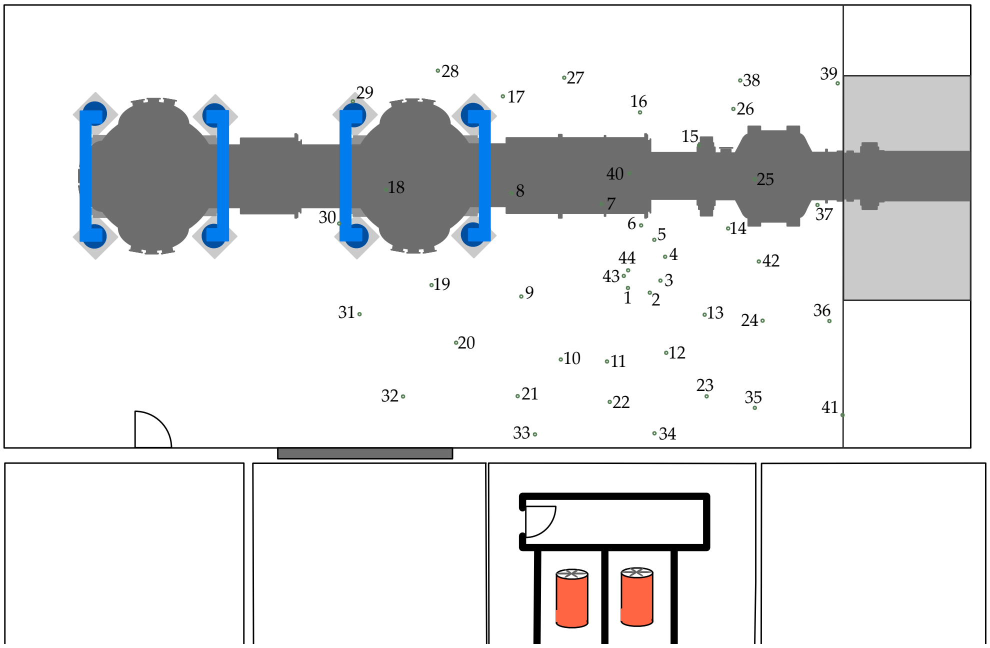

Figure 6.16: To-scale map of accelerometer placement, overlaid with the building walls and vacuum system. Top of diagram is north, and the chamber toward the middle of the room holds the Advanced LIGO ETMY. Numbered green circles are the accelerometer locations, and orange objects are large air handler fans for the HVAC system. Thick lines surrounding fans indicate acoustic isolation walls. The main instrument floor (where accelerometers are located) is isolated from the floor of other sections of the building, indicated by space between the walls.

To analyze the data from this array, we use Capon's method [162] to create "frequency-wavenumber" spectra. This will give us the wavevector of any seismic waves present, from which we can infer the speed of the wave and its direction of travel. This method requires that we analyze a single frequency at a time, so often, as in Figure 6.18, we plot several representative frequencies from the same data set.

We assume that we have $J$ sensors (where $J = 44$ in our case), each of which produces a discrete

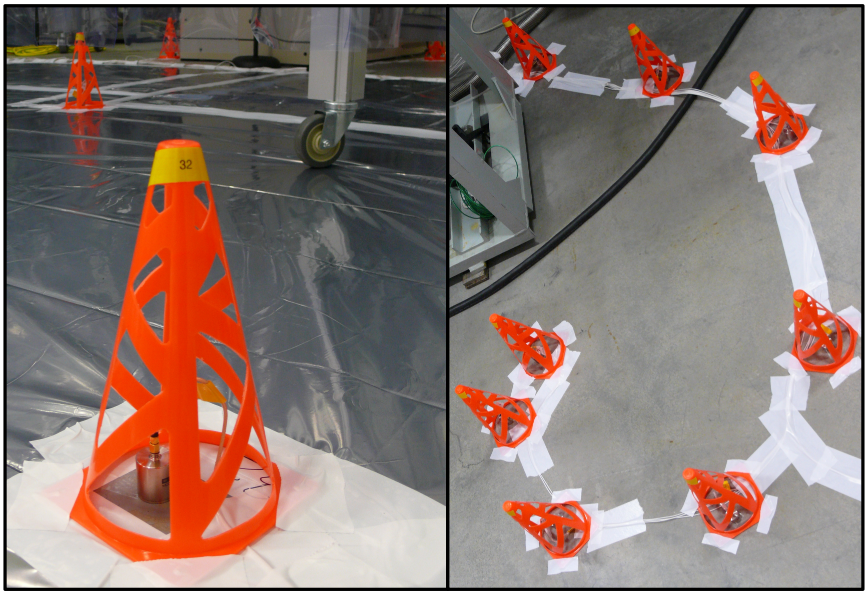

Figure 6.17: (Left) Accelerometer number 32, underneath protective cone, with other accelerometers' cones visible in the background. Small holes were cut in the ground covering such that accelerometers could be in direct contact with the concrete floor. (Right) Central spiral of array: sensors 1-6, 43, 44.

time series with sample rate $f_s$. For each stretch of data (300 seconds long, for our analysis), we apply a window and Fourier transform the data for each sensor. This gives us

$$a_j = \frac{1}{f_s} \sum_{n=1}^{N} d_{n,j}\, w_n\, e^{-i2\pi t_n f}, \tag{6.23}$$

where $n$ is the index of each point in the time series, $d_j$ is the data from the $j^{th}$ sensor, $f_s$ is the data's sample rate, $t$ is the time vector associated with the data, $f$ is our chosen frequency, and $w$ is the Nuttall window.

We create a normalized cross-power spectral density matrix

$$c_{j,l} = \frac{a_j a_l^*}{\sqrt{|a_j|^2 |a_l|^2}} \tag{6.24}$$

between the $j^{th}$ and $l^{th}$ sensors.

We use this matrix, along with the known sensor positions, to create a map, $P$, in wavenumber space. The map will show us, for our chosen frequency, what is the wavenumber of any seismic waves present,

$$P_{a,b} = \left( \sum_j \sum_l \left(c_{j,l}\right)^{-1} e^{-i\left(\vec{x}_l - \vec{x}_j\right)\cdot \vec{k}_{a,b}} \right)^{-1}, \tag{6.25}$$

where $\vec{x}_j$ is the position vector of the $j^{th}$ sensor and $\vec{k}_{a,b}$ is the wavenumber at point $(a, b)$ on the map. We assume that all of our sensors lie in a plane, so both $\vec{x}$ and $\vec{k}$ are two dimensional vectors for our analysis. Typically, the map is normalized so that the largest value is equal to 1.

Figure 6.18 is an example set of maps for 3 different frequencies (10 Hz, 25 Hz and 50 Hz). Note that

$$\vec{k} = \frac{2\pi f}{\vec{v}}, \tag{6.26}$$

where $\vec{v}$ is the apparent surface velocity of the wave. Since the frequency $f$ is fixed for each plot, the axes can also be interpreted as proportional to $1/|v|$. With this in mind, the innermost white circle represents wave speeds of 1,000 m/s. Other circles represent 500 m/s, 250 m/s with the outermost circle 100 m/s. Color scale is the normalized likelihood of the dominant wavevector of the seismic wave present. A vector drawn from the origin of each plot to the maxima of the likelihood is the wavevector of that plane wave. The size of the peak on each map is due to the array's finite resolution.

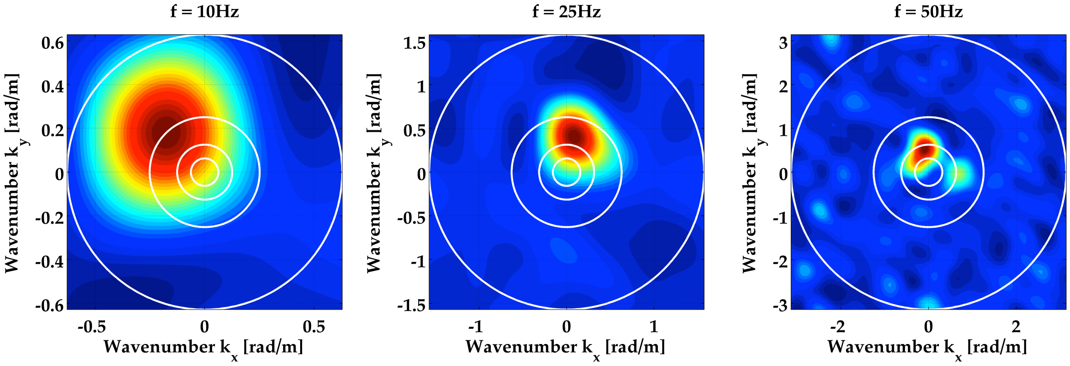

Figure 6.18: Dominant seismic waves at various frequencies (10 Hz, 25 Hz, and 50 Hz), as measured by the array. Innermost white circle represents wave speeds of 1,000 m/s. Other circles represent 500 m/s, 250 m/s, and 100 m/s. Color scale is the normalized likelihood of the dominant wavevector of the seismic wave present. A vector drawn from the origin of each plot to the maxima of the likelihood represents the wavevector of that plane wave.

Vectors pointing toward positive $k_y$ in Figure 6.18 correspond with waves travelling north in Figure 6.16. We see in Figure 6.19 that the air handler fans (shown in orange in Figure 6.16) are the dominant sources of seismic waves, particularly at 10 Hz. In the 50 Hz plot, we can clearly distinguish two waves of nearly equal amplitude present. At times such as these, we'd like to understand the source of each wave individually. Other features in the plots are due to aliasing, and the finite resolution of the analysis method.

While the maps such as that in Figure 6.18 are snapshots in time, we can extract the dominant wavevector from each map, and plot them as a function of time. Figure 6.19 shows the propagation direction of the dominant wave, as a function of time. Figure 6.20 shows the speed of the dominant wave, as a function of time, for each frequency. Our surface array of sensors can only measure the apparent horizontal wave speed of a seismic wave. If an incident wave is in fact a body wave, the horizontal speed would appear to be much larger than the speed of sound in the medium (e.g., concrete). The relatively few times that the apparent wave speed in Figure 6.20 is very high we infer that a body seismic wave was likely present. For most of the time however, the wave speeds indicate true surface waves. From two plots, we can see that most of the time, the source of the seismic waves are consistent over time, although there are regular transient events. Also, it is clear that there more transient events during the daytime (10 am set of plots) versus the night (2 am plots).

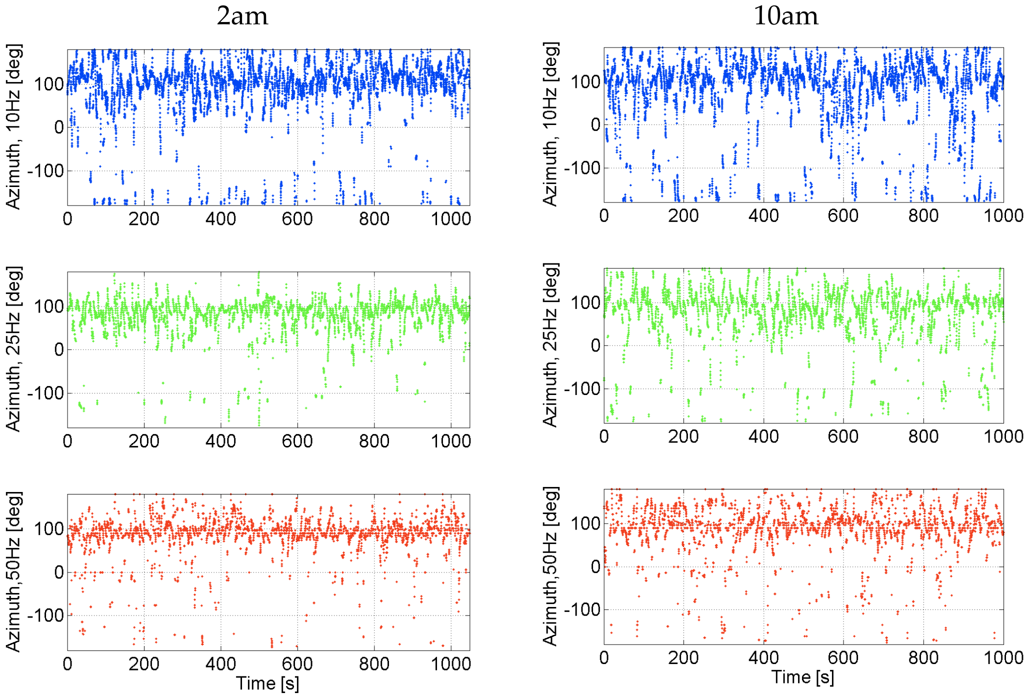

Figure 6.19: Seismic wave propagation direction vs. time.

While there is much more analysis that can be done with this data, and similar array measurements should be completed at each test mass of each site, we can make several important conclusions from the analysis that has been done. It appears that surface seismic waves are dominant, which indicates that we will likely not require sensors buried in the ground around the test masses. Also, scattering of seismic waves seems minimal, which greatly simplifies the analysis and later noise cancellation. Since Newtonian noise will be most significant at frequencies above a few Hz, low noise accelerometers such as the Wilcoxon 731-207s used here will be sufficient for measuring the seismic noise. We also expect to only need approximately 10 sensors per test mass, all within a few tens of meters of the suspensions. Together, this implies that sensor arrays for Newtonian noise cancellation will be much less expensive to implement than was originally feared.

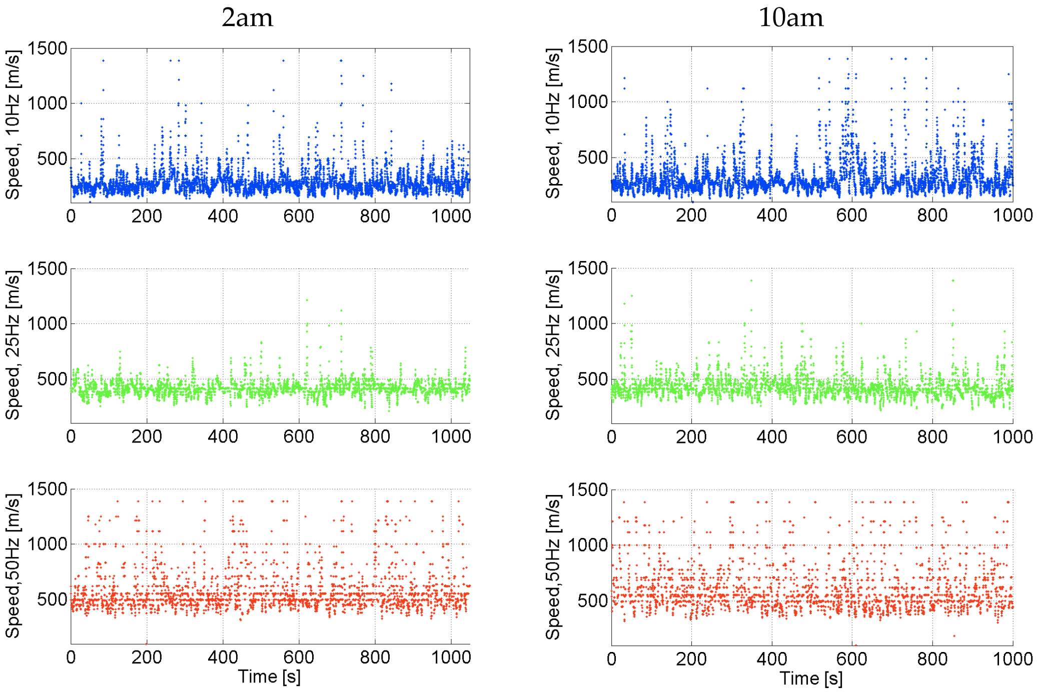

Figure 6.20: Apparent horizontal seismic wave speed vs. time.

## 6.5 Work remaining before Newtonian noise subtraction can be implemented

While much work has been done to ensure that Newtonian noise will not be an immediate limiting noise source for the current generation of gravitational wave detectors, much remains to be done. Perhaps most important will be measurements of the noise generators (seismic in particular) local to each site, so that our estimates can be made more accurate, and we will know at what level we expect Newtonian noise to appear at each mirror at each site. Also important will be time-domain simulations to ensure that we are able to subtract the noise to the level that will be required by future generations of detectors. This section describes some of the most pressing issues, with proposals of work that must be done.

### 6.5.1 Measurement and analysis of seismic field at each test mass

In the past, Newtonian noise estimates have been based on seismic measurements from only a few sensors per detector site, with the assumption that the seismic noise is roughly uniform for the whole site. However, ground-based gravitational wave detectors are very large, and there is no guarantee that the ground motion will be the same at the end stations as at the vertex. Also, we need to know what type of seismic waves are generating the motion at each test mass, which can only be determined by a full array measurement. So far, this has only been done at the Y-end station at the LIGO Hanford site (see Section 6.4 for details).

A single seismometer measurement cannot determine the difference between Rayleigh waves and Love waves; however, a carefully analyzed array measurement can, particularly if that array includes 3-axis sensors. Rayleigh waves directly contribute Newtonian noise while Love waves do not. As seen in Figure 6.21, Rayleigh waves move the surface of the ground up and down. The volume above the ground is primarily air, so displacing air and replacing it with soil or concrete significantly changes the density of that volume. Love waves do not alter the density in the area around the test mass, but rather cause only "traditional" seismic noise. However, since Love waves will couple to many sensors, if we naively use that information to feedforward to the test masses in hopes of canceling Newtonian noise, we will in fact be injecting unnecessary noise. It is inevitable that some Love waves are present in the vicinity of our test masses; however, their relative contribution to the seismic field is important. How much contribution we can tolerate from Love versus Rayleigh surface waves will depend on the level of Newtonian noise subtraction required for each generation of gravitational wave detector. If we only want a factor of a few subtraction, we can allow more contamination of the signals from Love waves. However, if we want a factor of 30 or more (the number often used in upgrades to Advanced LIGO), we may be required to do online separation of wavetypes before feeding the noise cancellation signal forward to the main data channels.

Similarly, if seismic body waves are significant contributors to the ground motion in the vicinity of the test masses, we may be confused between body shear and body pressure waves. Surface waves are expected to have much larger amplitudes than body waves near 10 Hz at the LIGO sites [141]. A borehole sensor buried a few meters underground can tell us if a wave is a body wave, so that we can exclude it (or use it differently). Simulations will need to be done, taking these wavetype confusions into account, to see how the confusion contamination factor affects the noise subtraction limits.

Once we thoroughly understand the seismic fields at each test mass, for each site, we must re-examine our optimal array designs. Do different array shapes offer better noise subtraction for different seismic fields? The optimal arrays in Section 6.2 were calculated using simulated seismic fields, and noise-free sensors. Calculating new optimal arrays can be started with the data from the Hanford 2012 array described in Section 6.4, and then once data has been collected from other locations the array calculations should be relatively easy — one can just plug new data into the established software.

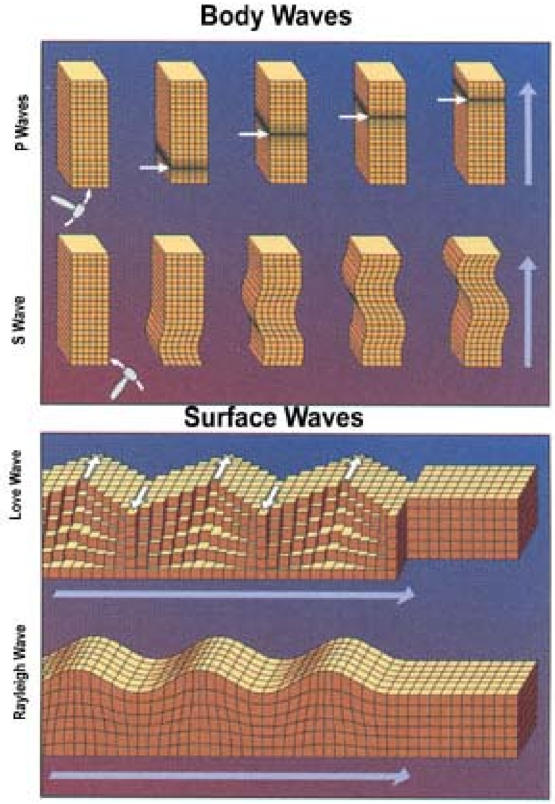

Figure 6.21: Diagram of seismic body and surface waves, from USGS glossary [163]. Wave is travelling in the direction of the grey arrows (upward for body waves, and to the right for the surface waves).

### 6.5.2 Newtonian noise budget improvements

In order to ensure that we fully understand our sources of Newtonian gravitational noise, so that we know what subtraction factor we will be able to achieve, we must improve upon the Newtonian "noise budget". By far the greatest challenge of Newtonian noise subtraction is to make sure that all relevant sources of gravity perturbations are identified, and represented at the appropriate level.

If we compare our seismic spectra from inside the LIGO buildings, to data from the vault seismometers buried in the ground elsewhere on the LIGO site, we have approximately a factor of 10 more seismic noise in the buildings than is typical for the areas surrounding our sites. We need to understand where this excess seismic activity is coming from, such that we can consider ways to prevent it from occurring or from reaching the area near our test masses. If it is coming from the HVAC fans discussed in Section 6.4, is there more that can be done to isolate the fans from the ground? If the ground motion is coming from the building tilts (for example, wind moving the building walls, which are anchored to the ground), is it possible to buttress the walls to stiffen them, or add "walls" of dense netting some distance from the buildings, to prevent the wind hitting the buildings? Understanding the sources of our seismic noise, such that we can break things down in the noise budget, will allow us to make decisions about where it might be worth attempting to shield our detectors from seismically sourced Newtonian noise, versus subtracting it after the fact.

The non-seismic traces in Figure 6.2, described more thoroughly in Section 6.1, may only be accurate to within about a factor of two or three since the noise coupling models are somewhat primitive. It is clear that these noise sources will not be significant problems for Advanced LIGO; however, for future generations of detectors, the gravitational wave strain sensitivity will approach our current estimates of these noises, and we need to know how close they will actually be.

We must build more accurate models of how things like air pressure fluctuations and building tilts will create Newtonian noise at the test masses. In particular, the previously used models assumed that the building walls remained planar, no matter how far they tilted. Understanding how close the walls get to the test masses at varying heights will likely require numerical finite-element modeling of the LIGO buildings.

Overall, a more detailed and accurate noise budget will help us to determine whether or not we are limited by Newtonian gravitational noise in our gravitational wave data. If we can see this comparison live, as we improve the strain sensitivity of our detectors, we can know with some warning when we must actively prepare to begin Newtonian noise cancellation.

Even though we tried to be conservative and to pick models in Section 6.1 such that Newtonian noise is overestimated rather than underestimated, this attempt may well have failed in some cases. The conclusion from these measurements is that we need further simulations and experiments targeting the sources that produce the strongest noise.

Given the results from Section 6.1, we identify the seismic Newtonian noise, the Newtonian noise from the building tilt and the chamber vibrations as main candidates. As previously explained, the Newtonian noise from fans could not be estimated well and may be stronger than shown in our results. Future experiments and simulations should target these four sources to further improve estimates. In order to improve the estimate of chamber Newtonian noise, we need a more accurate simulation that also includes tilts. The building tilt contributions can be improved by using better estimates of the mass of the walls, and by studying in detail how walls move to check if the simple uniform tilt motion is a valid assumption. The following list of tasks summarizes what should be done for the next phase of Newtonian noise investigations:

a) Improve numerical simulation of Newtonian noise from chambers (including near chamber mechanical structures);

b) Calculate/simulate Newtonian noise from the suspension cage near the test masses;

c) Measure building motion at end stations (for this, four or more accelerometers should be attached near roof height at all building walls, monitoring displacement normal to the wall);

d) Measure vibrations on large fans in mechanical rooms with strong motion sensor;

e) Construct theoretical model of Newtonian noise from the fan's rotating parts (in terms of mass multipoles).

If these analyses show that Newtonian noise has been underestimated significantly, then it will be easy to adapt the subtraction scheme for aLIGO to account for Newtonian noise from other sources.

### 6.5.3 Analysis of seismic shielding via excavations

An interesting idea that has come up over the years is whether or not it is possible to excavate a large area around the test mass, so there is (almost) no ground material to be displaced, and so nothing with which to create seismically generated Newtonian noise. Such an excavation might look like Figure 6.22, so that the void is nearly a hemisphere, of order 5 m in radius and 4 m depth [164]. J. Harms has shown that the excavations described in Figure 6.22 will provide approximately a factor of 3 reduction in Newtonian noise.

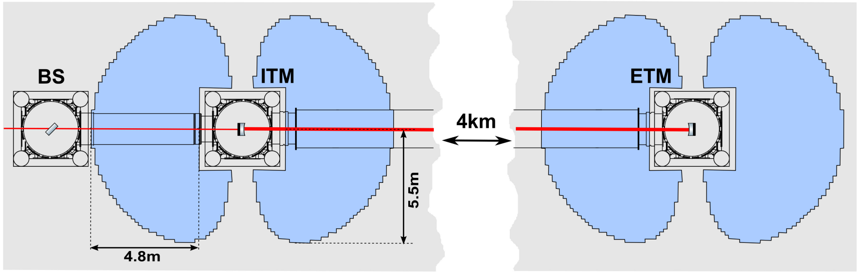

Figure 6.22: Excavation to shield test masses from Newtonian noise. Figure 4 from [164].

We must check though whether we can support the weight of the vacuum chambers and other equipment inside of the excavated area. One option (shown in Figure 6.22) is to leave a central wall dividing the hemisphere. If it is perpendicular to the arm cavity, vertical vibrations that move the surface will not to first order create Newtonian noise that will affect the cavity's length. However, the geometry of the structure may amplify some vibrational modes that cause direct ground motion to be a limiting noise source. If instead the entire hemisphere were excavated, and we create a support structure, we should consider whether it is possible to engineer something that will be sufficiently vibration-free, as well as strong enough to support the weight of the vacuum chamber and everything that goes inside and supports the chambers.

Based on the array data discussed in Section 6.4, we do not believe that scattering of seismic waves is a significant issue, in the current configuration. However, if we create enormous vacancies in the ground, it is possible that we could see coherent scattering. We must determine whether scattering off of the surface of the excavation will change the seismic field content significantly enough to undo any gains we might make. This could indeed pose a severe problem to Newtonian noise subtraction even if scattering is identified and fully characterized. Scattering can in principle make it impossible to estimate Newtonian noise from seismic measurements at the surface since it can lead to a more complex field structure that is not completely characterized by surface displacement. Moreover, it is possible that scattered waves have higher wave numbers compared to the freely propagating surface waves, so that the density of the seismic array would need to be increased to a point where it becomes very challenging to monitor the entire field accurately. Some work has been done to model the effect of seismic scattering at large length scales for choosing future detector sites [165], but in order to determine how significant scattering may or may not be at length scales relevant to existing detectors, we will likely require finite-element modelling to take into account the realities of the situation.

A potential benefit of having created a large excavation could be that we can easily place many sensors out of the plane of the ground surface, without significant extra effort and cost. We should investigate how much extra benefit we would achieve by extending the array. If body waves are shown to be a significant source of seismic noise near any of the test masses, this could be extremely important.

### 6.5.4 Newtonian noise cancellation using mass actuators

Much of the discussion surrounding Newtonian noise focuses on subtraction at the point where it affects detector sensitivity, and occasionally isolation. Rarely do we discuss noise cancellation at the source. However, it may be worthwhile to take this idea seriously, since it may turn out to be easier to implement than, say, the excavation isolation discussed in Section 6.5.3.

For this purpose, we posit that the primary source of Newtonian noise is Rayleigh surface seismic waves. These waves displace air with more-dense concrete (the floor material at the LIGO sites) in the vertical direction. See Figure 6.23 for a visualization. In the left panel, the test mass is shown as a free body in the presence of a perfectly flat floor. The force vector $\vec{F}_{\mathrm{air}}$ is shown for the gravitational attraction between the test mass and a volume of air (purple dashed object). The component of this force that is important for the gravitational wave detector is along the beam axis, shown as $\vec{F}^{\parallel}_{\mathrm{air}}$. Since the purple hemisphere is the only area that will change density, we ignore all other static gravitational forces. The right panel of Figure 6.23 shows a similar situation; however, a portion of the floor has been displaced, into the area of the orange object. New force vectors are shown in orange, with the previous purple vectors for reference. Note that the test mass now has an extra force of magnitude equal to the difference between the horizontal components,

$$\delta\vec{F}^{\parallel} = \vec{F}^{\parallel}_{\mathrm{ground}} - \vec{F}^{\parallel}_{\mathrm{air}}. \tag{6.27}$$

The question to answer for mass actuators is whether it is possible to place a mass on the

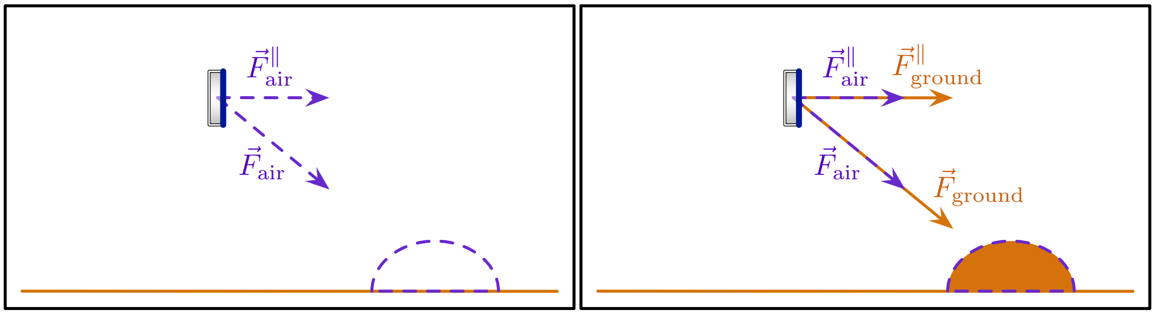

Figure 6.23: Problem setup for mass actuator noise cancellation. Left panel shows gravitational force vector between test mass and volume of air (purple), as well as the horizontal component. Right panel shows the new gravitational force vector if that volume is filled (orange).

opposite side of the test mass, and move it such that any extra horizontal component of the force vector is exactly cancelled by an opposing force vector. Figure 6.24 shows this setup. In the left panel, a block (grey) is located near the floor, and has a gravitational force vector $\vec{F}_{\rm mass1}$, as well as a small horizontal component $\vec{F}^{||}_{\rm mass1}$. In the right panel, when the volume of air is replaced by the floor material, the mass actuator block is moved upwards. The difference between the old and new horizontal components of the force due to the mass ($\vec{F}^{||}_{\rm mass2} - \vec{F}^{||}_{\rm mass1}$) should exactly cancel $\delta\vec{F}^{||}$. We do not care about motions of the test mass in the vertical direction at the level of this Newtonian noise, so the vertical components of the force are not required to cancel.

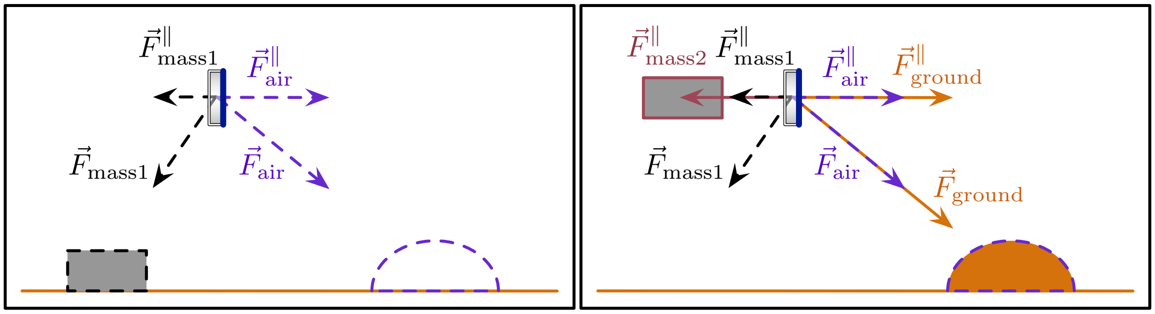

Figure 6.24: Mass actuator noise cancellation setup. Left panel now includes a block (grey) sitting on the floor, and the force vector. Right panel shows the block moved vertically, to counteract the effect of the change in density due to the air being replaced by floor material.

Such a mass actuator system would not be trivial to implement, however. Above, we have assumed that there is only a single area of ground that can move and so we can *a priori* set the location of the block. In reality, we would require a large array of mass actuators to cancel arbitrary ground motion that has a component along the laser beam axis. Simulations will be required to determine what kind of array of actuators would be required — are a few large actuators sufficient, or would we need many, many small actuators?

This simplified picture also does not take into account the fact that the ground will recoil when the block is moved. The block must be moved in such a way that the sum of the block motion plus the recoil exactly counteracts the original $\delta\vec{F}^{\parallel}$. The actuator must be able to move fast enough to counter seismic waves up to a few tens of Hz, but must also be stable enough that it does not cause extra ground vibrations above levels already present. Careful mechanical simulations of any mass actuator system must be done to understand all of these effects.

# Chapter 7

# Conclusion

We have introduced the design choices of LIGO interferometers, particularly in terms of low noise operation and noise couplings. Several new control schemes have been described that help alleviate challenges with commissioning such sophisticated gravitational wave detectors.

Optimal Wiener filters for feedforward noise cancellation have been introduced. Their efficacy for direct seismic vibration isolation has been shown, both for the LIGO sites in Hanford, WA and Livingston, LA, as well as special applications at the 40 m Lab.

Newtonian gravitational noise estimates, and their impact for future ground-based gravitational wave detectors have been updated using *in situ* measurements at the LIGO sites. We show that the subtraction of this noise will be critical for future detectors, but also that it will be simpler (requiring fewer sensors) than previously expected. Importantly, future plans for applying Newtonian noise cancellation are laid out in detail.

These techniques can be applied to Advanced LIGO and all other ground-based gravitational wave detectors to improve their sensitivity, or any other high-precision experiment.

# Appendix A

# Control theory

Broadly defined, controls are used to modify the behavior of a system. For example, an unstable system such as an inverted pendulum can be actuated upon to keep the pendulum upright. Another example is changing the behavior of already stable systems, such as damping a resonant mode in a pendulum system to keep the pendulum bob still.

A basic control system consists of the thing that will be controlled, often referred to as the "plant", a sensor that reports the current behavior of the plant, an actuator that can act on the plant, and a filter that will determine how to use the sensor signal to create an actuator signal. In the case of either type of pendulum, the pendulum itself is the plant, and one can sense the current angle of the pendulum relative to vertical, and actuate by moving the location of the support point in a horizontal direction. The control filters are used to determine how much to move the support point, and in which direction. Here I will briefly describe benefits and drawbacks of different kinds of controls, but for a more thorough introduction see, for example, [166].

Controls are used widely throughout the LIGO detectors. One of the most fundamental uses is to control the length of the laser cavities, to keep them within the linear range of the sensors, so that we can detect gravitational waves. They are also used to control the frequency and intensity of the main pre-stabilized laser, provide seismic isolation, control the angular motion of cavities, and in many other applications.

## A.1 Feedback controls

Feedback is arguably the most common form of control system. A sensor is used to determine the current position (or velocity, or other variable) of the system, and is compared to a desired setpoint. The difference between these is the error signal of the system. A control filter (in this thesis we only consider linear control filters, although in general they can be non-linear) is used to transform the error signal into a control signal. The control signal is fed to some kind of actuator that can affect the system. The system responds to this actuation, and the new position is determined by the sensor.

All of these components together make a “feedback loop”. Where the loop has high gain, the error signal will be forced to zero. Note that this nulling of the error signal implies that the loop will inject any noise from the sensor (e.g. electronics noise) into the loop.

Often, the sensor, actuator and plant are pre-determined and fixed, while the control filter must be chosen by the user. The total dynamics of the system must be “stable”. If the plant, sensor or actuator require a high-order set of equations to describe them, it can be challenging to create a control filter that stabilizes the system. On the other hand, feedback systems do not require exquisite knowledge of the fixed components of the system. A corollary to this is that feedback systems can be designed such that they are robust against small variations in the plant. Obviously this does not remain true if the plant changes drastically (such as a sign flip in the optical plant of the interferometers, as described in Section 4.4).

For many systems, feedback control is appropriate, but since it requires that the disturbance pass through the system before it is sensed and actuated upon, it may not be the best solution for all systems. For example, if noise is non-linearly coupled into the system (like angular fluctuations of a mirror causing scattered light, which will couple back into other areas of the interferometer in potentially unknown ways), we prefer to eliminate the noise before the coupling can take place. In cases such as this, feedforward (as described in Section A.2) may be a better choice.

### A.1.1 Calculating effect of feedback

Figure A.1 shows a very simple feedback system. The plant, or system dynamics (eg. pendulum) is labeled $P$ in the green block. The sensor (eg. measurement of pendulum angle) is labeled $S$ in the blue block. The actuator (eg. to move the pendulum) is labeled $A$ in the red block. The control filter, $G$ is represented by the yellow block, and is what will be used to decide how to create an actuation signal from the sensor.

We will assume that we are only dealing with linear time-invariant systems, and that our analysis will be entirely in the frequency domain, where we can find that the effect of two systems in series is the same as the product of the two frequency responses.

If location $a$ in Figure A.1 represents an input disturbance $\delta x$, location $b$ will represent the residual noise after stabilization $\delta x_s$. Location $b$ is referred to as the error point of the feedback loop, and location $c$ is the control point of the loop.

To determine the effect of the loop, we want to calculate $\frac{\delta x_s}{\delta x}$. The traditional way of calculating this quantity involves writing out the equation for each summing node, and solving. Here,

$$\delta x_s = \delta x + \delta x_s(GAPS). \tag{A.1}$$

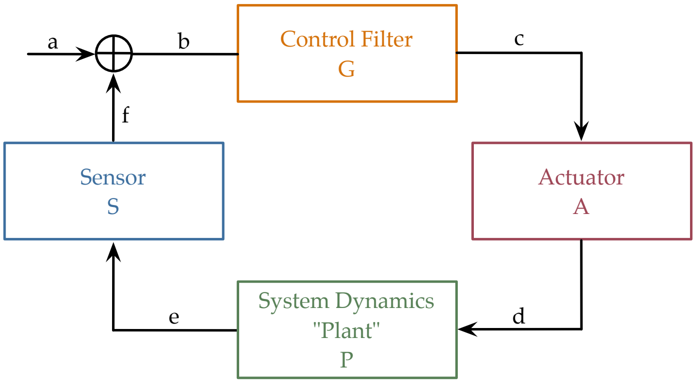

Figure A.1: Cartoon of how feedback control works. Sense disturbance after the plant, use a controller (designed by the user) to decide how to actuate on the plant.

We isolate $\delta x_s$,

$$\delta x_s(1 - GAPS) = \delta x \tag{A.2}$$

and then divide to find

$$\frac{\delta x_s}{\delta x} = \frac{1}{1 - GAPS}, \tag{A.3}$$

which is the canonical form for the closed loop effect of a feedback loop. While it is possible to solve more complicated systems (such as those shown in Figure 4.12 and Figure 4.15) with this technique, it can become very burdensome.

Instead, we will recognize that we can equivalently write out a system of equations, where each equation relates one numbered location in the loop diagram to each adjacent location.

$$\begin{aligned} b &= a + f \\ c &= Gb \\ d &= Ac \\ e &= Pd \\ f &= Se \end{aligned} \tag{A.4}$$

We can add a source term and write this system of equations in matrix form,

$$\vec{v}_{\text{steady state}} = M^T \vec{v}_{\text{steady state}} + \vec{v}_{\text{in}} \tag{A.5}$$

where

$$\vec{v} = \begin{pmatrix} a \\ b \\ c \\ d \\ e \\ f \end{pmatrix} \tag{A.6}$$

and

$$M^T = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & G & 0 & 0 & 0 & 0 \\ 0 & 0 & A & 0 & 0 & 0 \\ 0 & 0 & 0 & P & 0 & 0 \\ 0 & 0 & 0 & 0 & S & 0 \end{pmatrix}. \tag{A.7}$$

Note that other derivations may write $M^T$ as just $M$, but for consistency with Section 3.1.2, we will call it $M^T$ (the transpose of $M$). Solving Equation A.5 for $\frac{\vec{v}_{\text{steady state}}}{\vec{v}_{\text{in}}}$, we find

$$\frac{\vec{v}_{\text{steady state}}}{\vec{v}_{\text{in}}} = \left(\mathbb{1} - M^T\right)^{-1}. \tag{A.8}$$

This "transfer matrix" tells us the transfer function from any point to any other point in the loop (even if we do not have the ability to actually measure that particular transfer function). For example, to compare to Equation A.3, we can extract the $[2, 1]$ element from Equation A.8 which will give us the transfer element from location $a$ to location $b$. This is

$$\frac{\vec{v}_{\text{steady state}}}{\vec{v}_{\text{in}}}[2, 1] = \frac{1}{1 - GAPS}, \tag{A.9}$$

which is the same result as the original method of calculating the closed loop transfer function. While the system shown in Figure A.1 is simple to solve, this nodal matrix technique is easily extended to much more complex systems, such as those shown in Figure 4.12 or Figure 4.15.

## A.2 Feedforward controls

As discussed above, some systems may require a different type of control system to prevent disturbances from propagating through the system. Figure A.2 shows a basic block diagram of a feedforward system that can accomplish this. An external disturbance affect the plant through a coupling transfer function, but if we are able to witness that disturbance independently from the plant, we can construct a filter to send through an actuator and cancel the disturbance's effects before it affects the plant. Note that for some noises (such as quantum shot noise or radiation pressure noise) we do not have the ability to independently measure the noise, and so cannot apply feedforward to suppress those noise contributions.

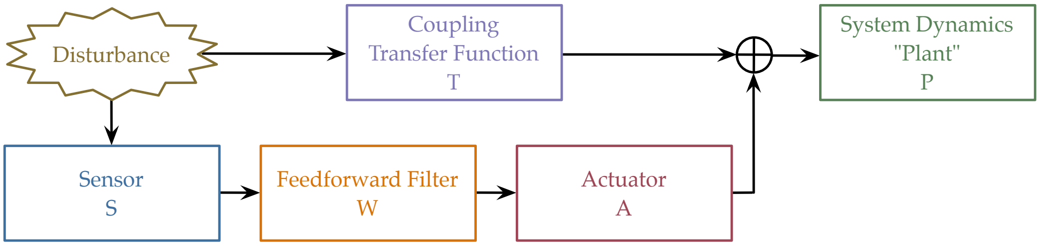

Figure A.2: Cartoon of how feedforward control works. The disturbance directly affects the plant, but it is also sensed by an external "witness" sensor. The output of the sensor is fed through a filter to the actuator, and cancels the effect of the disturbance before it arrives at the plant. Dynamics between the disturbance and the plant are the primary features that the feedforward filter must account for.

A challenge with the feedforward control topology is that it requires very precise knowledge of the system. In particular, one must very carefully measure the transfer function from the input of the actuator to the input of the plant, so that can be removed from the calculation of the feedforward filter. The ability to measure this transfer function can be a limiting factor in the ability of the feedforward system to subtract the noise. The error goes as

$$\sigma = \sqrt{\frac{1-C}{2NC}}, \tag{A.10}$$

where $C$ is the coherence of the measurement, and $N$ is the number of averages in that measurement. So, if we only measure the actuator to 10 %, we will only be able to subtract approximately a factor of 10 of the noise. If we want to subtract a factor of 100, we must measure to better than 1 % (assuming no other limitations, such as the noise of the witness sensor, become limiters).

# Appendix B

# Review of Pound-Drever-Hall locking

This Appendix sketches out the method for deriving the canonical Pound-Drever-Hall ("PDH") error signal for reflection locking a Fabry-Pérot cavity. Here we only derive the error signal assuming a single radio frequency modulation, although in general several may be applied in series. To expand this derivation for a full interferometer rather than a simple Fabry-Pérot cavity, one can follow a derivation such as that described in K. Arai's thesis [31].

We will write the laser output as the electric field,

$$E_{\text{laser}} = E_0 e^{i\omega t} \tag{B.1}$$

where $E_0$ is the amplitude of the electric field and $\omega$ is the frequency of the laser. We phase modulate the light at frequency $\Omega$ by passing through an electro-optic modulator (EOM), which gives us the total electric field incident on the cavity,

$$E_{\text{inc}} = E_0 e^{i\omega t} e^{i\Gamma \cos \Omega t}, \tag{B.2}$$

where $\Gamma$ is called the "modulation depth". We expand this using the Jacobi-Anger expansion to get

$$E_{\text{inc}} = E_0 e^{i\omega t} \sum_{n=-\infty}^{\infty} i^n (-1)^n J_{|n|}(\Gamma) e^{in\Omega t} \tag{B.3}$$

where the functions $J_n$ are Bessel functions of $n^{\text{th}}$ order. We approximate this by only keeping terms up to $|n| = 1$,

$$\begin{aligned} E_{\text{inc}} &\simeq E_0 e^{i\omega t} \sum_{n=-1}^{1} i^n (-1)^n J_{|n|}(\Gamma) e^{in\Omega t} \\ &= E_0 e^{i\omega t} \left( J_0(\Gamma) - iJ_1(\Gamma) e^{i\Omega t} + iJ_1(\Gamma) e^{-i\Omega t} \right) \\ &= E_0 \left( J_0(\Gamma) e^{i\omega t} - iJ_1(\Gamma) e^{i(\omega+\Omega)t} + iJ_1(\Gamma) e^{i(\omega-\Omega)t} \right). \end{aligned} \tag{B.4}$$

The final form of this equation makes it clear that (to first order) we now have carrier light of amplitude $E_0 J_0(\Gamma)$ at the original carrier frequency $\omega$ as well as two sidebands at frequencies $(\omega + \Omega)$ and $(\omega - \Omega)$, each of amplitude $E_0 J_1(\Gamma)$.

The light that we will be present at the reflection port of the cavity will be given by

$$E_{\text{refl}} = E_{\text{inc}} r_{\text{cav}} \tag{B.5}$$

where $r_{\text{cav}}$ is the amplitude reflectivity of the cavity, including both the prompt reflection and leakage of the circulating power in the cavity. $r_{\text{cav}}$ is a function of the frequency $\omega$ of the light, and is given by

$$r_{\text{cav}}(\omega) = -r_i + \frac{t_i^2 r_e e^{-i\frac{2L\omega}{c}}}{1 - r_i r_e e^{-i\frac{2L\omega}{c}}}. \tag{B.6}$$

As defined in Chapter 3, $r_i$ is the amplitude reflectivity of the input mirror of the cavity, $r_e$ is the amplitude reflectivity of the end mirror, and $t_i$ is the amplitude transmission of the input mirror. $L$ is the length of the cavity and $c$ is the speed of light. Equation B.5 expands to

$$E_{\text{refl}} = E_0 \left( J_0(\Gamma) e^{i\omega t} r_{\text{cav}}(\omega) - i J_1(\Gamma) e^{i(\omega+\Omega)t} r_{\text{cav}}(\omega + \Omega) + i J_1(\Gamma) e^{i(\omega-\Omega)t} r_{\text{cav}}(\omega - \Omega) \right). \tag{B.7}$$

Photodiodes are not capable of measuring the electric field directly. Rather, they detect the power of the light. For example, at the reflected port,

$$P_{\text{refl}} = E^*_{\text{refl}} E_{\text{refl}}. \tag{B.8}$$

where $E^*$ is the complex conjugate of $E$. Explicitly, this expands to

$$\begin{aligned} P_{\text{refl}} = E_0^2 & \left( J_0(\Gamma) e^{-i\omega t} r^*_{\text{cav}}(\omega) + i J_1(\Gamma) e^{-i(\omega+\Omega)t} r^*_{\text{cav}}(\omega + \Omega) - i J_1(\Gamma) e^{-i(\omega-\Omega)t} r^*_{\text{cav}}(\omega - \Omega) \right) \\ & \times \left( J_0(\Gamma) e^{i\omega t} r_{\text{cav}}(\omega) - i J_1(\Gamma) e^{i(\omega+\Omega)t} r_{\text{cav}}(\omega + \Omega) + i J_1(\Gamma) e^{i(\omega-\Omega)t} r_{\text{cav}}(\omega - \Omega) \right). \end{aligned} \tag{B.9}$$

Multiplying this out, and dropping the $2\Omega$ terms (because, as described below, we will demodulate at $\Omega$ and then low-pass the resulting signal) will leave us with

$$\begin{aligned} P_{\text{refl}} = E_0^2 \, [ & J_0^2(\Gamma) r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega) \\ & + J_1^2(\Gamma)(|r_{\text{cav}}(\omega + \Omega)|^2 + |r_{\text{cav}}(\omega - \Omega)|^2) \\ & - i J_0(\Gamma) J_1(\Gamma) e^{i\Omega t} r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega + \Omega) + i J_0(\Gamma) J_1(\Gamma) e^{-i\Omega t} r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega - \Omega) \\ & - i J_0(\Gamma) J_1(\Gamma) e^{i\Omega t} r_{\text{cav}}(\omega) r^*_{\text{cav}}(\omega - \Omega) + i J_0(\Gamma) J_1(\Gamma) e^{-i\Omega t} r_{\text{cav}}(\omega) r^*_{\text{cav}}(\omega + \Omega) ] \end{aligned} \tag{B.10}$$

Rearranging, this is

$$\begin{aligned} P_{\text{refl}} = E_0^2\, [&J_0^2(\Gamma) r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega) \\ &+ J_1^2(\Gamma)(|r_{\text{cav}}(\omega+\Omega)|^2 + |r_{\text{cav}}(\omega-\Omega)|^2) \\ &+ iJ_0(\Gamma)J_1(\Gamma)e^{-i\Omega t}(r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega-\Omega) + r^*_{\text{cav}}(\omega+\Omega)r_{\text{cav}}(\omega)) \\ &- iJ_0(\Gamma)J_1(\Gamma)e^{i\Omega t}(r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega+\Omega) + r^*_{\text{cav}}(\omega-\Omega)r_{\text{cav}}(\omega))] \end{aligned} \tag{B.11}$$

We would like to simplify this equation into something proportional to either cos(Ω*t*) or sin(Ω*t*). For this, we need only examine the cross terms on the last 2 lines of Equation B.11. For ease of notation, we will let

$$A \equiv r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega-\Omega) \qquad \text{and} \qquad B \equiv r^*_{\text{cav}}(\omega) r_{\text{cav}}(\omega+\Omega). \tag{B.12}$$

The cross terms of Equation B.11 are then

$$\text{c.t.} = iJ_0(\Gamma)J_1(\Gamma)e^{-i\Omega t}\,(A + B^*) - iJ_0(\Gamma)J_1(\Gamma)e^{i\Omega t}\,(B + A^*) \tag{B.13}$$

which, moving around and expanding gives

$$\text{c.t.} = J_0(\Gamma)J_1(\Gamma)\left(-\frac{1}{i}Ae^{-i\Omega t} - \frac{1}{i}B^*e^{-i\Omega t} + \frac{1}{i}Be^{i\Omega t} + \frac{1}{i}A^*e^{i\Omega t}\right). \tag{B.14}$$

Note that, since *A* and *B* are complex numbers,

$$A \equiv |A|e^{i\alpha}, \qquad A^* \equiv |A|e^{-i\alpha}, \qquad B \equiv |B|e^{i\beta}, \qquad B^* \equiv |B|e^{-i\beta}. \tag{B.15}$$

Utilizing this,

$$\text{c.t.} = J_0(\Gamma)J_1(\Gamma)\left(|A|\frac{1}{i}\left(e^{i(\Omega t-\alpha)} - e^{-i(\Omega t-\alpha)}\right) + |B|\frac{1}{i}\left(e^{i(\Omega t+\beta)} - e^{-i(\Omega t+\beta)}\right)\right). \tag{B.16}$$

Recall that

$$\sin(x) = \frac{e^x - e^{-x}}{2i} \tag{B.17}$$

to simplify the cross terms to

$$\text{c.t.} = J_0(\Gamma)J_1(\Gamma)\left[2|A|\sin(\Omega t-\alpha) + 2|B|\cos(\Omega t+\beta)\right]. \tag{B.18}$$

Using the fact that

$$\sin(u \pm v) = \sin(u)\cos(v) \pm \cos(u)\sin(v) \tag{B.19}$$

we can expand the cross terms to

$$\text{c.t.} = J_0(\Gamma)J_1(\Gamma)\left[2|A|\left(\sin(\Omega t)\cos(\alpha) - \cos(\Omega t)\sin(\alpha)\right) + 2|B|\left(\sin(\Omega t)\cos(\beta) + \cos(\Omega t)\sin(\beta)\right)\right] \tag{B.20}$$

which rearranges to

$$\text{c.t.} = J_0(\Gamma)J_1(\Gamma)\left[2\sin(\Omega t)\left(|A|\cos(\alpha) + |B|\cos(\beta)\right) + 2\cos(\Omega t)\left(|B|\sin(\beta) - |A|\sin(\alpha)\right)\right]. \tag{B.21}$$

Recalling that, for an arbitrary complex number $z = |z|e^{i\zeta}$,

$$\mathrm{Re}(z) = |z|\cos(\zeta) \qquad \text{and} \qquad \mathrm{Im}(z) = |z|\sin(\zeta), \tag{B.22}$$

we can simplify one final time to find

$$\text{c.t.} = 2J_0(\Gamma)J_1(\Gamma)\left[\sin(\Omega t)\mathrm{Re}(A + B) + \cos(\Omega t)\mathrm{Im}(-A + B)\right]. \tag{B.23}$$

Putting this back together with Equation B.11, we have

$$P_{\text{refl}} = E_0^2 \begin{bmatrix} J_0^2(\Gamma)r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega) \\ +J_1^2(\Gamma)(|r_{\text{cav}}(\omega+\Omega)|^2 + |r_{\text{cav}}(\omega-\Omega)|^2) \\ +2J_0(\Gamma)J_1(\Gamma)\begin{bmatrix} \sin(\Omega t)\mathrm{Re}(r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega-\Omega) + r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega+\Omega)) \\ +\cos(\Omega t)\mathrm{Im}(r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega+\Omega) - r^*_{\text{cav}}(\omega)r_{\text{cav}}(\omega-\Omega)) \end{bmatrix} \end{bmatrix} \tag{B.24}$$

Recall that for Equation B.4 we have restricted ourselves to $|n| < 1$. In general, and in particular for the 3f analysis discussed in Section 4.4, we will need to keep several more terms in the series. In that case, Equation B.24 must be expanded to include these terms.

Equation B.24 is the power incident on the photodiode. Some amount of photons will, via the photoelectric effect create a “photocurrent”, which is just the current of electrons ejected by the photons. The fraction of photons that will eject electrons is described by the quantum efficiency of the photodiode, and the responsivity of the photodiode which has units of Amps/Watt. The transimpedance of the electronics surrounding the diode will convert the current to a voltage. The transimpedance need not be the same for all frequencies (and in general LIGO uses resonant photodiodes to emphasize the $1\Omega$ components, and de-emphasize other order harmonics), but for the rest of this analysis we will ignore this frequency dependence. So, the output of the photodiode and accompanying electronics will be

$$V_{\text{refl}} \propto P_{\text{refl}}. \tag{B.25}$$

We are next interested in demodulating this voltage to a lower frequency, so that it is a more tractable error signal. We use the same frequency as the original RF modulation for to create a "local oscillator" reference signal. To maintain generality, we must multiply by either a sine or a cosine, to extract all of the information from $V_{\mathrm{refl}}$. To create both the sine and cosine we will split the local oscillator signal, and phase shift one component by 90°. We lowpass the resulting signal to eliminate the high frequency terms, to get (assuming the constants are absorbed in the proportionality factor of Equation B.25)

$$V^{I}_{\mathrm{refl}} \simeq J_0(\Gamma)J_1(\Gamma)\,\mathrm{Re}(r^*_{\mathrm{cav}}(\omega)r_{\mathrm{cav}}(\omega+\Omega)+r^*_{\mathrm{cav}}(\omega)r_{\mathrm{cav}}(\omega-\Omega)) \tag{B.26}$$

and

$$V^{Q}_{\mathrm{refl}} \simeq J_0(\Gamma)J_1(\Gamma)\,\mathrm{Im}(r^*_{\mathrm{cav}}(\omega)r_{\mathrm{cav}}(\omega+\Omega)-r^*_{\mathrm{cav}}(\omega)r_{\mathrm{cav}}(\omega-\Omega)). \tag{B.27}$$

The version of $V_{\mathrm{refl}}$ that contains the real parts of $r_{\mathrm{cav}}$ is referred to as "in-phase", or the "I-phase", which has been denoted $V^{I}_{\mathrm{refl}}$. The other component which contains the imaginary parts of $r_{\mathrm{cav}}$ is referred to as the "quadrature phase", or the "Q-phase", which has been denoted $V^{Q}_{\mathrm{refl}}$. Note that often, the component resulting from $V_{\mathrm{refl}}$ being multiplied by a cosine will result in the I-phase component of the error signal, and the component resulting from $V_{\mathrm{refl}}$ being multiplied by a sine will result in the Q-phase. This difference in notation is a result of the choice to use sine as the modulation in Equation B.2 rather than cosine. If the phase of the local oscillator is set properly, the Q-phase signal will vanish, and all of the information about the length of the cavity relative to the laser frequency will be contained in the I-phase signal.

# Appendix C

# Layout of the 40 m Interferometer

This appendix contains to-scale CAD drawings of the optical layout of the 40m Prototype, for components within the vacuum system. Also included are photos of various optical tables, in order to record the current configuration of the interferometer.

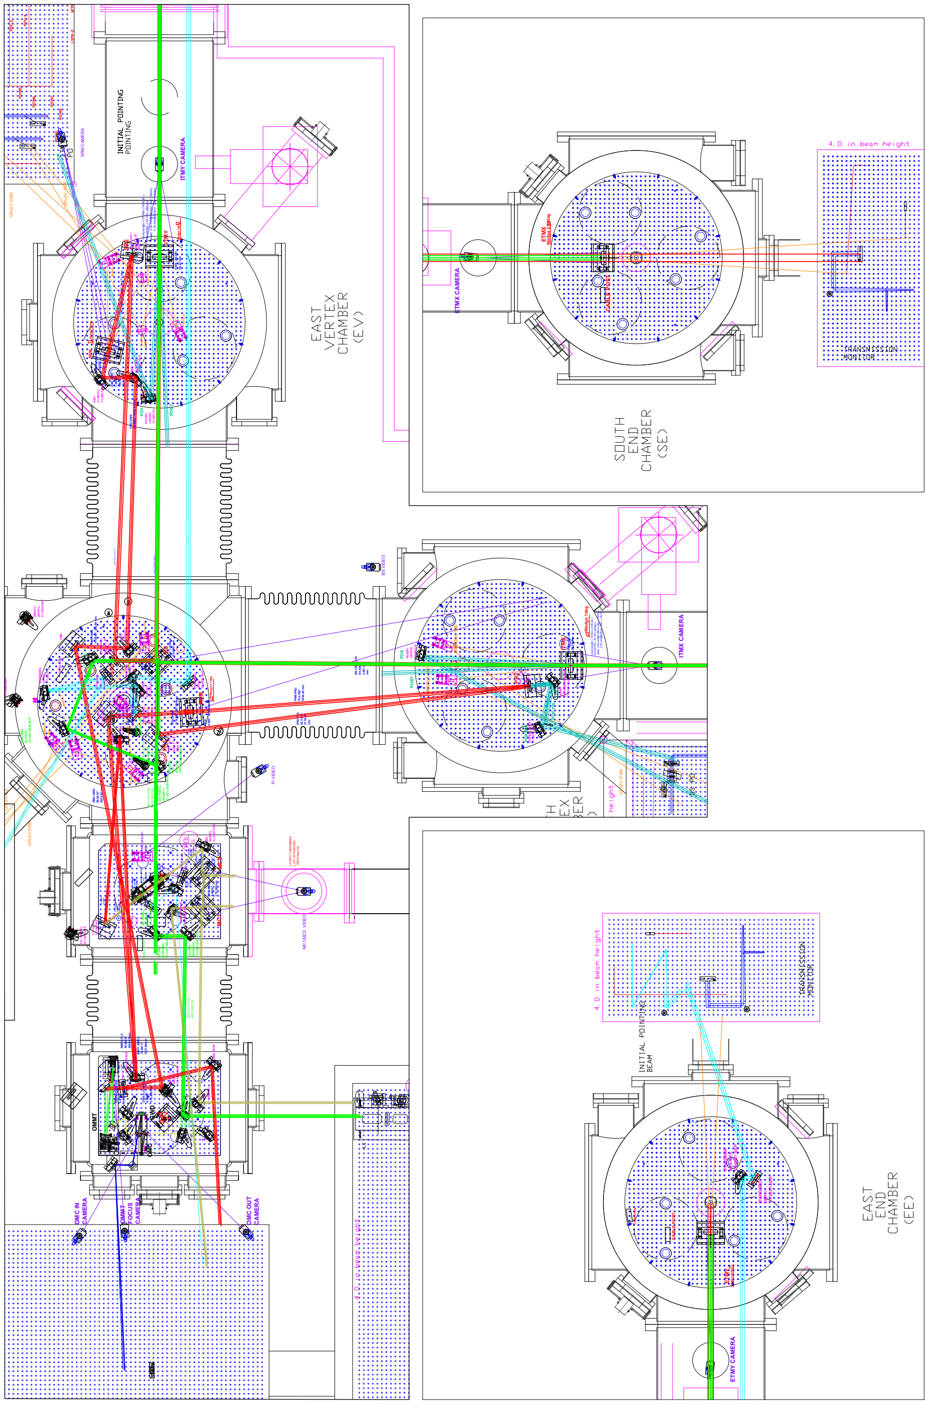

Figure C.1: Overview of optical layout.

## C.1 In-vacuum optical tables

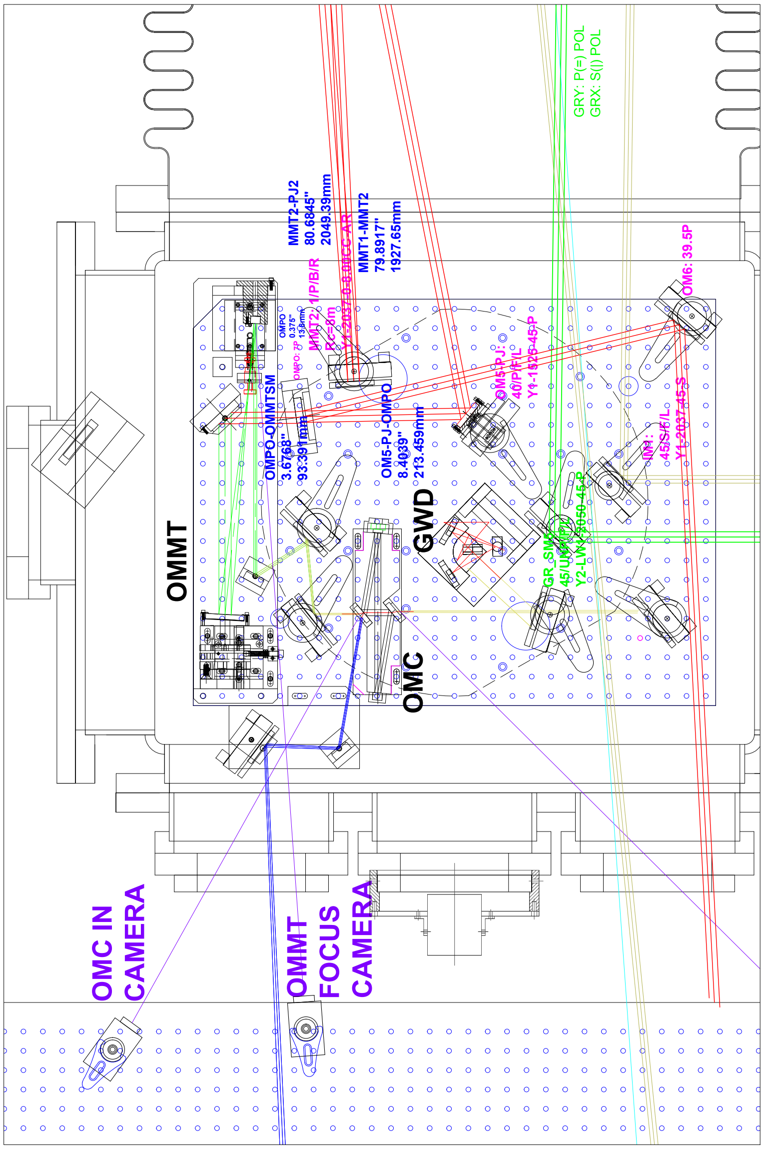

Figure C.2: Detail of output mode cleaner table. The output mode cleaner (and associated optical path) is not utilized; however, this table holds one of the curved input mode matching telescope mirrors, as well as steering to direct the main beam in and out of the vacuum envelope.

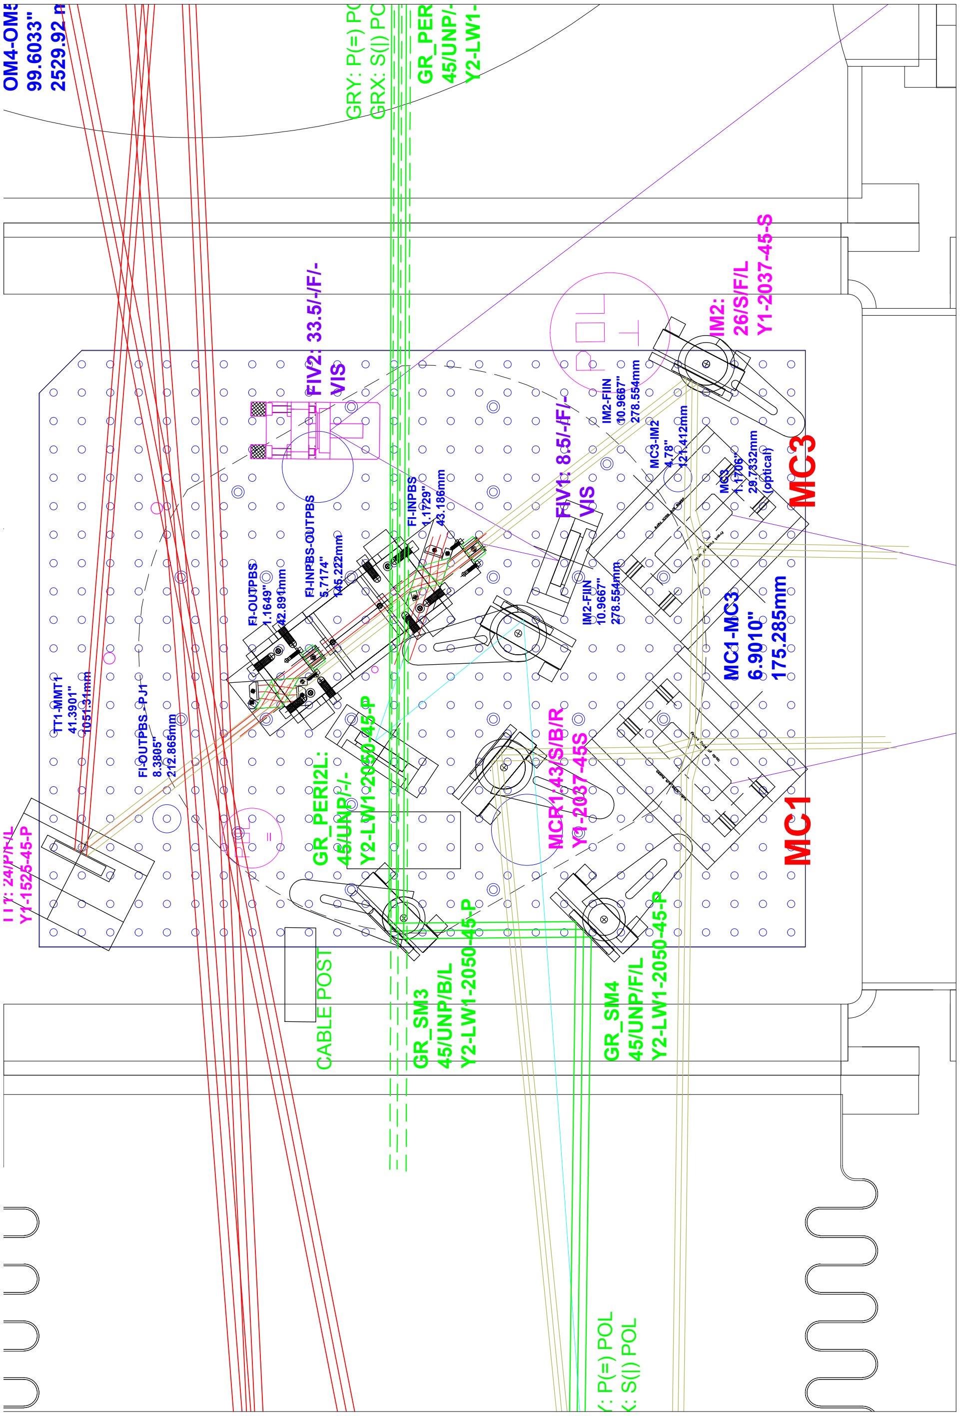

Figure C.3: Detail of input mode cleaner table. MC1 and MC3 sit on this table, while the curved mode cleaner mirror, MC2, sits 13.5 m away (not shown, to the right).

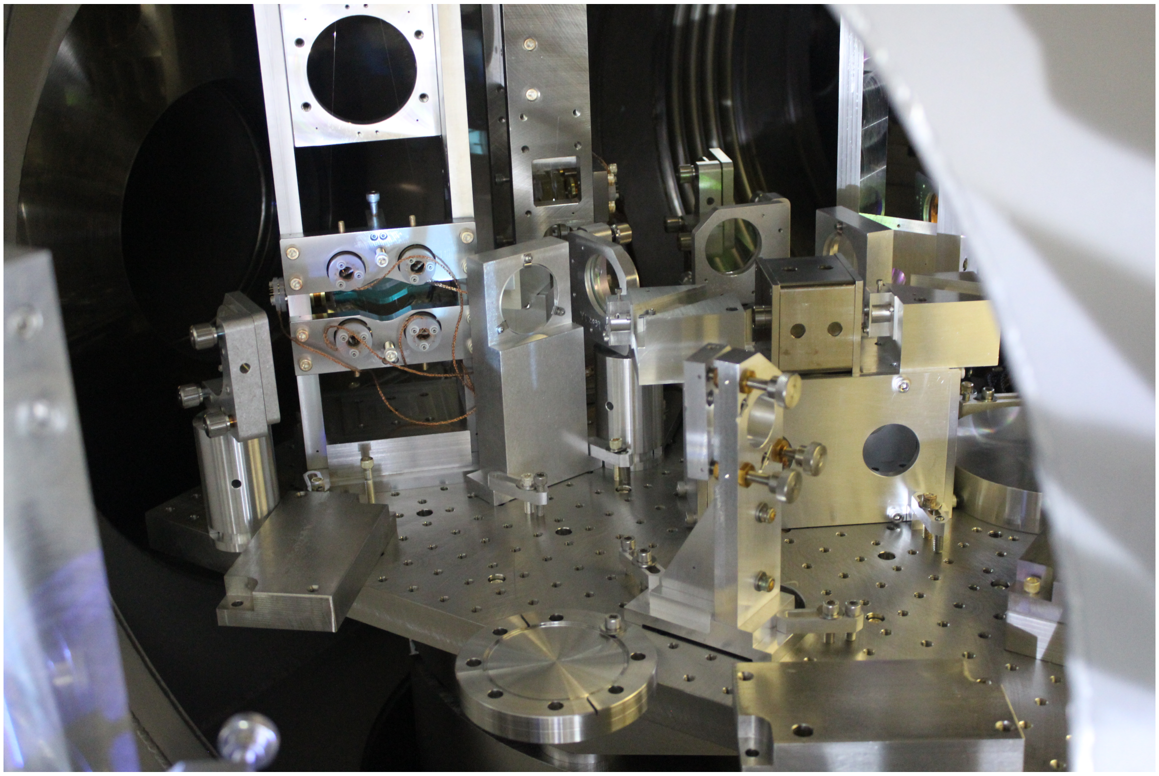

Figure C.4: Input mode cleaner table, as seen from the beam splitter chamber.

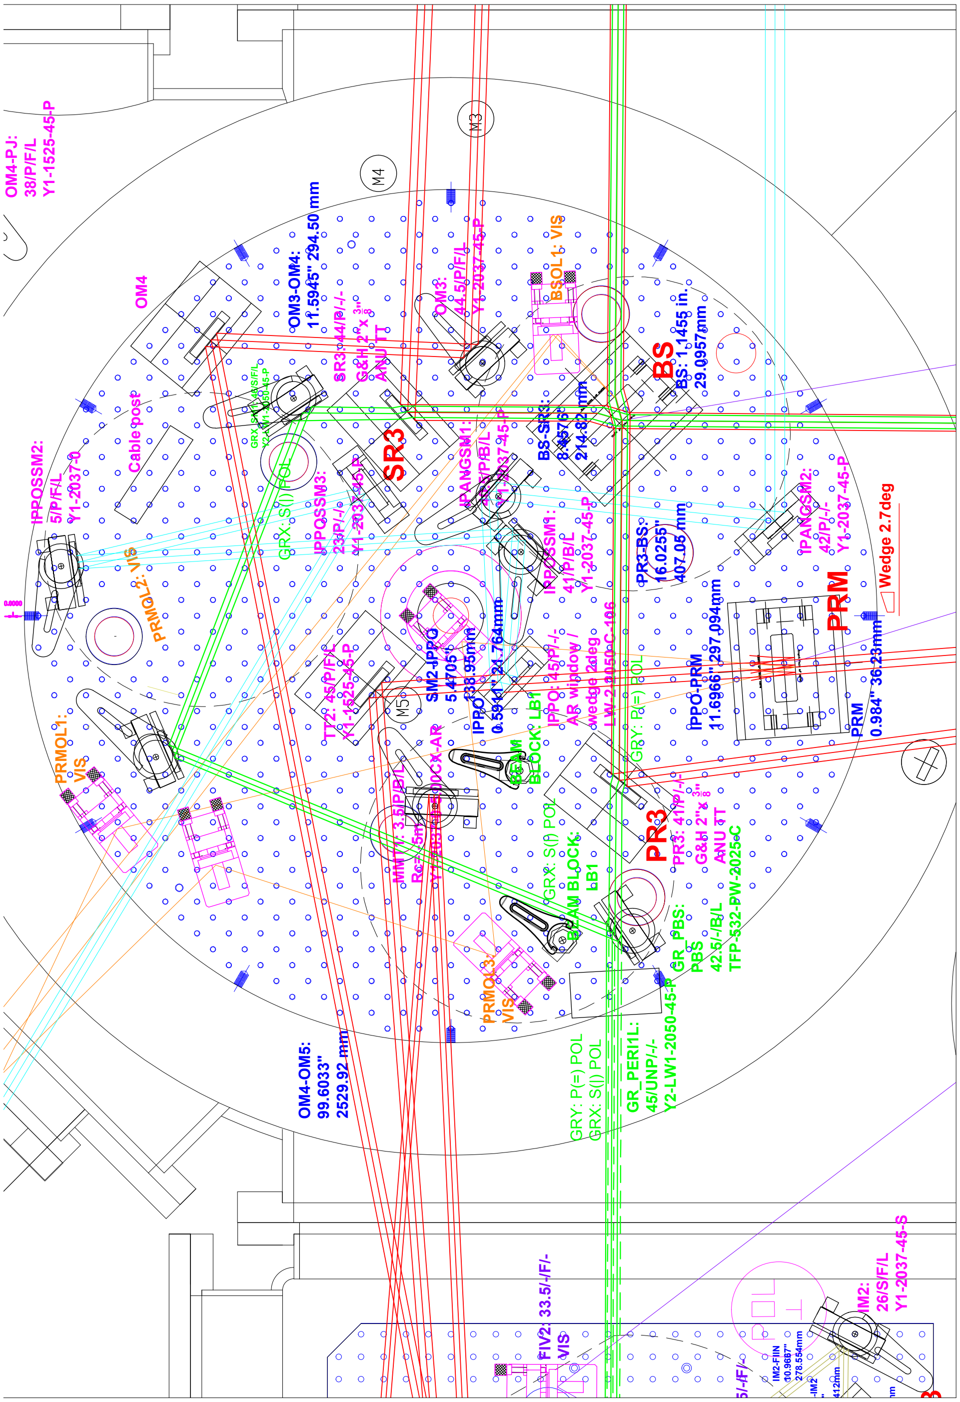

Figure C.5: Detail of beam splitter table. The beam splitter table is the most crowded in-vacuum table at the 40m, holding input optics (steering and mode matching), the power recycling mirror, a power recycling cavity folding mirror, a signal recycling cavity folding mirror, output steering optics, auxiliary optics for optical lever, and main beam pointing monitoring, as well as the beam splitter itself.

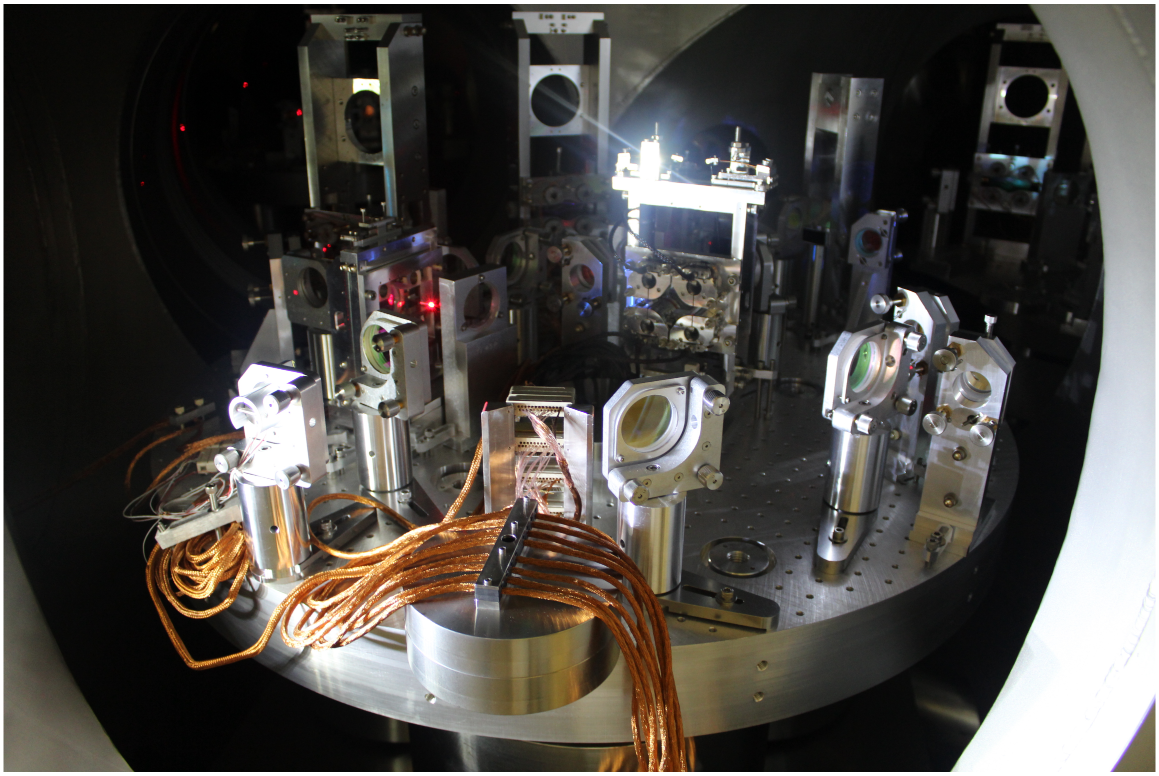

Figure C.6: Photo of beam splitter table.

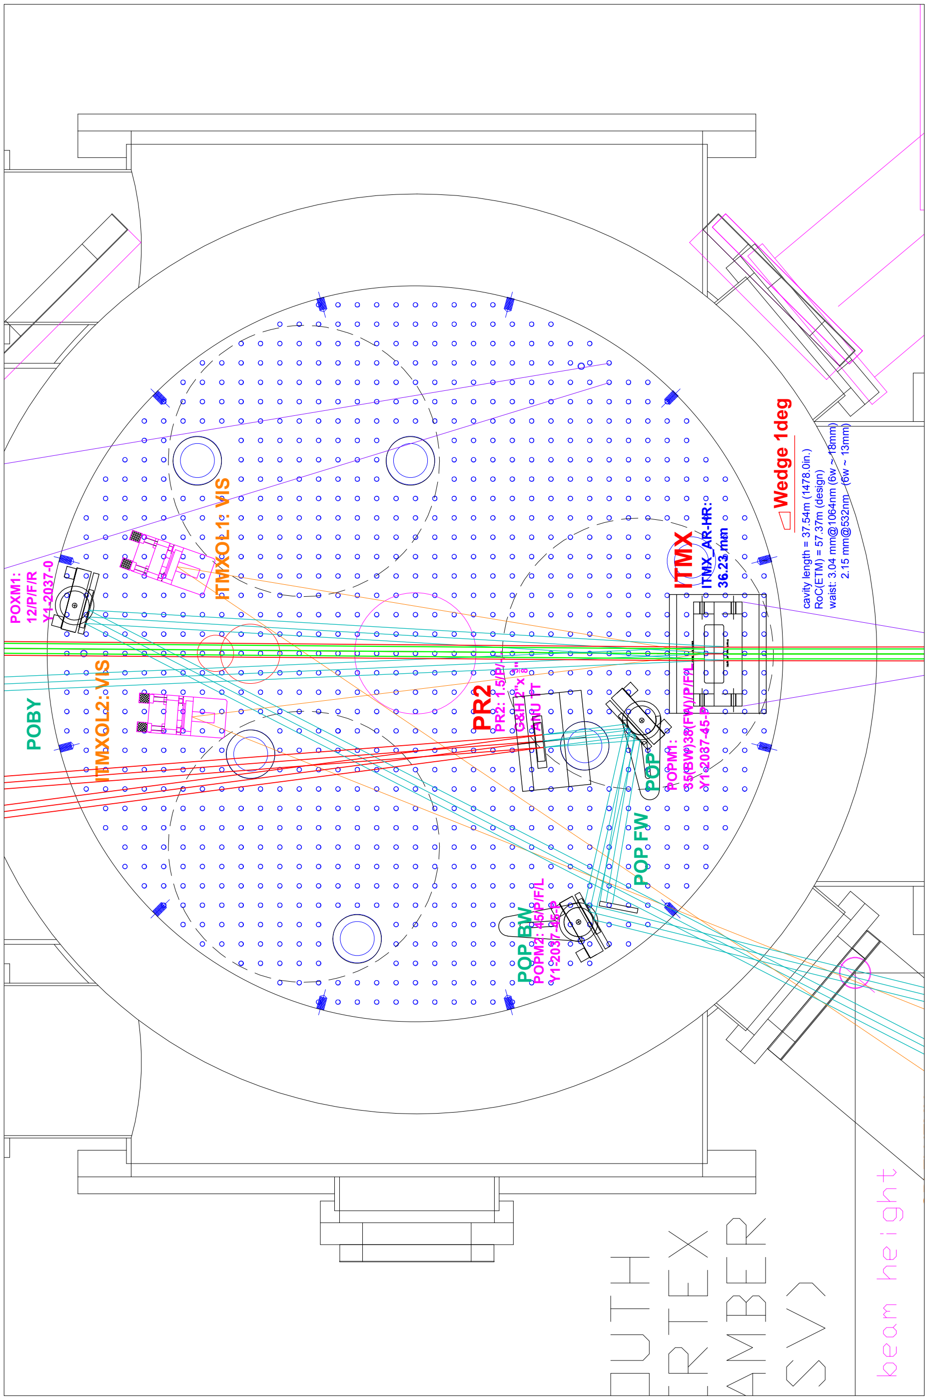

Figure C.7: Detail of the ITMX table. In addition to the input test mass for the X-arm, this table holds a power recycling cavity folding mirror, and optics to extract the power recycling cavity pick-off beam ("POP").

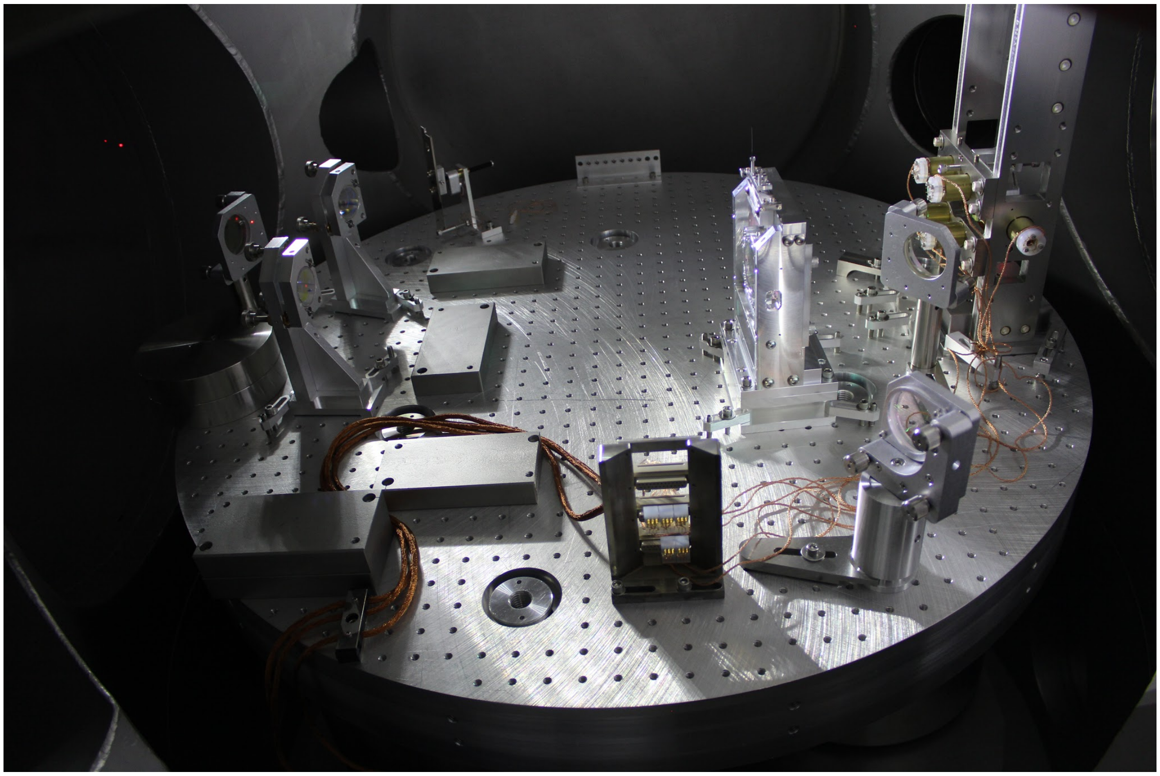

Figure C.8: Photo of ITMX chamber.

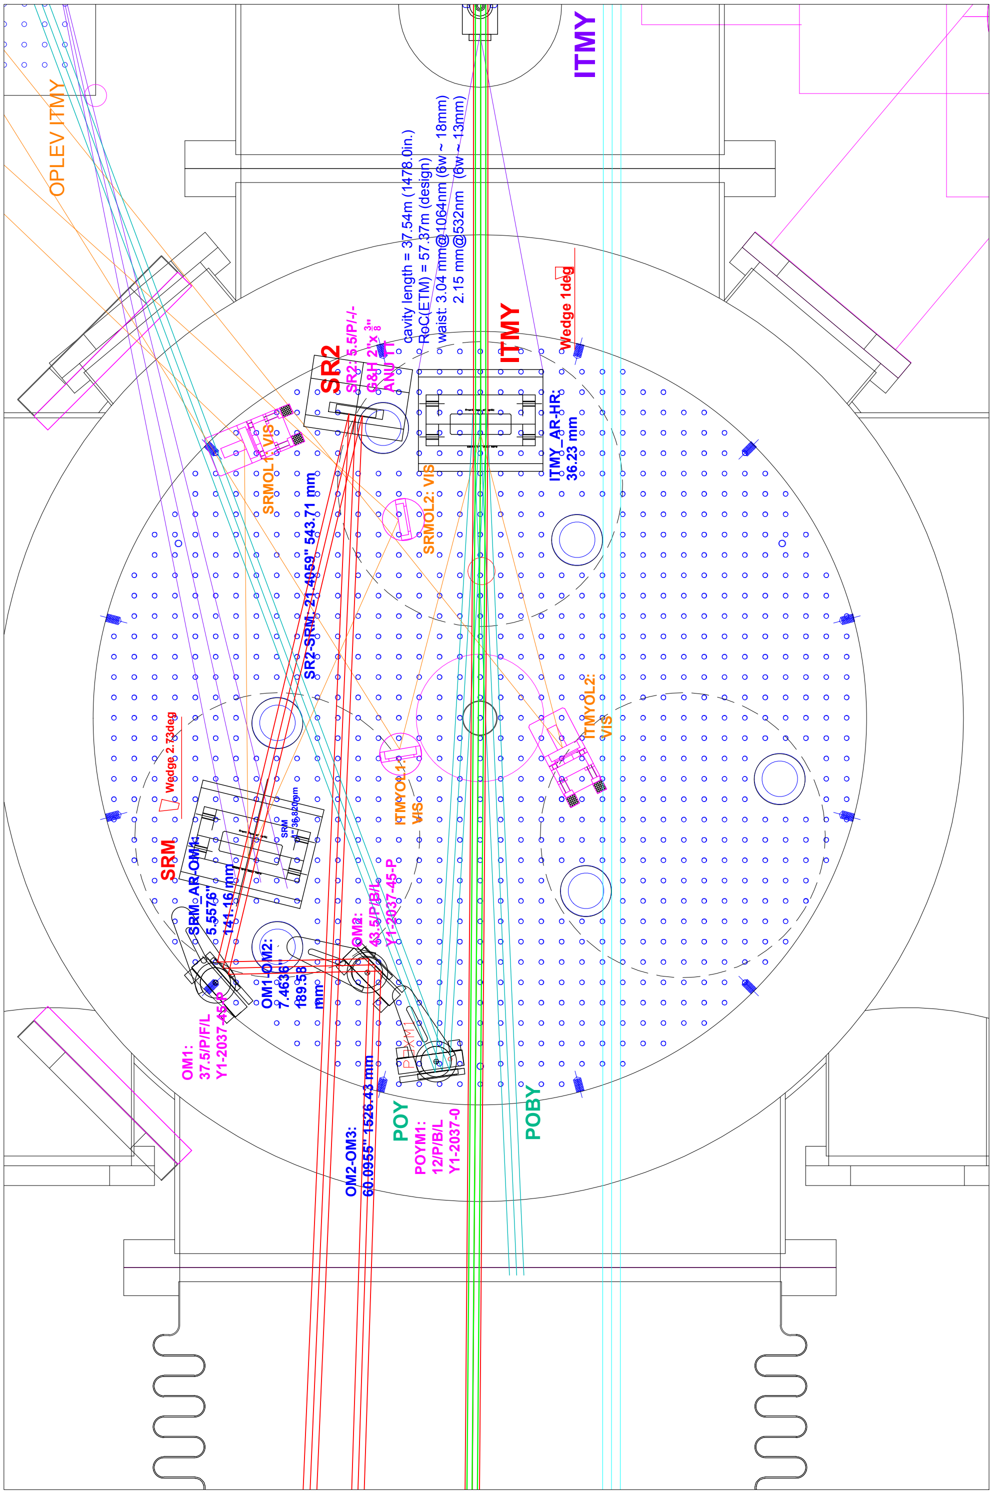

Figure C.9: Detail of the ITMY table. This table holds the signal recycling mirror and a signal recycling cavity folding mirror, as well as output path steering optics and the input test mass for the Y-arm.

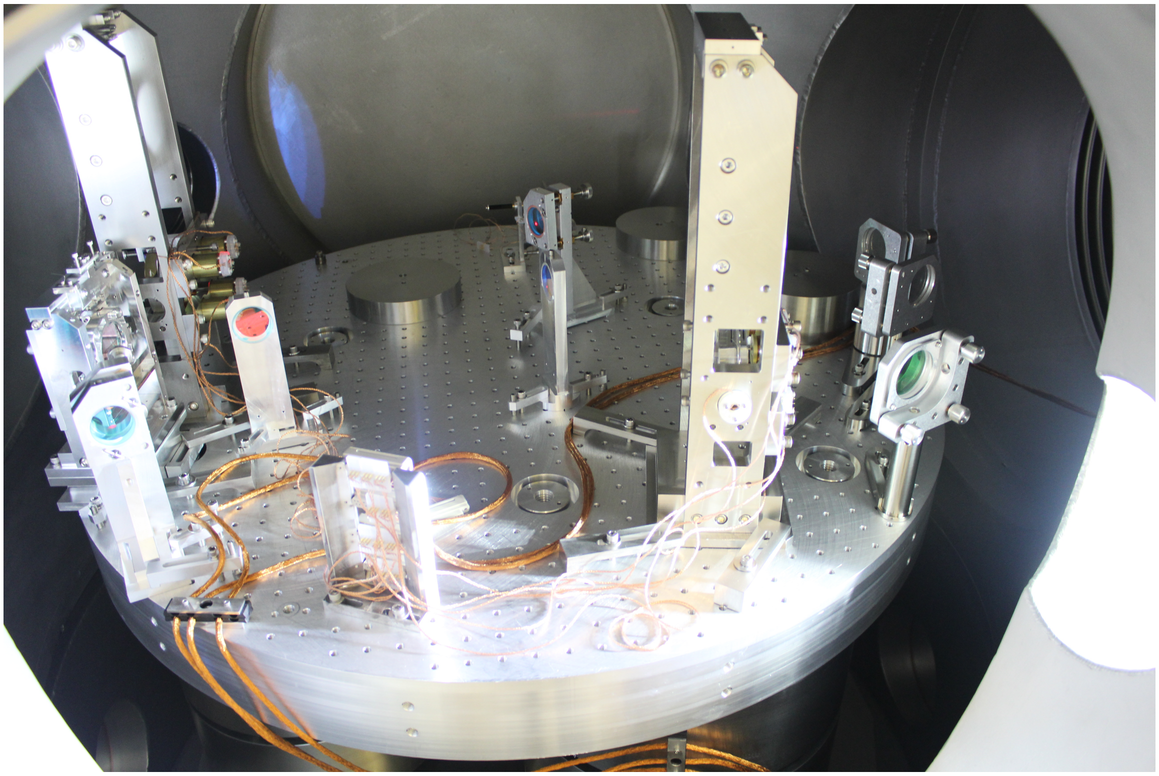

Figure C.10: Photo of ITMY chamber.

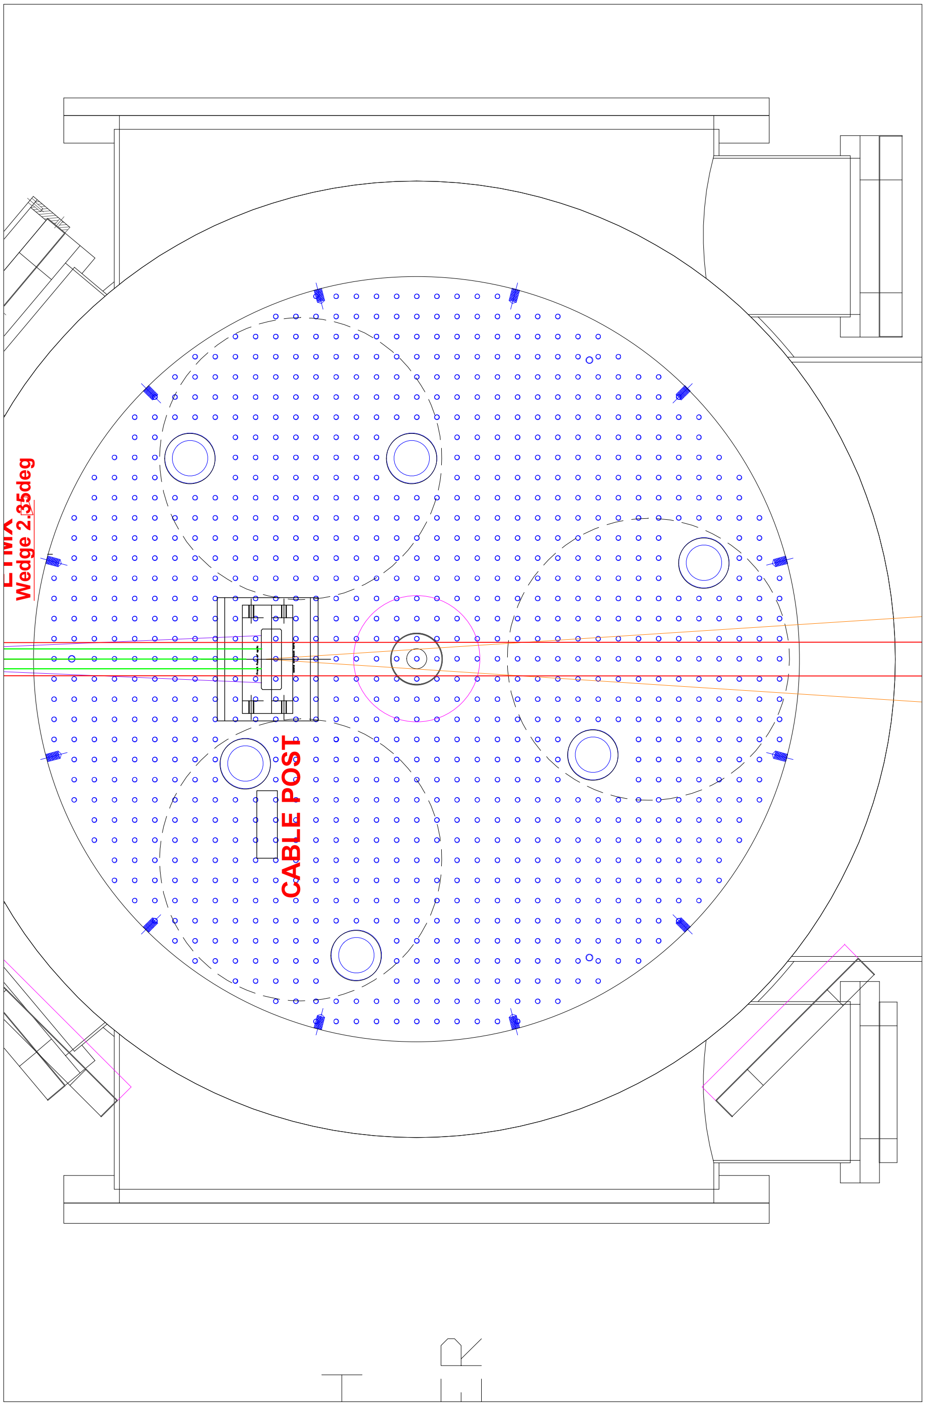

Figure C.11: Detail of ETMX table. Not shown is a black glass baffle, located several inches in front of the ETM to prevent stray scattered light from hitting the test mass.

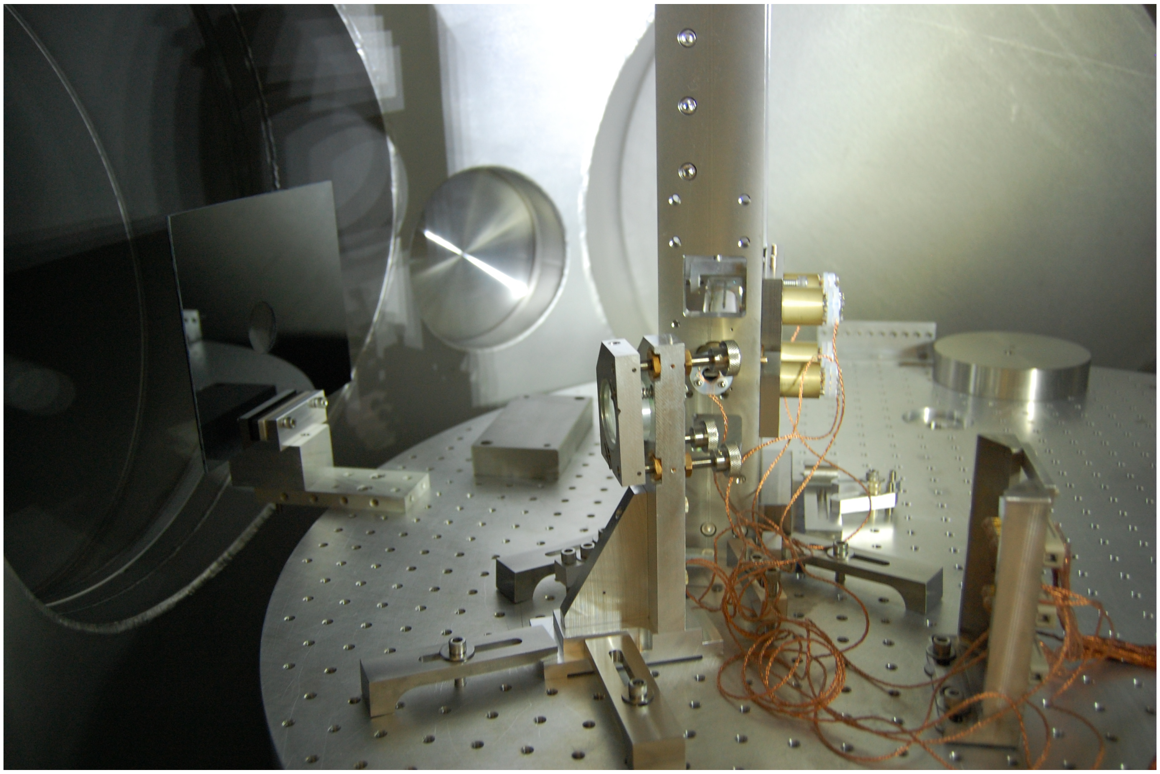

Figure C.12: Photo of ETMX chamber.

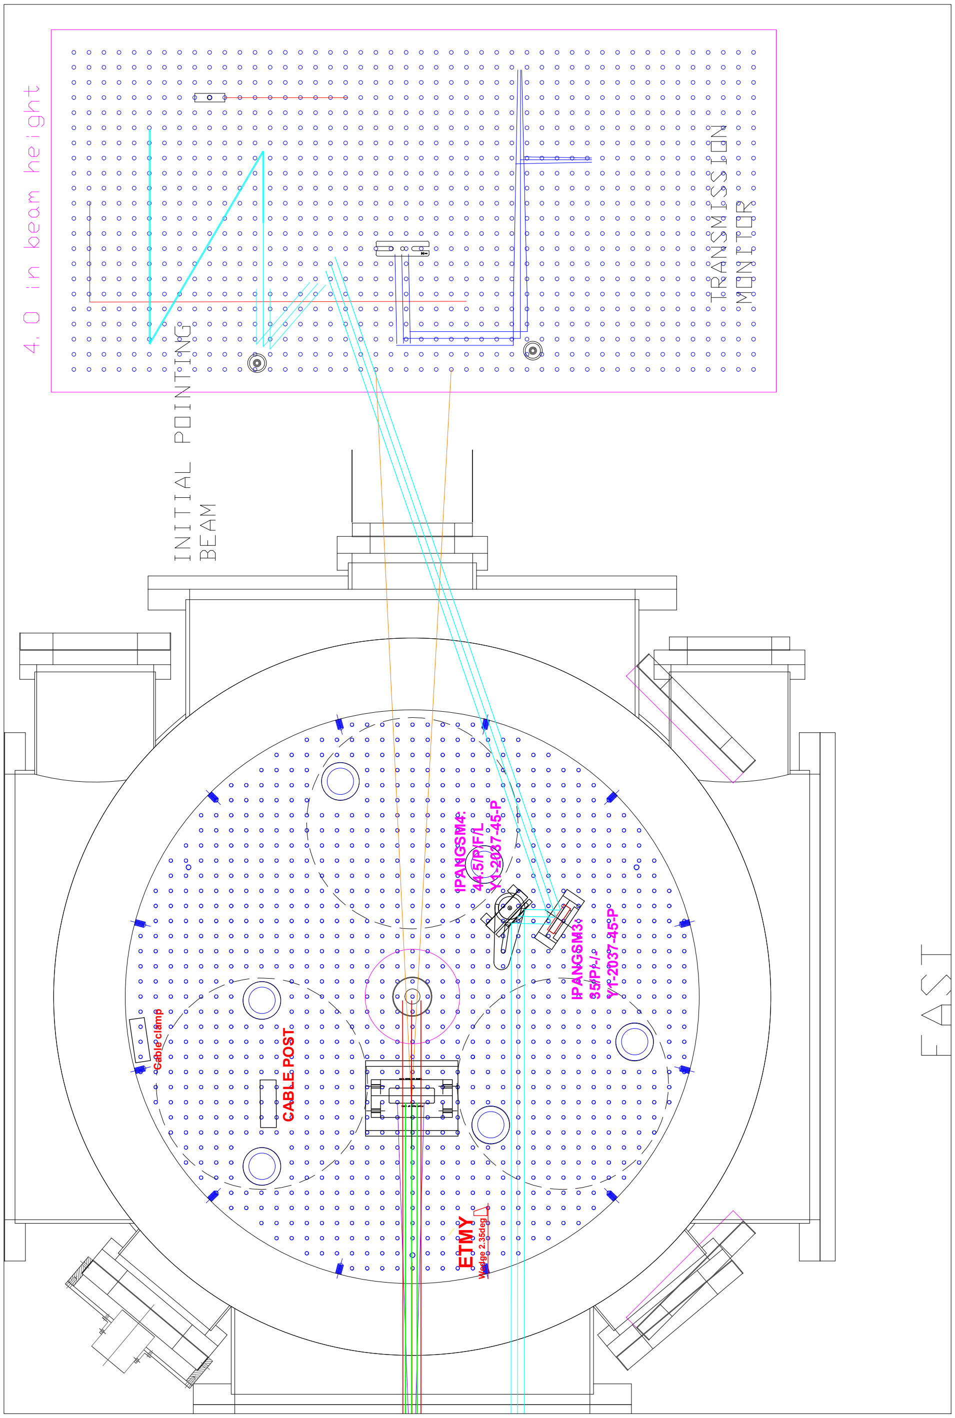

Figure C.13: Detail of ETMY table. This table holds steering optics for the long lever arm input beam pointing monitor. Not shown is a black glass baffle, located several inches in front of the ETM to prevent stray scattered light from hitting the test mass.

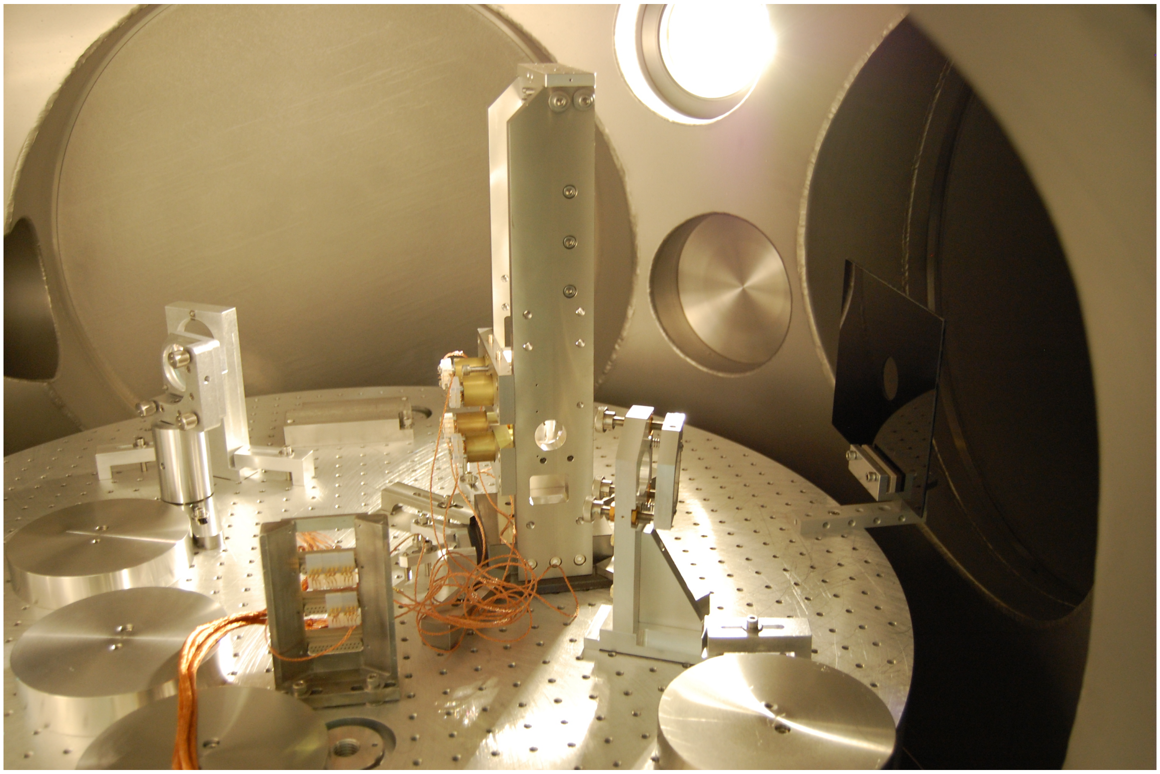

Figure C.14: Photo of ETMY chamber.

## C.2 In-air optical tables

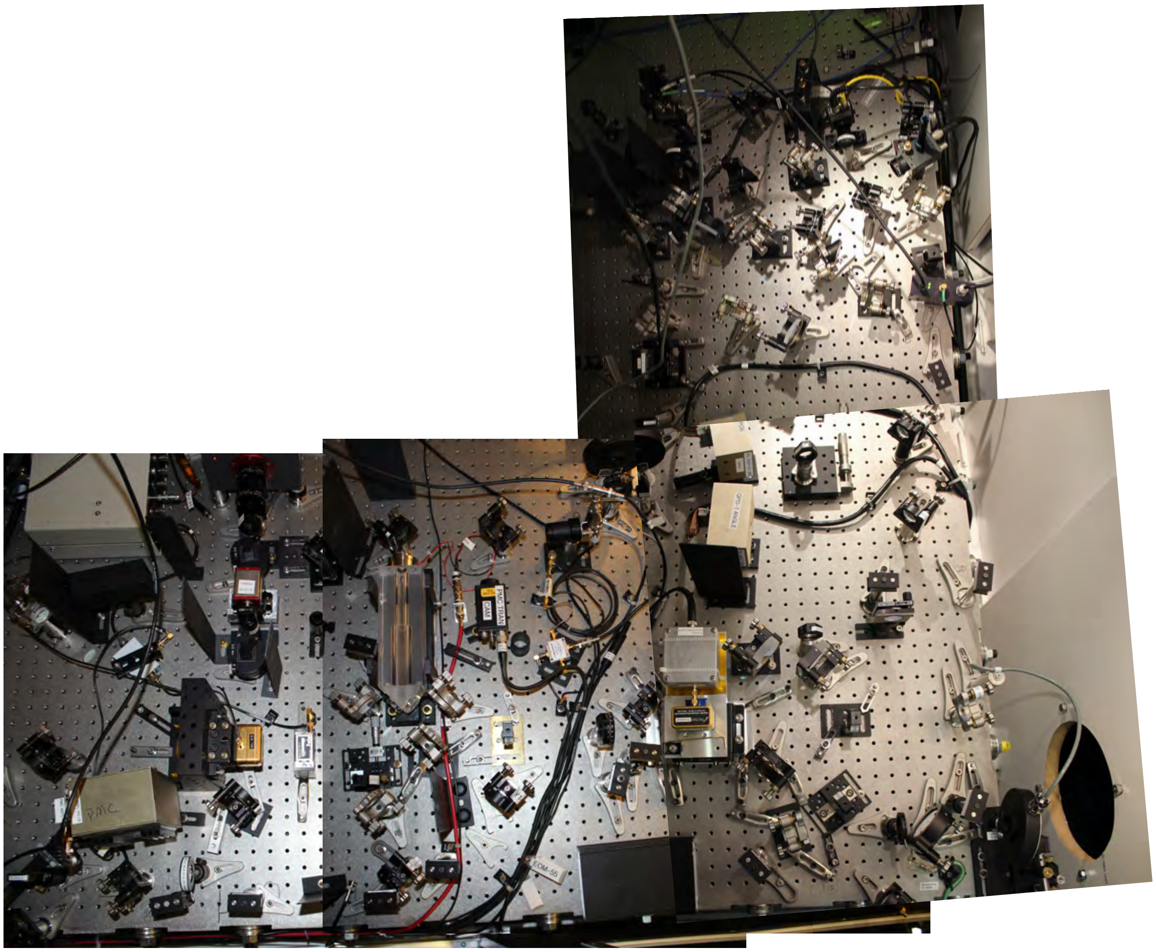

Figure C.15: Pre-stabilized laser table.

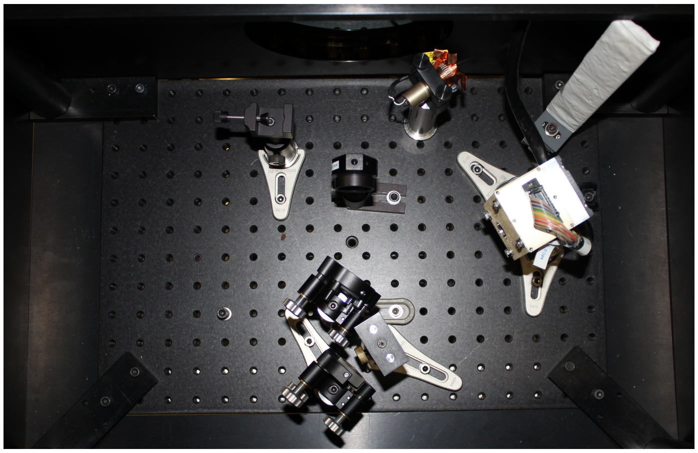

Figure C.16: MC2 transmission table.

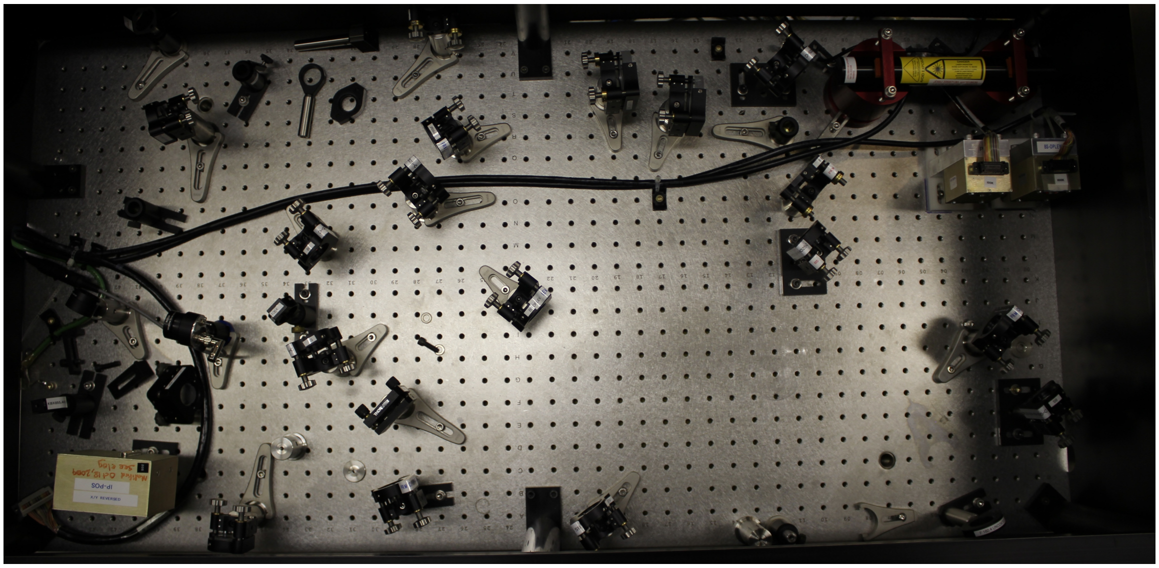

Figure C.17: BS-PRM optical lever table.

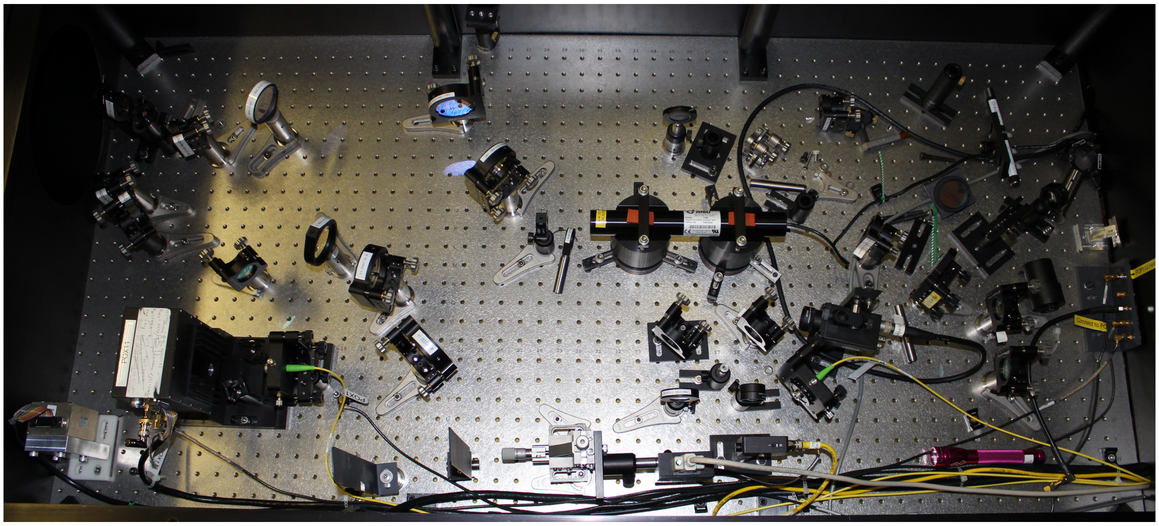

Figure C.18: ITMX optical lever table. Also holds POP photodiodes.

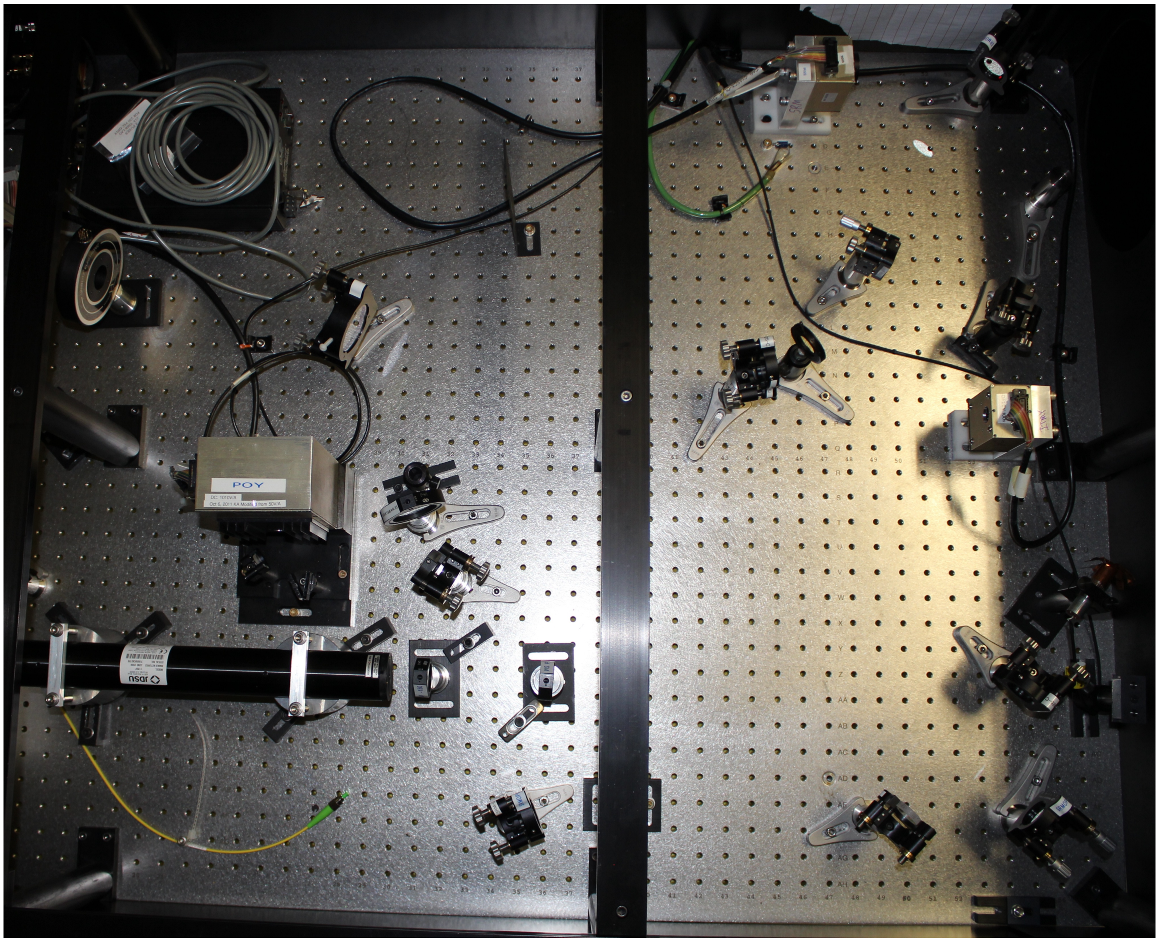

Figure C.19: ITMY optical lever table which also holds SRM optical lever.

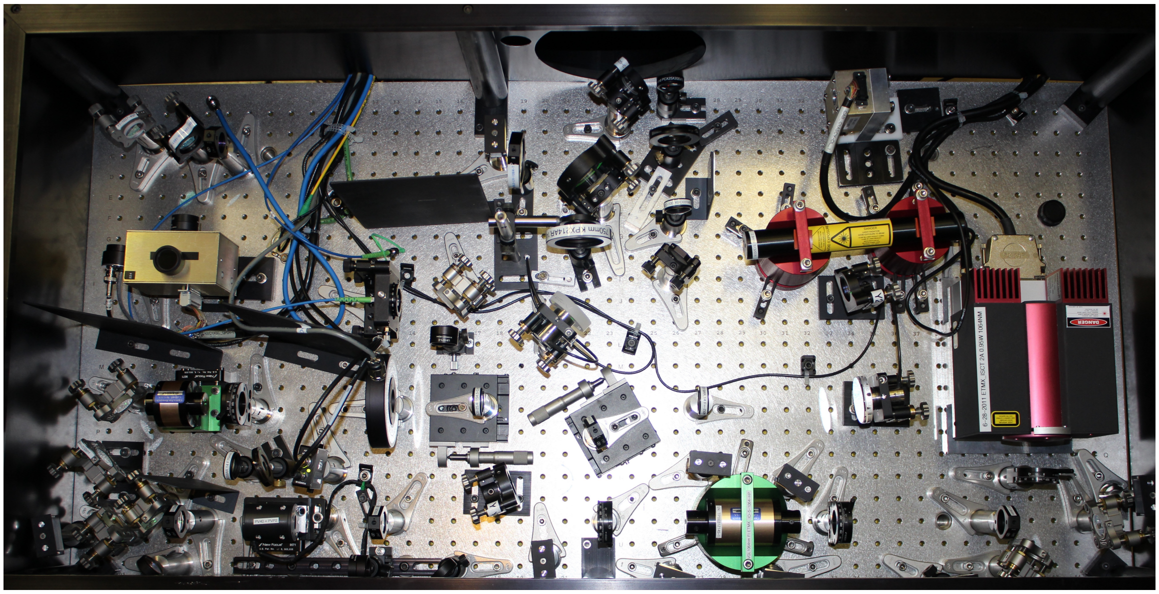

Figure C.20: Transmission table for X-arm, which also holds injection optics for green laser. Table enclosure is prototype of acoustic and air current isolation.

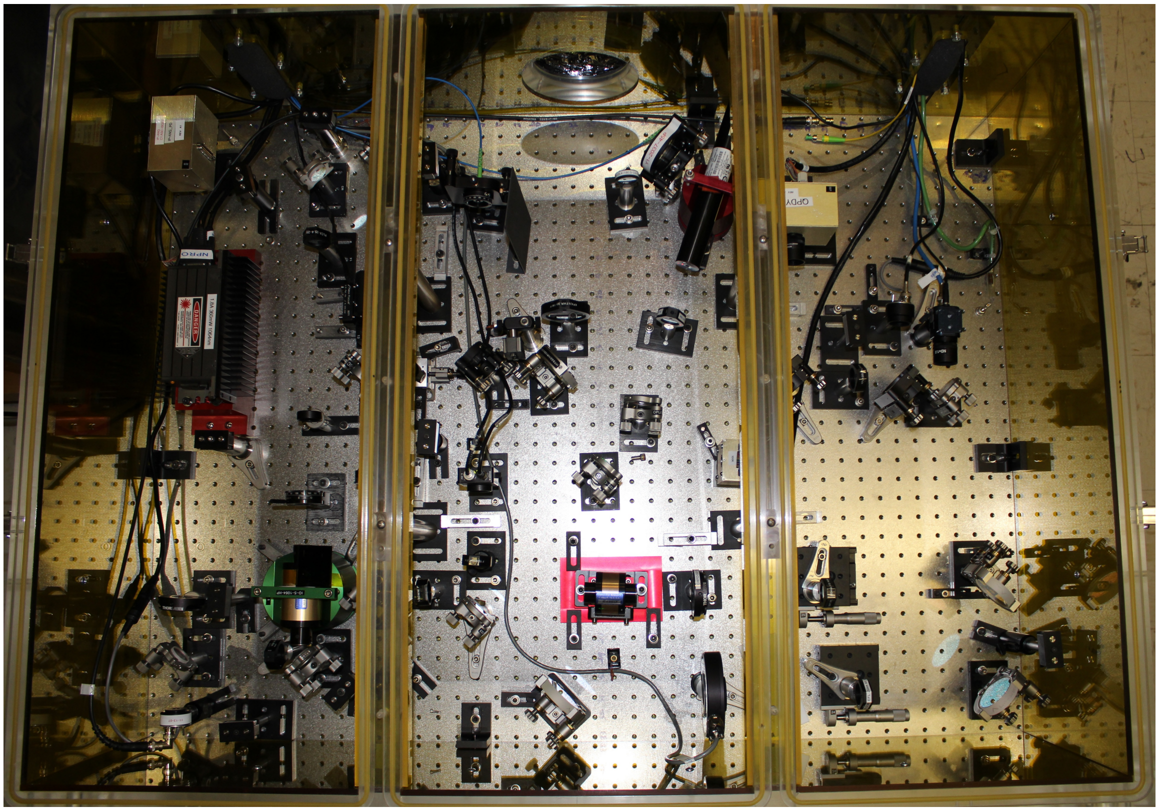

Figure C.21: Transmission table for Y-arm, which also holds injection optics for green laser.

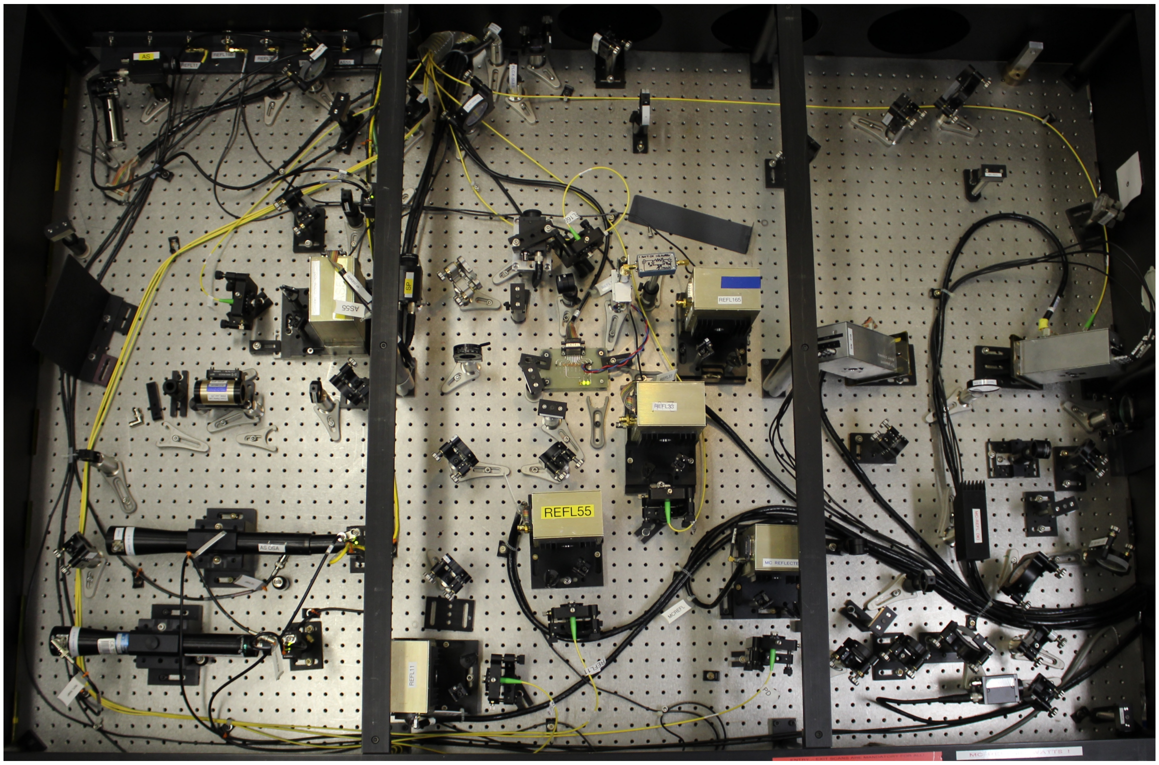

Figure C.22: Main detection table, containing input mode cleaner reflected port and main interferometer reflected port, as well as anti-symmetric port.

# Bibliography

[1] A. Einstein. Die grundlage der allgemeinen relativitätstheorie (The foundation of the general theory of relativity). *Annalen der Physik*, 354(7):769–822, 1916.

[2] J. Weber. Gravitational-wave-detector events. *Phys. Rev. Lett.*, 20(1307), 1968.

[3] Robert L. Forward. Wideband laser-interferometer graviational-radiation experiment. *Phys. Rev. D*, 17(379), 1978.

[4] G. Hobbs, et al. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Class. Quantum Grav.*, 27:084013, 2010.

[5] B. P. Abbott, et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Rep. Prog. Phys.*, 72:076901, 2009.

[6] Ryan DeRosa, Jennifer C. Driggers, Dani Atkinson, Haixing Miao, Valery Frolov, Michael Landry, Joseph A. Giaime, and Rana X Adhikari. Global feed-forward vibration isolation in a km scale interferometer. *Class. Quantum Grav.*, 29:215008, 2012.

[7] Jennifer C. Driggers, Matthew Evans, Keenan Pepper and Rana Adhikari. Active noise cancellation in a suspended interferometer. *Rev. Sci. Instrum.*, 83:024501, 2012.

[8] J. Driggers and J. Harms. Results of phase 1 Newtonian noise measurements at the LIGO sites, February-March 2011. Technical report, LIGO-T1100237, 2011.

[9] Jennifer C. Driggers, Jan Harms, and Rana X Adhikari. Subtraction of Newtonian noise using optimized sensor arrays. *Phys. Rev. D*, 86:102001, 2012.

[10] J. Aasi, et al. Prospects for localization of gravitational wave transients by the Advanced LIGO and Advanced Virgo observatories. *LIGO-P1200087*.

[11] F. Acernese, et al. Advanced Virgo: a 2nd generation interferometric gravitational wave detector. *Class. Quantum Grav.*, 32:024001, 2015.

[12] Kentaro Somiya. Detector configuration of KAGRA – the Japanese cryogenic gravitational-wave detector. *Class. Quantum Grav.*, 29:124007, 2012.

[13] J. Abadie, et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Class. Quantum Grav.*, 27:173001, 2010.

[14] B. Allen, et al. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85:122006, 2012.

[15] Luc Blanchet, Bala R. Iyer, Clifford M. Will and Alan G. Wiseman. Gravitational waveforms from inspiralling compact binaries to second-post-Newtonian order . *Class. Quantum Grav.*, 13:575, 1996.

[16] C. W. Misner, K. S. Thorne and J. A. Wheeler. *Gravitation*. Freeman, 1973.

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

[18] Eva Noyola, Karl Gebhardt, and Marcel Bergmann. Gemini and Hubble Space Telescope evidence for an intermediate-mass black hole in Omega Centauri. *ApJ*, 676(1008), 2008.

[19] Jay Anderson and Roeland P. van der Marel. New limits on an intermediate-mass black hole in Omega Centauri II: Dynamical models. *ApJ*, 710:1063, 2010.

[20] W. David Arnett and Richard L. Bowers. A microscopic interpretation of neutron star structure. *Astrophys. J. Suppl.*, 33:415, 1977.

[21] Christian D. Ott. The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26:063001, 2009.

[22] Luca Baiotti, Ian Hawke and Luciano Rezzolla. On the gravitational radiation from the collapse of neutron stars to rotating black holes. *Class. Quantum Grav.*, 24:S187, 2007.

[23] J. Abadie, et al. All-sky search for gravitational-wave bursts in the second joint LIGO-Virgo run. *Phys. Rev. D*, 85:122007, 2012.

[24] Rai Weiss. Collection of reports on Barkhausen noise. Technical report, LIGO-T0900061, 2008.

[25] H. Barkhausen. Zwei mit Hilfe der Neuen Verstärker entdeckte Erscheinungen (Two phenomena discovered using new amplifiers). *Physikalische Zeitschrift*, 20:401, 1919.

[26] Eric D. Black and Ryan N. Gutenkunst. An introduction to signal extraction in interferometric gravitational wave detectors. *Am. J. Phys.*, 71:365, 2003.

[27] Drever et al. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31:97, 1983.

[28] Eric D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *Am. J. Phys.*, 69:79, 2001.

[29] L. Schnupp. talk at European Collaboration Meeting on Interferometric detection of gravitational waves [*sic*]. Sorrent, 1988.

[30] Nicolás de Mateo Smith-Lefebvre. *Techniques for improving the readout sensitivity of gravitational wave antennae*. PhD thesis, Massachusetts Institute of Technology, LIGO-P1200052, 2012.

[31] K. Arai. *Robust extraction of control signals for power-recycled interferometric gravitational-wave detectors*. PhD thesis, University of Tokyo, 2001.

[32] Martin W. Regehr. *Signal extraction and control for an interferometric gravitational wave detector*. PhD thesis, California Institute of Technology, LIGO-P940002, 1995.

[33] M. Evans. Optickle. Technical report, October 2012, LIGO-T070260.

[34] Rochus E. Vogt, Ronald W. P. Drever, Frederick J. Raab, Kip S. Thorne and Rainer Weiss. Proposal to the National Science Foundation: The construction, operation and supporting research and development of a Laser Interferometer Gravitational-wave Observatory (LIGO). Technical report, 1989.

[35] Aidan F. Brooks. *Hartmann wavefront sensors for advanced gravitational wave interferometers*. PhD thesis, The University of Adelaide, 2007.

[36] Stefan Ballmer, Valery Frolov, Ryan Lawrence, William Kells, Gerado Moreno, Ken Mason, David Ottaway, Mike Smith, Cheryl Vorvick, Phil Willems and Mike Zucker. Thermal compensation system description. Technical report, LIGO-T050064, 2005.

[37] John Zweizig. Download site for archived publication-quality interferometer sensitivity curves. Technical report, 2010.

[38] J. Aasi, et al. Advanced LIGO. *Class. Quantum Grav.*, 32(7):074001, 2015.

[39] Robert L. Ward. *Length sensing and control of a prototype advanced interferometric gravitational wave detector*. PhD thesis, California Institute of Technology, LIGO-P1000018, 2010.

[40] S. M. Aston, et al. Update on quadruple suspension design for Advanced LIGO. *Class. Quantum Grav.*, 29:235004, 2012.

[41] Jeffrey S. Kissel. Calibrating and improving the sensitivity of the LIGO detectors. Technical report, Louisiana State University, 2010.

[42] Shyang Wen. *Improved seismic isolation for the laser interferometer gravitational wave observatory with hydraulic external pre-isolator system*. PhD thesis, Louisiana State University, LIGO-P0900021, 2008.

[43] J. A. Sidles and D. Sigg. Optical Torques in Suspended Fabry-Perot Interferometers. *Physics Letters A*, 354(3):167–172, 2006.

[44] Tobin Thomas Fricke. *Homodyne detection for laser-interferometric gravitational wave detectors*. PhD thesis, Louisiana State University, LIGO-P1000010, 2011.

[45] Kenneth A. Strain, Guido Müller, Tom Delker, David H. Reitze, David B. Tanner, James E. Mason, Phil A. Willems, Daniel A. Shaddock, Malcolm B. Gray, Conor Mow-Lowry, and David E. McClelland. Sensing and control in dual-recycling laser interferometer gravitational-wave detectors. *Appl. Opt.*, 42(7):1244–1256, 2003.

[46] Kiwamu Izumi, Koji Arai, Bryan Barr, Joseph Betzwieser, Aidan Brooks, Katrin Dahl, Suresh Doravari, Jennifer C. Driggers, W. Zach Korth, Haixing Miao, Jameson Rollins, Stephen Vass, David Yeaton-Massey, and Rana X. Adhikari. Multi-color cavity metrology. *JOSA A*, 29(10):2092–2103, 2012.

[47] Kiwamu Izumi. *Multi-color interferometry for lock acquisition of laser interferometric gravitational-wave detectors*. PhD thesis, University of Tokyo, LIGO-P1300001, 2012.

[48] J. Hazel, S. Kawamura. Small optics suspension (SOS) final design (mechanical system). Technical report, LIGO-T970135, 1997.

[49] Joseph A. Giaime. *Studies of laser interferometer design and a vibration isolation system for interferometric gravitational wave detectors*. PhD thesis, Massachusetts Institute of Technology, LIGO-P950008, 1995.

[50] B. Willke, N. Uehara, E. K. Gustafson, R. L. Byer, P. J. King, S. U. Seel and R. L. Savage Jr. Spatial and temporal filtering of a 10-W Nd:YAG laser with a Fabry-Perot ring-cavity premode cleaner. *Opt. Lett.*, 23(21):1704, 1998.

[51] PSL PMC cavity, as built at 40m. Technical report.

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

[53] Koji Arai. Mode coupling of two astigmatic Gaussian beams. Technical report, LIGO-T1300364, 2009.

[54] K. Izumi. Mode profile after mode matching telescope (round2). Technical report, 40m elog # 3088, 18 June 2010.

[55] Bram J. J. Slagmolen, Adam J. Mullavey, John Miller, David E. McClelland and Peter Fritschel. Tip-tilt mirror suspension: Beam steering for advanced laser interferometer gravitational wave observatory sensing and control signals. *Rev. Sci. Instrum.*, 82:125108, 2011.

[56] J. Hazel Romie. Small optic suspension (SOS) assembly. Technical report, LIGO-D960001, 1995.

[57] 40m Wiki. Summary of mechanical resonances at the 40 m lab. Technical report.

[58] Janeen Hazel Seiji Kawamura and Jay Heefner. Design report for the 40m test mass suspension prototype. LIGO-T950058, 1995.

[59] E. Quintero. Relative ETM calibration (Rough MC2 calibration). Technical report, 40m elog # 11215, 10 April 2015.

[60] Y. Michimura. Calibration of PRM actuator . Technical report, 40m elog # 8255, 8 March 2013.

[61] K. Arai. SRM calibration. Technical report, 40m elog # 10664, 3 November 2014.

[62] M. Thirugnanasambandam. Calibration of BS, ITMX and ITMY actuators . Technical report, 40m elog # 8242, 6 March 2013.

[63] M. Rakhmanov, D. Reitze, D. Tanner, S. Yoshida and H. Yamamoto. Dynamical properties of LIGO single-loop suspended mirrors. Technical report, LIGO-T000134, 2000.

[64] K. Arai. On the accumulated round-trip Gouy phase shift for a general optical cavity. Technical report, LIGO-T1300189, 2013.

[65] Keiko Kokeyama, Kiwamu Izumi, William Z. Korth, Nicolas Smith-Lefebvre, Koji Arai, and Rana X Adhikari. Residual amplitude modulation in interferometric gravitational wave detectors. *JOSA A*, 31(1):81, 2014.

[66] Jameson Rollins. 40m ALS “beatbox” delay-line frequency discriminator. Technical report, LIGO-D1102241, 2011.

[67] M. Thirugnanasambandam. Beatbox gets a makeover. Technical report, 40m elog # 8818, 10 July 2013.

[68] K. Arai. Beatbox XARM whitening modified. Technical report, 40m elog # 8855, 16 July 2013.

[69] M. Evans and K. Arai. Broadband photodetector (BBPD) design. Technical report, LIGO-T1100467, 2013.

[70] R. Adhikari and J. Driggers. ALS Fool filter updated for more cancellation . Technical report, 40m elog # 11041, 17 February 2015.

[71] E. Quintero. ALS Fool impulse response . Technical report, 40m elog # 11047, 18 February 2015.

[72] M. Flanigan, G. Billingsley. Advanced LIGO beam splitter (BS). Technical report, LIGO-E080514, 2009.

[73] G. Billingsley. Specification - substrate, beam splitter 40M RSE experiment. Technical report, LIGO-E010101, 2001.

[74] 40m Wiki. Measured beam splitter radius of curvature. Technical report.

[75] R. Dannenberg. Advanced LIGO beam splitter coating specification. Technical report, LIGO-E0900073, 2009.

[76] G. Billingsley. Specification: coating, beam splitter, 40m RSE experiment. Technical report, LIGO-E010194, 2001.

[77] M. Flanigan, G. Billingsley. Advanced LIGO input test mass. Technical report, LIGO-E080511, 2009.

[78] R. Adhikari, Y. Aso. Substrate, input test mass, 40m upgrade. Technical report, LIGO-E080454, 2008.

[79] 40m Wiki. Pre-installation evaluation of 40m optics. Technical report, 2010.

[80] R. Dannenberg. Advanced LIGO input test mass coating specification. Technical report, LIGO-E0900041, 2009.

[81] R. Adhikari, Y. Aso, K. Arai. Coating, input test mass, 40m upgrade. Technical report, LIGO-E080457, 2009.

[82] G. Billingsley. Specification: coating, input test mass, 40m RSE experiment. Technical report, LIGO-010195, 2001.

[83] G. Billingsley M. Flanigan. Advanced LIGO end test mass (ETM). Technical report, LIGO-E080512, 2009.

[84] K. Arai A. Brooks. Substrate, end test mass, 40m upgrade. Technical report, LIGO-E0900379, 2009.

[85] R. Dannenberg. Advanced LIGO end test mass coating specification. Technical report, LIGO-E0900068, 2009.

[86] R. Adhikari A. Brooks, K. Arai. Coating, end test mass, 40m upgrade. Technical report, LIGO-E0900378, 2010.

[87] G. Billingsley. Specification: coating, end test mass, 40m RSE experiment. Technical report, LIGO-E010196, 2001.

[88] Rodica Martin. Substrate, aLIGO power recycling mirror (PRM) for L1 and H1 interferometers - unfolded. Technical report, LIGO-E0900087, 2009.

[89] Y. Aso R. Adhikari. Substrate, power recycling mirror, 40m upgrade. Technical report, LIGO-E080452, 2009.

[90] Rodica Martin. Coated substrate, aLIGO power recycling mirror. Technical report, LIGO-E0900245, 2010.

[91] K. Arai R. Adhikari, Y. Aso. Coating, power recycling mirror, 40m upgrade. Technical report, LIGO-E080458, 2010.

[92] G. Billingsley. Specification: coating, power recycling mirror, type B, 40m RSE experiment. Technical report, LIGO-E010198, 2001.

[93] Rodica Martin. Substrate, aLIGO power recycling mirror 2 (PR2) for L1 and H1 interferometers - unfolded. Technical report, LIGO-E0900091, 2009.

[94] Jameson Rollins. Updated calculations of PRC/SRC g-factors and arm mode matching. Technical report, 40m elog # 8316, 20 March 2013.

[95] G. Billingsley M. Flanigan. Advanced LIGO power recycling mirror PR3. Technical report, LIGO-E080516, 2009.

[96] Rodica Martin. Substrate, aLIGO signal recycling mirror (SRM) for L1 and H1 interferometers - unfolded. Technical report, LIGO-E0900089, 2009.

[97] Y. Aso. Substrate, signal recycling mirror, 40m upgrade. Technical report, LIGO-E080453, 2009.

[98] Rodica Martin. Coated substrate, aLIGO signal recycling mirror. Technical report, LIGO-E0900246, 2012.

[99] K. Arai R. Adhikari, Y. Aso. Coating, signal recycling mirror, 40m upgrade. Technical report, LIGO-E080459, 2009.

[100] G. Billingsley. Specification: Coating, signal recycling mirror, 40m RSE experiment. Technical report, LIGO-E010193, 2001.

[101] Rodica Martin. Substrate, aLIGO signal recycling mirror 2 (SR2) for L1 and H1 interferometers - unfolded. Technical report, LIGO-E0900093, 2009.

[102] G. Billingsley M. Flanigan. Advanced LIGO signal recycling mirror SR3. Technical report, LIGO-E080518, 2009.

[103] Jameson Rollins. How to calculate effective RoC of flipped TT. Technical report, 40m elog # 8352, 26 March 2013.

[104] Muzammil A. Arain and Guido Mueller. Optical layout and parameters for the Advanced LIGO cavities. Technical report, LIGO-T0900043, 2010.

[105] K. Arai. Modulation depth. Technical report, 40m elog # 8264, 9 March 2013.

[106] M. Thirugnanasambandam. 40m arm length measurements. Technical report, 40m elog # 9804, 11 April 2014.

[107] E. Quintero. Recycling cavity lengths. Technical report, 40m elog # 10453, 4 Sept 2014.

[108] E. Quintero. Vent update. Technical report, 40m elog # 10564, 26 Sept 2014.

[109] J. Driggers. Remeasuring the 40m Schnupp asymmetry. Technical report, 40m elog # 8528, 3 May 2013.

[110] P. R. Saulson. Terrestrial gravitational noise on a gravitational wave antenna. *Phys. Rev. D*, 30:732, 1984.

[111] S. A. Hughes and K. S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Phys. Rev. D*, 58:122002, 1998.

[112] M. Beccaria, et al. Relevance of Newtonian seismic noise for the VIRGO interferometer sensitivity. *Class. Quantum Grav.*, 15:3339, 1998.

[113] J. A. Giaime, E. J. Daw, M. Weitz, R. Adhikari, P. Fritschel, R. Abbott, R. Bork, and J. Heefner. Feedforward reduction of the microseism disturbance in a long-base-line interferometric gravitational-wave detector. *Rev. Sci. Instrum.*, 74, 2003.

[114] E. J. Daw, J. A. Giaime, D. Lormand, M. Lubinski, and J. Zweizig. Long-term study of the seismic environment at LIGO. *Class. Quantum Grav.*, 21:2255, 2004.

[115] D. J. Ottaway, P. Fritschel, and S. J. Waldman. Impact of upconverted scattered light on advanced interferometric gravitational wave detectors. *Optics Express*, 20:8329–8336, 2012.

[116] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series*. M.I.T. Press, 1964.

[117] J. Durbin. The fitting of time-series models. *Rev. Inst. Int. stat.*, 28:233, 1960.

[118] Y. Huang, J. Benesty, and J. Chen. *Acoustic MIMO signal processing*. Signals and Communication Technology. Springer, 2006.

[119] Güralp Systems Ltd. CMG-40T triaxial broadband seismometer, 2006.

[120] Wilcoxon Research. 731A ultra-quiet, ultra low frequency seismic accelerometer.

[121] J. Giaime, P. Saha, D. Shoemaker, and L. Sievers. A passive vibration isolation stack for LIGO: Design, modeling, and testing. *Rev. Sci. Instrum.*, 67:208, 1996.

[122] MathWorks. MATLAB 2010.

[123] N. A. Robertson, et al. Seismic isolation and suspension systems for Advanced LIGO. *Proc. SPIE, Gravitational Wave and Particle Astrophysics Detectors*, 5500:81, 2004.

[124] Richard A. Haubrich and Keith McCamy. Microseisms: Coastal and pelagic sources. *Rev. Geophysics*, 7(3):539–571, 1969.

[125] B. Gustavsen and A. Semlyen. Rational approximation of frequency domain responses by vector fitting. *IEEE Transactions on Power Delivery*, 14:1052–1061, 1999.

[126] R. Abbott *et al.* Seismic isolation enhancements for Initial and Advanced LIGO. *Classical and Quantum Gravity*, 21:915, 2004.

[127] C. Hardham, et al. Multi-DOF isolation and alignment with quiet hydraulic actuators. *Proceedings of ASPE Spring 2004 Topical Meeting on Control of Precision Systems*, 2004.

[128] A. M. Gretarsson, E. D'Ambrosio, V. Frolov, B. O'Reilly, and P. K. Fritschel. Effects of mode degeneracy in the LIGO Livingston Observatory recycling cavity. *J. Opt. Soc. Am. B*, 24, 2007.

[129] Trillium 240 broadband seismometer.

[130] J. Driggers. First try with PRCL ASC Wiener filtering. Technical report, 40m elog # 10959, 30 January 2015.

[131] J. Driggers. Preliminary PRMI angular Wiener results. Technical report, 40m elog # 11186, 31 March 2015.

[132] Bernard Widrow and Marcian E. Hoff. Adaptive switching circuits. *IRE WESCON Conv. Rec., Part 4*, pages 96–104, 1960.

[133] Ali H. Sayed. *Fundamentals of adaptive filtering*. Wiley-IEEE Press, 2003.

[134] Simon Haykin. *Adaptive filter theory*. Prentice Hall, 4th edition, 2002.

[135] Max Kamenetsky and Bernard Widrow. A variable leaky LMS adaptive algorithm. *Conference Record of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers*, 1:125, 2004.

[136] Bernard Widrow and Peter N. Stearns. *Adaptive signal processing*. Prentice Hall, 1985.

[137] National Instruments. Filtered-x LMS algorithms.

[138] D. Martynov. Calibrated spectra from OAF test. Technical report, 40m elog # 9606, 5 February 2014.

[139] R. Weiss. Electromagnetically coupled broadband gravitational antenna. Technical report, LIGO-P720002, 1972.

[140] P. R. Saulson. Terrestrial gravitational noise on a gravitational wave antenna. *Phys. Rev. D*, 30:732, 1984.

[141] Scott A. Hughes and Kip S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Phys. Rev. D*, 58:122002, 1998.

[142] Giancarlo Cella. Underground Reduction of Gravity Gradient Noise. Technical report, LIGO-G060311, 2006.

[143] T. Accadia *et al.* (VIRGO Collaboration). The seismic superattenuators of the Virgo gravitational waves interferometer. *Low Frequency Noise, Vibration and Active Control*, 30(1):63–79, 2011.

[144] The LIGO Scientific Collaboration. Instrument Science white paper. Technical report, LIGO-T1200199, 2012.

[145] S. Hild *et al.* Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28:094013, 2011.

[146] Keenan Pepper. Newtonian noise simulation and suppression for gravitational-wave interferometers. Technical report, LIGO-T070192, 2007.

[147] Jan Harms, Riccardo DeSalvo, Steven Dorsher, and Vuk Mandic. Simulation of underground gravity gradients from stochastic seismic fields. *Phys. Rev. D*, 80:122001, 2009.

[148] Wilcoxon Research. 731-207 low frequency seismic accelerometer.

[149] T. Creighton. Tumbleweeds and airborne gravitational noise sources for LIGO. *Class. Quantum Grav.*, 25:125011, 2008.

[150] Streckeisen STS-2 Broadband Sensor.

[151] Jan Harms and Brian O'Reilly. Velocity and attenuation characterization of the LIGO site near Livingston, Louisiana. *BSSA*, 101:1478, 2011.

[152] M. N. Toksöz and R. T. Lacoss. Microseisms: Mode Structure and Sources. *Science*, 159:872, 1968.

[153] M. G. Beker, G. Cella, R. DeSalvo, M. Doets, H. Grote, J. Harms, E. Hennes, V. Mandic, D. S. Rabeling, J. F. J. van den Brand and C. M. van Leeuwen. Improving the sensitivity of future GW observatories in the 1–10 Hz band: Newtonian and seismic noise. *Gen. Relativ. Gravit.*, 43(2):623–656, 2011.

[154] G. A. Prieto, J. F. Lawrence, and G. C. Beroza. Anelastic Earth structure from the coherency of the ambient seismic field. *J. Geophys. Res.*, 114:B07303, 2009.

[155] W. J. Stephenson, S. Hartzell, A. D. Frankel, M. Asten, D. L. Carver, and W. Y. Kim. Site characterization for urban seismic hazards in lower Manhattan, New York City, from microtremor array analysis. *Geophys. Res. Lett.*, 36:L03301, 2009.

[156] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of IEEE International Conference on Neural Networks*, volume 4, pages 1942–1948, Perth, WA, Australia, 1995.

[157] S. Chen, MATLAB File Exchange. Another particle swarm toolbox. 2009.

[158] M. B. Priestley. *Spectral analysis and time series*. Academic Press, 2001.

[159] Freeman Dyson. Seismic response of the Earth to a gravitational wave in the 1-Hz band. *Astrophysics Journal*, 156:529, 1969.

[160] Michael Coughlin and Jan Harms. Upper limit on a stochastic background of gravitational waves from seismic measurements in the range 0.05–1 Hz. *Phys. Rev. Lett.*, 112:101102, 2014.

[161] Michael Coughlin and Jan Harms. Constraining the gravitational wave energy density of the Universe using Earth's ring. *Phys. Rev. D*, 90:042005, 2014.

[162] J. Capon. High-resolution frequency-wavenumber spectrum analysis. *Proc. IEEE*, 57:1408, 1969.

[163] United States Geological Survey (USGS) glossary. Diagram of body and surface waves.

[164] Jan Harms and Stefan Hild. Passive Newtonian noise suppression for gravitational-wave observatories based on shaping of the local topography. *Class. Quantum Grav.*, 31(18):185011, 2014.

[165] M. Coughlin and J. Harms. Seismic topographic scattering in the context of GW detector site selection. *Class. Quantum Grav.*, 29(7):075004, 2012.

[166] John Bechhoefer. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.*, 77:783, 2005.