LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T1200462–x0-v1** 2012/09/28

Real-time Simulation of a Suspended Cavity with the Advanced LIGO Digital Controls System - Final Report

Alexandra Zhdanova

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

1 Introduction

1.1 This Report

My project consisted of two parts: simulating a Fabry Perot cavity with Matlab, and modeling the same Fabry-Perot cavity in real-time. One cavity consists of two freely hanging, opposing mirrors called the "input test mass" (the mirror closest to the laser) or ITM and the "end test mass" (the mirror farthest from the laser) or ETM. The mirrors at each end are single-suspension mirrors, consisting of a reasonably sized mirror supported by a metal fiber in a vacuum. The cavity is meant to be part of a larger system of simulations that model the 40m interferometer in real time; part of my simulation was setting up the infrastructure for these. The system would eventually serve as a way to compare the reliability of the data and our understanding of the noise sources. Theoretically, if the noise between the real and fake interferometers were exactly the same, the data they produce would also be the same (within the same time frame).

1.2 LIGO and Gravitational Waves

Looking at the larger picture, Albert Einstein's theory of General Relativity predicts that especially massive, quickly-moving objects send out waves in the fabric of space-time. These waves are called gravitational waves, and it is the goal of the Laser Interferometer Gravitational Wave Observatory (or LIGO) to detect them. To understand how it does (or plans) to do so, some discussion on the effects of gravitational waves is necessary. When a gravitational wave passes through, say, a yard stick, it stretches that yard stick by a very, very small amount, proportional to the length of the yard stick. With two identical yard sticks perpendicular to each other, the gravitational wave will simultaneously stretch one and shorten the other. If the difference in the lengths of the yard sticks were measured (at the exact moment the gravitational wave passed through), one would be longer than the other and we would know that we had seen a gravitational wave [1].

With interferometers, one can recreate that effect on a significantly larger scale using laser beams in Fabry-Perot cavities. The LIGO interferometers located at Hanford, WA and Livingston, LA essentially consist of two perpendicular 4km long arms, with mirrors at each end and a beam splitter and detector in the middle (at the cross of the arms). We measure the difference in length of the arms by observing the interference between the two beams of light returning from the arms. If one beam of light (coming from one of the arms) takes a longer time to return than the other, then they will not perfectly cancel each other out and light is sent to the detector. This is done (on a smaller scale) at Caltech as a 40m interferometer consisting of optical Fabry-Perot cavities, so that the science behind data analysis and interferometery can be studied.

2 The Interferometer

While entire theses have been written on the subject of the interferometers (in particular, the ones at Livingston and Hanford), of most concern to us are three parts: one Fabry-Perot

cavity, the suspended mirrors that make up the cavity, and the Pound-Drever-Hall technique used to control the laser that feeds light into the cavity. The behavior of each of these is essential in fully modelling a real-life Fabry-Perot cavity at the 40m interferometer. I will begin with a discussion of the cavity, as it is possible to build a basic simulation using only its response, then move on to the suspensions and, lastly, the Pound-Drever-Hall technique.

2.1 The Cavity

Fabry-Perot cavities are relvant to LIGO because light can build or constructively interfere with itself inside such a cavity (i.e. between the two mirrors). This condition, called "resonance", traps light and forces it to bounce back and forth between the mirrors multiple times. This, in turn, allows our 4km arms to seem like 40 km arms (recall that the strength of a gravitational wave signal is proportional to the length of each arm) while simultaneously improving the power and strength of any signal since each signal build on the signal from the previous bounce, which then builds on the bounce before that, etc..

Resonance occurs when an integer number of wavelengths of light fit into the cavity. Mathematically, the cavity is in resonance when

$$L = n\lambda_0/2, \tag{1}$$

where L is the length of the cavity, n is an integer, and λ_0 is the wavelength of the light coming in. Changing L or the frequency of the light coming in will bring the cavity in and out of resonance. For now, my simulation only takes into account changes in length, not changes in frequency (though this will be added - see the section 4.1 for further details). In theory, we could just define our simulation from here, knowing (for a given λ_0) the resonance condition for changes in length. However, a simpler way to simulate the cavity response is in the frequency domain, using the cavity's linewidth. The linewidth is defined as the frequency between resonance peaks. For example, if our cavity resonates at 20 Hz, then drops out of resonance, then is back again at 40 Hz, our linewidth is $40 - 20 = 20$ Hz.

A general formula for the linewidth can be derived in multiple ways; in [6] it is done by transforming the cavity fields into the frequency domain and solving for the linewidth from there. It is also done in [7] by fitting a Lorentzian peak to the power in the cavity and solving numerically. Either way, the expression derived for the linewidth is fairly simple:

$$\gamma = cT_1/(4L) \tag{2}$$

Where c is the speed of light in a vacuum, T_1 is the transmission constant of the input mirror, and γ is the cavity's pole, or linewidth in Hz. If the frequency of the incoming laser light is an integer multiple of γ, our cavity resonanates. This relatively simple mathematical behavior is modeled as a filter with a single pole at $\omega = \gamma$, which is also why the cavity linewidth is occasionally referred to as the cavity pole[1].

[1]Note that ω is sometimes expressed as the angular frequency, in which case γ must be multiplied by 2π. In general, be very careful when converting from Hz to radians.

The optical spring effect can also affect resonance conditions. The two suspended mirrors that form the cavity act like a single coupled harmonic oscillator, connected by the light between them. When the mirrors are pushed outward, away from resonance, the power in the cavity decreases and so the force (caused by radiation pressure from the light building up in the cavity) pushing out is diminished. This brings the mirrors back into resonance, which then increases the power and the force pushing outward dominates the interaction again. The formula describing this effect (sometimes referred to as optical rigidity) is given in a variety of places. The derivation is given in [7], while an easy-to-work-with form is given in [8] as:

$$K(\omega) = K_0 \frac{1 + (\delta/\gamma)^2}{(1 + i\omega\gamma)^2 + (\delta/\gamma)^2} \tag{3}$$

where

$$K_0 = 64 w_0 I_0 \frac{c^2 T_1^2}{(\delta/\gamma)^2/(1 + (\delta/\gamma)^2}. \tag{4}$$

Where $K(\omega)$ is the transfer function between the frequency of the light on the test mass (or masses) and their resultant displacement due to radiation pressure. Defining the various constants: I_0 is the power of the laser incident on the cavity (5W in my simulation), and δ is the detuning parameter - it describes how far off resonance the laser light currently is. For example, if our cavity resonantes at $30Hz$, and our laser light is incoming at $40Hz$, $\delta = 40 - 30 = 10Hz$. δ/γ then is called the normalized detuning parameter and is unitless.

The most important feature of these equations is that past the cavity pole (γ), $K(\omega)$ dies off fairly quickly, which makes sense given that the power in the cavity dies fairly quickly past resonance. A more interesting aspect of these equations is the detuning parameter, δ/γ , since it provides a convenient way to express the transfer functions in the cavity. A positive δ/γ corresponds to a restoring optical spring effect, while a negative δ/γ corresponds to an anti-restoring optical spring (otherwise known as an anti-spring). Mathematically, δ/γ is given as:

$$\delta/\gamma = (\omega_0 - \pi c n/L)/\gamma. \tag{5}$$

Where n is an integer from 1, 2, 3... For those not convinced that, together, the cavity resonance condition and the optical spring effect completely describe the response of the cavity, see [9] for a similar, but slightly more technical breakdown of the cavity response.

We described the definitions of γ and $K(\omega)$ in terms of changes in the frequency of the laser light with ideally stationary mirrors. In reality, it is usaully the other way around - changes in resonance are brought about by changes in the position of the mirrors, not changes in the frequency of the laser light. However, these changes can be described identically, with the frequency of the motion of the mirrors adding to the frequency of the laser light. Indeed, if we assume that the cavity is by default in resonance and that the frequency of the laser is stable, then the motion of the mirrors is the most significant source of change in the cavity's state. This leads us into the next section - describing the motion of the mirrors.

2.2 The Suspensions

The LIGO mirror, or test mass, is a small optic (or a large optic, in the case of the larger interferometers) hanging from piano wire attached to the ceiling of the arm. We can describe its motion using three equations for displacement (x), pitch (θ), and yaw (ϕ) [12][2]:

$$\ddot{x} + \gamma_x \dot{x} + \omega_x^2 x = \omega_x^2 (x_{sus} + b\theta) \tag{6}$$

$$\ddot{\theta} + \gamma_\theta \dot{\theta} + \omega_\theta^2 \theta = \frac{\omega_\theta^2 (x - x_{sus})}{l + b} \tag{7}$$

$$\ddot{\phi} + \gamma_\phi \dot{\phi} + \omega_\phi^2 \phi = \omega_\phi^2 \phi_{sus} \tag{8}$$

Where γ_x, γ_θ, and γ_ϕ are damping coefficients each equal to $\frac{1}{2Q}$ where Q is the quality factor of their respective modes. x_{sus} and ϕ_{sus} are the position and yaw of the suspension point, l is the distance from the suspension point to the center of mass, and b is the distance from the point where the wires are attached to the center line of the mirror (see Figure 1).

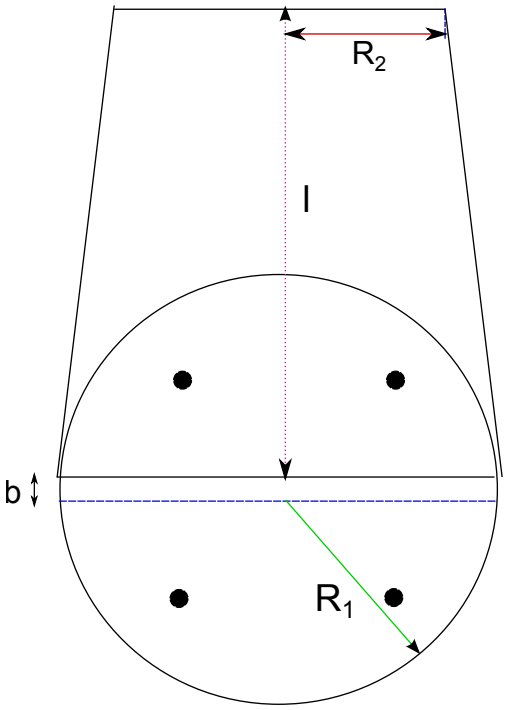

Figure 1: The LIGO mirror as a rigid body. The center line is red, while the line of attachment is black. Note the four OSEMs - Upper Left (UL), Upper Right (UR), Lower Left (LL), Lower Right (LR).

[2]Note that a treatment of the side degree of freedom is missing from our treatment of the general degrees of freedom. This is because the equations that describe motion of side are exactly the same as those that describe motion of position.

If we transform $\dot{x}$ (our time domain variable) into s (our Laplace variable), we can express the coupled pitch/position equations as:

$$Q(s)\begin{bmatrix} x(s) \\ \theta(s) \end{bmatrix} = \begin{bmatrix} \omega_x^2 \\ \frac{-\omega_\theta^2}{l+b} \end{bmatrix} x_{sus}(s) \tag{9}$$

where $Q(s)$ is just the left-hand side of our previous mechanics equations, transformed into the frequency domain and with $x_{sus}(s)$ isolated on the right-hand side. To put it more plainly;

$$Q(s) = \begin{bmatrix} s^2 + \gamma_x s + \omega_x^2 & -b\omega_x^2 \\ \frac{-\omega_\theta^2}{l+b} & s^2 + \gamma_\theta s + \omega_\theta^2 \end{bmatrix} \tag{10}$$

.

We can express this equation in a more convenient form as:

$$\begin{bmatrix} x(s) \\ \theta(s) \end{bmatrix} = Q^{-1}(s) \begin{bmatrix} \omega_x^2 \\ \frac{-\omega_\theta^2}{l+b} \end{bmatrix} x_{sus}(s) \tag{11}$$

where $Q^{-1}(s)$ is the inverse of $Q(s)$. This gives us frequency-domain equations for position and pitch, which we can then use.

This was a major part of my simulation of the suspensions, in addition to conversion between the degree of freedom basis (position, pitch, yaw, side) to the OSEM basis (UL, UR, LL, LR). One OSEM consists of a coil actuator and a shadow sensor. The coil actuator translates the electronic signals it receives into an acutation force, while the shadow sensor uses the shadow of the suspension to detect its displacement and then output that as a voltage. There is an OSEM on each quadrant of the mirror and one on the side for a total of five OSEMs. The digital control system interprets data in terms of the OSEM basis, while our mechanics are done in the DOF basis, so it is necessary to convert from one to the other. [13] has a clear-cut treatment of how to do so; suffice to say that we employ an extra DOF called pringle (or butterfly or bend) to form a complete DOF basis. We then use the resulting matrix to convert from DOF to OSEM, then invert it to get the matrix that converts from OSEM to DOF.

This matrix, in reality, is not constant in time. Since there are differences between the two test masses (small ones, but still there), we have to generate the matrices ourselves based on data from the suspensions. For more information on this topic, see Section 4.2.

2.3 Pound-Drever-Hall

The Pound-Drever-Hall (or PDH) technique is the way we detect changes in the state of the cavity - either through changes in position of the mirrors or changes in the laser's frequency. This section will focus specifically on how we detect changes in the position of the mirrors, but (like our discussion of the cavity) the treatment of changes in the laser's frequency is identical.

To begin with, the electric field outputted from the laser is a field which is ideally resonant in the cavity and has a wavelength λ of 1064 nanometers. Onto that field (henceforth referred to as the carrier field), we add sidebands. The sidebands are less powerful electric fields with frequencies slightly higher ("upper sideband") and slightly lower ("lower sideband") than the main carrier. How much higher and how much lower is defined by a process called "modulation", typically done through a Pockels Cell, a discussion of which is outside the scope of this paper. Usually, the sidebands are modulated to be within radio-frequency of the carrier, around 12.5 MHz higher and lower. The modulation frequency will be referred to as Ω, while the carrier frequency will be ω to make the math simpler. The sidebands and carrier are collectively referred to as E_{inc}, the incoming field on the input test mass. P_c and P_s are the power in the cavity and the power in the sidebands, respectively, defined as follows:

$$P_c = J_0^2(\beta)P_0, \tag{12}$$

$$P_s = J_1^2(\beta)P_0, \tag{13}$$

where P_0 is the total power in the beam (typically around 5W, though aLIGO will use lasers on the order of 200W) and β is the modulation depth (we can set whatever β we like in the modulation process). When β is small, almost all of the power is in the carrier and first-order sidebands [11]:

$$P_0 \approx P_c + 2P_s. \tag{14}$$

Basic optics suggests that the field reflected (E_{refl})from the input miror will have the same frequency as the incoming field, and its amplitude will be proportional to the amplitude of the incoming field. We will label this constant of proporionality as the reflection coefficient $F(\omega)$:

$$F(\omega) = E_{refl}/E_{inc} = \frac{-r_1 + r_2(r_1^2 + t_1^2)e^{i\omega/\Delta\nu_{fsr}}}{1 - r_1 r_2} \tag{15}$$

where $\Delta\nu_{fsr}$ is the cavity free-spectral-range $\pi c/L/2\pi = cL/2$, r_1 and t_1 are the square roots of the reflection constant and transmission constant (typically referred to as R_1 and T_1 respectively) of the input test mass, and r_2 and t_2 are the equivalents of r_1 and t_1 for the end test mass, and i is the square root of -1.

We can now solve forE_{refl} in terms of $F(\omega)$ and E_{inc}, which will then give us P_{refl} (and hence the signal we receive at the photodetector).

$$E_{refl} = E_0[F(\omega)J_0(\beta)e^{i\omega t} + F(\omega+\Omega)J_1(\beta)e^{i(\omega+\Omega)t}]F(\omega-\Omega)J_1(\beta)e^{i(\omega-\Omega)t} \tag{16}$$

$$\begin{aligned} P_{refl} = |E_{refl}|^2 &= P_c|F(\omega)|^2 + P_s|F(\omega+\Omega)|^2 + |F(\omega-\Omega)|^2 \\ &+2\sqrt{P_c P_s}(\mathrm{Re}[F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)]\cos(\Omega t) \\ &+\mathrm{Im}[F(\omega)F^*(\omega+\Omega) - F^*(w)F(\omega-\Omega)]\sin(\Omega t)) \\ &+(2\Omega\,\mathrm{terms}) \end{aligned}$$

The equation for P_{refl} represents the complete DC signal from the photodiode. The term proportional to $\cos(\Omega t)$ will become the quadrature signal, while the $\sin(\Omega t)$ term will become the in-phase signal. Either way, we want to measure $F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)$. The 2Ω terms can usually be discarded. When the sidebands are exactly anti-resonant, $F(\omega \pm \Omega) = -1$, and the term we're interested in reduces to:

$$F(\omega)F^*(\omega+\Omega)F^*(\omega)F(\omega+\Omega) = -2i\,\mathrm{Im}(F(\omega)) \tag{17}$$

A purely imaginary signal! This leaves our in-phase signal as 0, and so we focus on the quadrature. The signal from the photodiode is then mixed (i.e. multiplied) with the original modulation signal $\sin(\Omega t)$. Our basic trigonometry tells us:

$$\sin(\omega t)\sin(\omega' t) = 1/2[\cos((\omega-\omega')t)\cos((\omega+\omega')t)] \tag{18}$$

In the case where $\omega = \omega'$, that leaves us with a non-oscillatory term and an oscillatory term! What really concerns us is how all of this relates to changes in the position of the test masses. We get there by defining a variable for the phase light picks up in a single trip. This variable is usually called ϕ, but to avoid confuion with yaw, let's call it ζ. ζ is usually related to the cavity length by:

$$\zeta = 2\pi\frac{2L}{\lambda} \tag{19}$$

If we plug this into our original expression for $F(\omega)$ (Equation 15) as ω, and then solve for P_{refl}, we will have an expression for the cavity length in terms of the power of the light at the photodiode. We start by writing ζ in a more convenient form; near resonance, we can write ζ as:

$$\zeta \approx 2\pi n + 4\pi dL/\lambda_0. \tag{20}$$

Where dL is the displacement from resonance. Now we plug this expression for ζ into our expression for $F(\omega)$ (Equation 15) and use the small-angles approximation to get:

$$F(dL) \approx (-r_1 + r_2(r_1^2 + t_1^2)) * (i4\pi dL/\lambda_0) \tag{21}$$

Finally, plugging in to P_{refl}...

$$P_{refl} \approx P_c|F(dL)|^2 + 2P_s - 4\sqrt{P_c P_s}\,\mathrm{Im}[F(dL)]\sin(\Omega t) \tag{22}$$

After mixing, the quadrature signal is low-passed (to fully isolate the non-oscillatory terms) and then converted into a form the computer can read (see Section 3.2 for more on this).

An understanding of the Pound-Drever-Hall Technique (abbreviated as PDH) is important in any serious simulation of the optical cavity, since this is where a good deal of the sensing noise manifests itself (noise in the suspensions, otherwise known as "displacement noise, is discussed further in the paper). The only information we technically have (at least in the case of a single optical cavity) is the light on the photodetector at the signal port, and from that we have to deduce everything we possibly can. [10] and [11] are basic, well-rounded introductions to PDH. Our simulation took the displacement as an input, but it is easy to conceive of a simulation that instead took the light on the detector.

3 The Initial Simulation

At the beginning of the project, our knowledge of Fabry-Perot cavities, suspensions, and the interferometer were very limited. In order to learn (and to develop an initial prototype), we started by developing a simulation of the Fabry-Perot cavity in Simulink, Matlab's simulaiton environment. Noises were not the major concern here, rather, we focused on the mechanics and general description of the cavity and control system, a discussion of which we shall begin with.

3.1 The Control System

The control system is what keeps the cavity in resonance; it receives signals from the OSEMs and the photodiodes and accordingly sends out voltages to the actuators in the OSEMs to move the mirrors back into place. The control system operates on two levels. The first level is simply responsible for damping out resonant motion in the mirror, so it is usually referred to as the local damping system. The second level is what actively controls the length of the cavity and the frequency of the laser light. It essentially feeds back the motion of either instrument; for example, if the mirror moves two inches to the left, the control system will send out a voltage that causes the acutators on the mirror to output a force that moves it two inches to the right.

Transitioning data to the control system requries an analog-to-digital conversion (ADC) from volts (analog units) to counts (digital units), while transitioning data from the control system requires a digital-to-analog conversion (DAC) from counts to volts.

3.2 DAC and ADC

There is a simple algorithm for converting a digital signal to an analog force (F) output in Newtons:

$$F = \frac{20\text{ V}}{32\text{ kcounts}} * \frac{1\text{ Amp}}{400\text{ Volts}} * 0.016\frac{\text{N}}{\text{Amp}} \tag{23}$$

To convert from an analog force input to a digital output, we just implement the algorithm in reverse. However, there is a limit on how large a signal can be in order to be digitized, because the largest number of counts that we can output with our current method is 2^{16}. The physical ADC is essentially made of 16 voltage dividers, and because of this, the signal must first be compacted before it is sent through the ADC. This is achieved through a whitening filter that flattens a signal so that it has roughly the same strength at all frequencies.

3.3 Electronic Filters

A major part of the simulation was finding and implementing the various electronic filters - whitening, dewhitening, anti-aliasing, anti-imaging, and general low-passing. Whitening, anti-aliasing, and dewhitening filters will be discussed in the next section. Anti-imaging filters are used after Digital to Analog Conversion (DAC) to construct a continuous analog signal from a discrete digital input. Anti-aliasing filters are used after whitening filters and before Analog to Digital (ADC) filters to remove data above the Nyquist frequency. The Nyquist frequency is half the sampling frequency, and is the maximum frequency of an incoming signal. Any signal going above the Nyquist frequency is moving too fast for us to adequately sample two points (hence the factor of 1/2). For a complete simulation, it was also necessary to include DAC and ADC, which will be discussed in the next section.

3.4 Minutia of Simulink

For those already familiar with Simulink, feel free to skip this part and go directly to the description of the Real-time Model (section 4). Otherwise, we will discuss how Simulink actually works.

Simulink and Matlab have a programming environment extremely well suited for dynamic systems. Transfer functions are easy to deduce and implement, and Matlab takes care of most of the math. The initial simulation went through five iterations, each handling the cavity in a completely different but hopefully more accurate way. The fifth iteration handles the suspensions and local damping of the motion as library parts (i.e. the simulations for each suspension are identical). It successfully handled four degrees of freedom - pitch, position, yaw, and side (as well as pitch/x coupling). It also included a Simulink version of the local damping system.

In Simulink, each block represents a different function, with arrows representing the path of the input. So "mod" blocks perform the "modulus" function on the input, and sum blocks sum their inputs together. With that in mind, we can present the ADC function in Simulink as in Figure 2.

We can also display a Simulink diagram of the full cavity simulation (Figure 3).

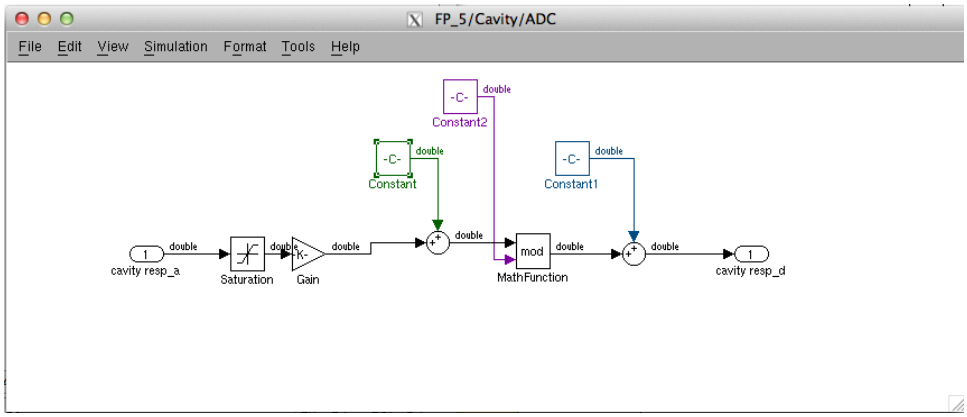

Figure 2: A typical ADC. The constants are defined as follows: Green C = 32768.5, Purple C = 100000, Blue C = -32768.

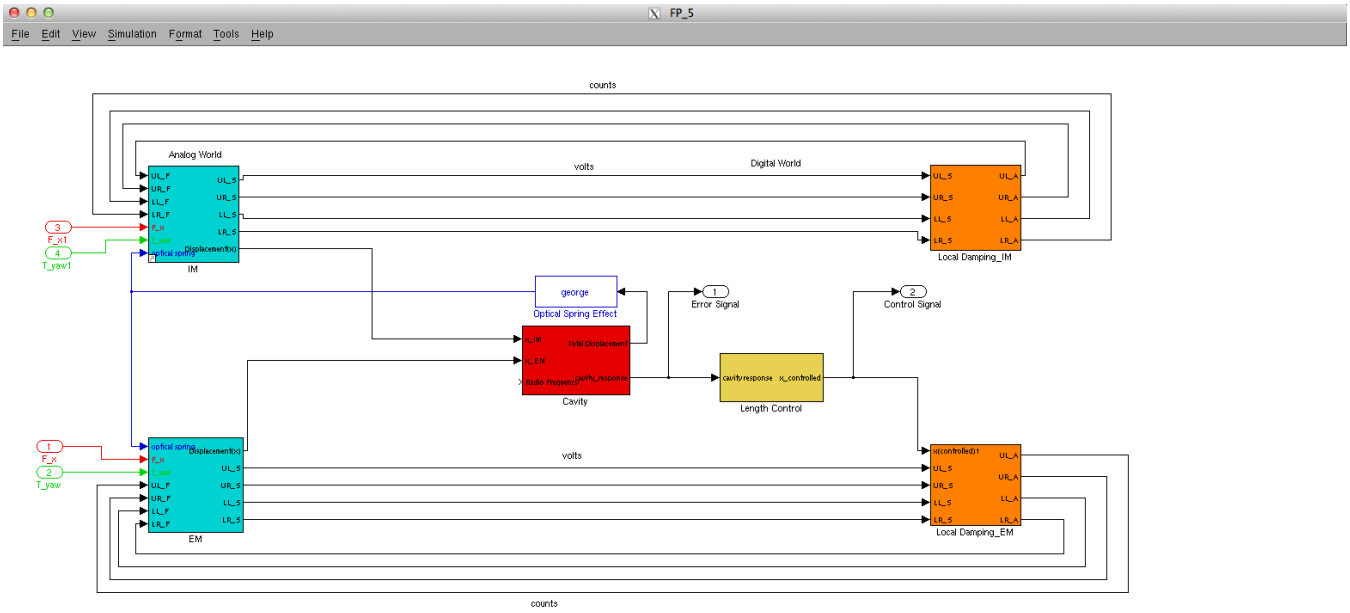

Figure 3: The final version of the full Simulink simulation. IM and EM are the Input Test Mass and the End Test Mass, respectively. In those blocks, the mechanics of the suspensions (discussed in section 2.) are fully implemented. The cavity block represents the full response of the cavity (discussed in section 2.), including PDH. The local damping and length control blocks take input in the form of counts.

4 The Real-time Model

A working real-time model was the end goal of the project; the Simulink simulation was preparation for it. The real-time model gave us our first opportunity to test our knowledge of the noise sources, and we spent a significant amount of time studying them. However, our discussion of the real-time model will be split between the things inherent to the complete real-time model (this section) and the noises (next section).

4.1 Differences Between The Initial Simulation and the Real-time Model

The most significant difference is that Simulink is an environment where any math needed can be done within the simulation itself. For example, Simulink comes with Bessel functions and integration, so there was no trouble using those in our original simulation. The real-time model, however, does not. As a result, we represented the mechanics as a frequency domain filter (as done in Section 2.2) and all the noises as filtered white noise (see Section 5).

Another salient difference between the Simulink and real-time models is that the Simulink model was offline; it depended on nothing except what was contained in the Simulink simulation. The real-time models, on the other hand, took data from the actual, real-time control system as an input. As a result, the real-time control system sent its output to the models instead of to the actual suspension, and the real suspension was left to swing free. This meant that testing of the real-time model had to be confined to times when other people did not need to use the interferometer. In other words, there was much less opportunity to check our work.

Besides synchronization with the control system, the real-time system involved several different models working together. Essentially, every block in Figure 3 was a seperate model, and each model had to sync up with the others. So, our model of the cavity took its input from the suspension model, and then it sent its output to the real-time control system, which then sent its input back to the suspension model. See Figure 4 for a visual realization of this idea.

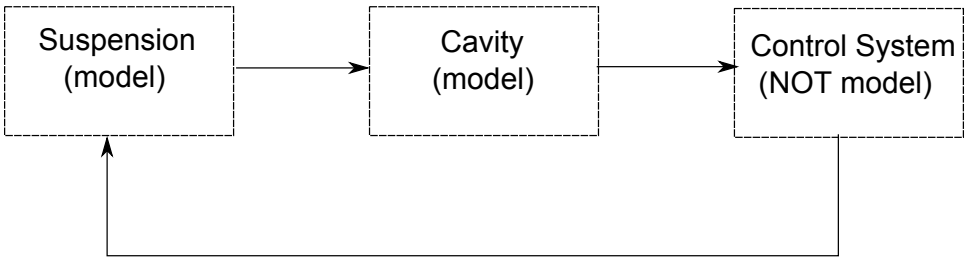

Figure 4: A simple block diagram of the closed-loop system, with blocks representation the cavity and suspensions.

4.2 Conversions

The OSEM to degree of freedom (i.e. pitch, yaw, side, position) conversion matrix given in [13] is not valid for our real-time model. Indeed, it is not even close to a true conversion matrix, due to strange couplings (like yaw and side), wire inconsistencies, and optical deformations. As such, it is necessary to constantly measure how displacements in degrees of freedom (i.e. a direct push on the mirror) result in signals from the OSEMs and then derive the matrix numerically. The procedure for doing so is outlined in [14], though what really concerns us is the end result. We used the matrix the control system used for conversion between OSEMs and DOF, though it is important to keep in mind that this matrix may not be correct in the future when another measurement is made (see Section 7).

In addition, the ADC/DAC conversion algorithm outlined in section 4. had to be condensced into one numerical factor, which we then incorporated as a filter gain. This brings us to our next topic of discussion, the noises.

5 The Noises

The noises can be divided into three broad categories: displacement noises (noises that actually move the mirrors), sensing noises (noises that prevent us from precisely measuring the movement of the mirrors), and electronic noises (noises that prevent the data the control system sees from being exactly the same as the data the sensors send out). We included one significant noise source from each of the three categories, and set up the basis for future noise sources.

5.1 Displacement Noise

Seismic noise is the most significant noise source for our model. To model it, we started by obtaining a power spectrum of the seismic noise from a seismometer. In general, a power spectrum of seismic noise wiill be rather high at low frequencies and exponentially lower past 1 Hz until it hits a floor at $3 * 10^2$ Hz. Seismometers are less sensitive and much noisier at low frequencies, and seismic noise isn't nearly as high as seismometer noise is below 1 Hz. Therefore, in reality, seismic noise is much lower below 1 Hz than can be implied from a basic power spectrum. Significant peaks include a peak at 0.33 Hz from water hitting shores all over the world (water waves have a period of three seconds) and a peak at 1 Hz from people walking.

After obtaining a power spectrum, we used a previously written program to create a rough filter directly from the spectrum. We then refined this filter, to make the fake seismic noise look more like real seismic noise, and obtained an acceptable seismic noise spectrum from white noise (Figure 5).

5.2 Sensing Noise

Shot noise is caused by fluctuations in the vacuum field beating against the radio frequency sidebands. It is proportional to the light inside the cavity, as the larger light level enhances

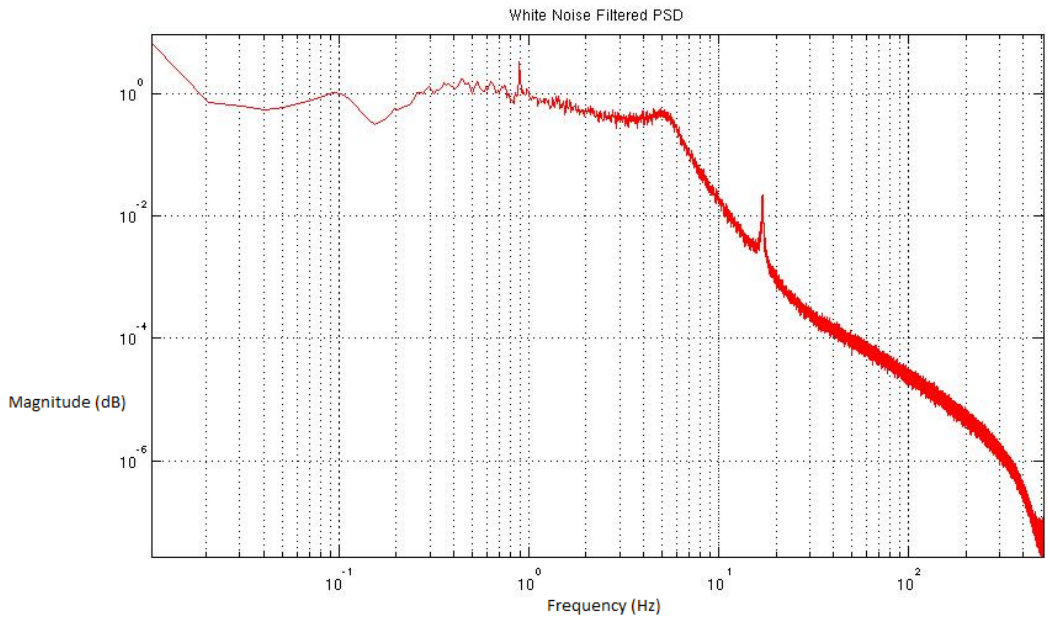

Figure 5: A seismic noise power spectrum obtained from filtered white noise.

fluctuations in the vacuum. However, a more fitting metaphor is rain beating on a tin roof. Each rain drop causes a noise, and the more rain drops there are, the greater the noise. Hence, for the purposes of the real-time model, we estimated shot noise as being white noise noise roughly proportional to the power in the cavity (obtained mathematically in section 2.3). A more thorough, mathematical treatment of the shot noise for the differential light signal is given in [1].

Other major sensing noise sources include laser amplitude/frequency noise and local oscillator phase/amplitude noise. The first noise we completely neglected, as a good deal of the math done in section 2.3 would have to be edited to include fluctuations in laser frequency. The second is part of the more complicated process of modulation, and so was outside the scope of the project. In addition, these noises are relatively minor in comparison to the electronic noise and seismic noise. However, they do provide an area for further development (see Section 7). For a good overview of these noises, see [9].

5.3 Electronic Noise

A major source of ADC noise is the so called "60 Hz" noise, where 60 Hz is the standard frequency at which power is transmitted through power lines. The 60 Hz noise is essentially a series of large spikes of noise at 60 Hz and its harmonics (30, 120, 180...) mixed with a low white noise level. We developed an electronic noise filter from scrath, though the extremely sharp peaks at 60 Hz (and 30, 120, 180 Hz...) that are characteristic of 60 Hz noise were difficult to reproduce in a filter. The result can be seen in Figure 6.

There are, obviously, other electronic noises. A whole summer project ([15]) in Livingston was done to detect them. However, for the purposes of this project, ADC noise was sufficient

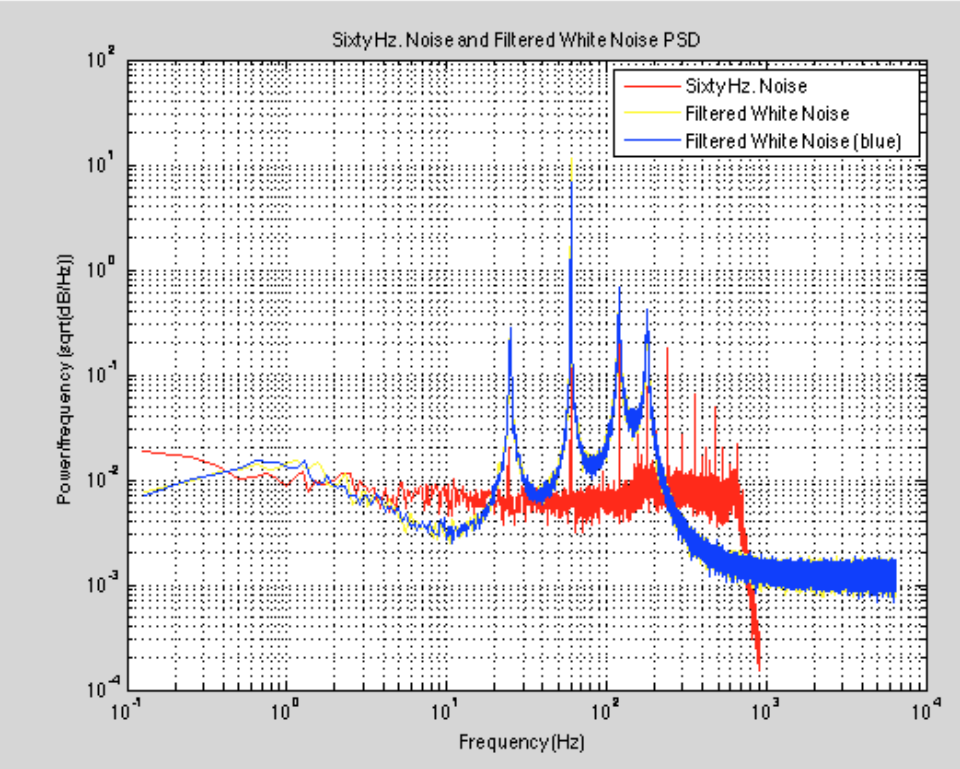

Figure 6: A Matlab plot comparing the power spectral density (PSD) of 60 Hz noise (taken from a terminated data channel) with the power spectral density of white noise with my electronic noise filter applied. The peaks are much more spread out in the fake noise, though the behavior is roughly similar.

in describing electronic noise. Indeed, Figure 6 takes into account general electronic noise as well as 60 Hz noise, but the 60 Hz noise clearly produced the most salient characteristics of the plot.

Now, lastly, we come to the results.

6 Results

All plots were done in Matlab.

6.1 Undamped

The undamped power spectrums match up quite well (though the resolution is not high, this is just an examination of the general properties). The dashed curves in Figure 7 represent the modelled power spectrum, while the solid curves are the power spectrums of the various degrees of freedom in the real suspensions. The biggest differences are the graphs of yaw and side. This is not surprising, given that yaw and side are coupled together and we did not take this coupling into account when laying out the mechanics. However, this shows that yaw/side coupling is a significant part of the power spectrum and can not be ignored. This is one basic example of how we further our own understanding with this project!

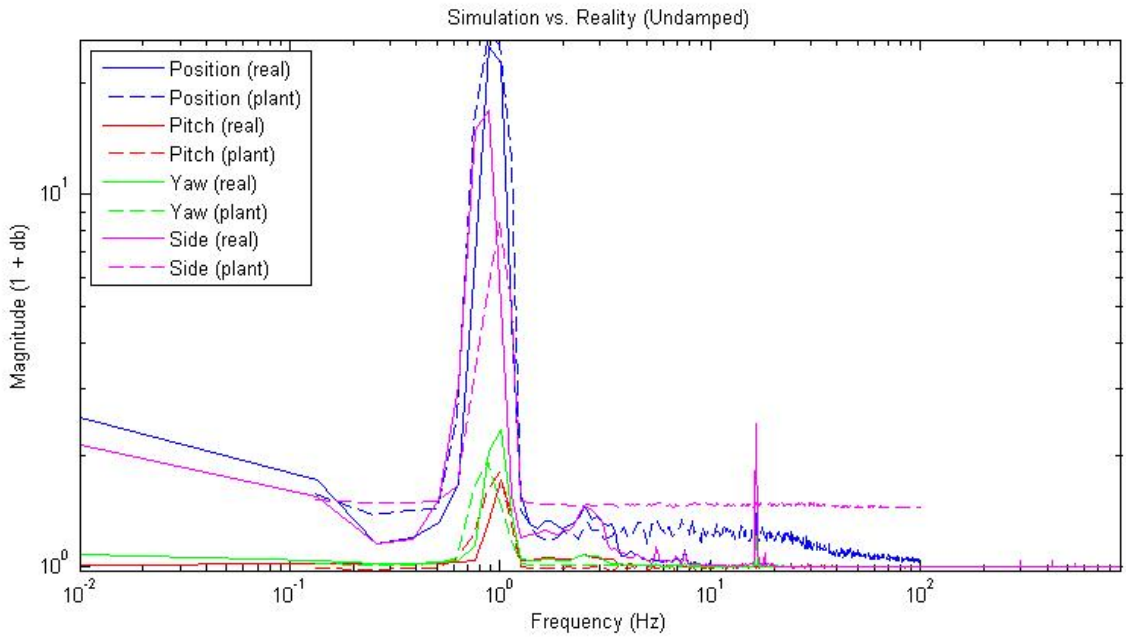

Figure 7: The undamped noise spectra of the model and reality. Note the order of magnitude difference between position/side and pitch/yaw.

6.2 Damped

On the other hand, the damped power spectrum doesn't match up quite so nicely, especially in the lower frequencies; see Figure **??**. There seems to be a significant order of magnitude difference between our pitch/yaw and the real pitch/yaw, an order of magnitude difference that did not appear in our first graph. This is intriguing, as the data for this graph was taken before we attempted to compensate for the $2pi$ factor in some of our conversions from Hz to radians!

In addition, the noise level falls much quickly in real-life than it does in our simulation - this is likely because our seismic noise curve is too steep. Since seismic noise is the most significant source of noise, any inconsistency in noise level between the two power spectra likely comes from it. We can continue experimenting with the seismic noise to make it match up better with reality.

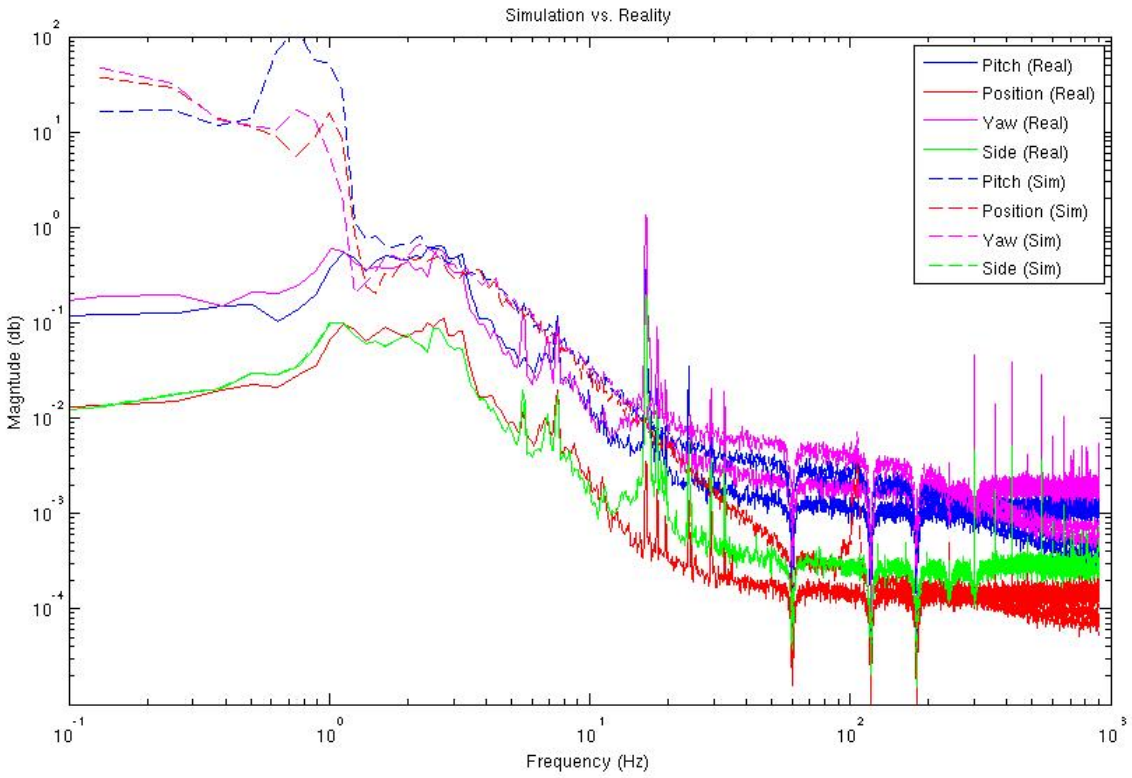

Figure 8: The damped noise spectra of the model and reality. Note the high low-frequency noise (in the case of the model) and the steep noise curve thereafter. Also note the effects of the notch filter meant to surpress 60 Hz noise at 60 Hz.

7 Future Developments

There is still plenty of room for development in this project. A future direction could be a network of simPlants representing the whole interferometer. In order to do this truly accurately, the noise sources I neglected (thermal, local oscillator, laser) would need to be given a thorough treatment. In addition, there are two degrees of freedom that we can not actuate on (so I did not include them) but are still there: bounce (motion in the z direction, or up and down) and roll. My model also only takes data from the OSEMs into consideration while in actuality data from optical levers is used to detect angular motion, since they have been known to be less noisy and more reliable. The matrix that is used to convert from the OSEM basis to the DOF basis and back is periodically updated to take new placement of suspensions into account. Every time it updates, the model has to be manually corrected in order to take these changes into account. However, a program can and and should be written to sync up the changes in the OSEM to DOF matrix with the model automatically. Lastly, there is coupling between side (y) and yaw that I did not take into account. These are all viable areas for future improvements.

8 Acknowledgments

First of all, I would like to thank my mentors, Jamie Rollins and Rana Adhikari, for their support, encouragement, and wisdom throughout the summer. Thanks to my fellow SURFs - Masha, Liz, Eric, and Yaka we had a wonderful summer together. Thanks to Jenne for the much-needed help, Koji for the kindness, Seiji for the jazz, Den for the Russian, Steve for introducing me to soccer talk, and the rest of the 40m lab for making it an amazing place to do research. Also thanks to Joseph B., the father of simPlants. Lastly, thanks to the NSF and the LIGO SURF program at large - I can't imagine what kind of disruptions SURF students bring year after year, yet you still support us. Thank you!

References

[1] Adhikari, Rana. *Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae*. http://www.ligo.caltech.edu/docs/P/P040032-00.pdf.

[2] Heefner, J., R. Bork, and R. Abbott, *THE LIGO INTERFEROMETER SENSING AND CONTROL SYSTEM*. LIGO.

[3] Yamamoto, Hiro, *Simulation tools for future interferometers*. LIGO. http://iopscience.iop.org/1742-6596/32/1/061/pdf/1742-6596_32_1_061.pdf.

[4] Bork, R., and M. Aronsson, *AdvLigo CDS Real-time Code Generator (RCG) Application Developers Guide*. LIGO. https://dcc.ligo.org/public/0001/T080135/003/T080135-v3.pdf.

[5] Saulson, Peter R., *Fundamentals of Interferometric Gravitational Wave Detectors*. Singapore: World Scientific, 1994.

[6] Sigg, Daniel, *Frequency Response of the LIGO Interferometer*. LIGO. https://dcc-llo.ligo.org/public/0028/T970084/000/T970084-00.pdf.

[7] Pelc, James S., *Radiation Pressure Effects in a Suspended Fabry-Perot Cavity*. MIT: Department of Physics, 2006.

[8] Corbitt, Thomas, et. al., *Measurement of radiation-pressure-induced optomechanical dynamics in a suspended Fabry-Perot cavity*. Physical Review A 74, 021802(R) (2006).

[9] Ward, Robert L., *Length Sensing and Control of a Prototype Advanced Interferometric Gravitational Wave Detector*. Caltech: Department of Physics, 2010.

[10] Black, Eric D., *An introduction to PoundDreverHall laser frequency stabilization*. American Journal of Physics 69, Issue 1 (2001).

[11] Black, Eric D., *Notes on the Pound-Drever-Hall technique*. LIGO DCC T980045- 00- D (1998).

[12] Findley, T., *Analysis of the Pendular and Pitch Motions of a Driven Three-Dimensional Pendulum*. LIGO DCC P060034-00-R (2006).

[13] Rakhmanov, Malik, *Dynamics of Laser Interferometric Gravitational Wave Detectors*. Caltech: Department of Physics, 2000.

[14] Driggers, Jenne, *Advanced Suspension Diagnostic Procedure*. LIGO DCC G1101001-v1 (2011).

[15] Betzwieser, Joseph, *Investigating Electromagnetic Interference in the Control Electronics of the Advanced LIGO Detectors*. LIGO. http://www.its.caltech.edu/~kgl/REU2012/REU2012-preliminary.pdf.