Adaptive Modal Damping for Advanced LIGO Suspensions

by

Brett Noah Shapiro

B.S.,Engineering Science (2005)
The Pennsylvania State University

S.M., Mechanical Engineering (2007)
Massachusetts Institute of Technology

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

Author .....
Department of Mechanical Engineering
May 11, 2012

Certified by .....
Nergis Mavalvala
Professor of Physics
Thesis Supervisor

Certified by .....
Kamal Youcef-Toumi
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by .....
David E. Hardt
Chairman, Committee on Graduate Students

Adaptive Modal Damping for Advanced LIGO Suspensions

by

Brett Noah Shapiro

Submitted to the Department of Mechanical Engineering
on May 11, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

Gravitational waves are predicted to exist by Einstein's Theory of General Relativity. The waves interact extremely weakly with the surrounding universe so only the most massive and violent events such as supernovae and collisions of black holes or neutron stars produce waves of sufficient amplitude to consider detecting. The Laser Interferometer Gravitational-Wave Observatory (LIGO) aims to pick up the signals from these very faint waves.

LIGO directs much of its effort to the areas of disturbance rejection and noise suppression to measure these waves. The work in this thesis develops an adaptive modal damping control scheme for the suspended optics steering the laser beams in the LIGO interferometers. The controller must damp high quality factor mechanical resonances while meeting strict noise and disturbance rejection requirements with the challenges of time varying ground vibrations, many coupled degrees of freedom, process noise, and nonlinear behavior. A modal damping scheme is developed to decouple the complex system into many simpler systems that are easily controlled. An adaptive algorithm is then built around the modal damping scheme to automatically tune the amount of damping applied to each mode to achieve the optimal trade-off between disturbance rejection and noise filtering for all time as the non-stationary stochastic disturbances evolve. The adaptation is tuned to provide optimal sensitivity to astrophysical sources of gravitational waves. The degree of sensitivity improvement is analyzed for several classes of these sources.

Thesis Supervisor: Nergis Mavalvala
Title: Professor of Physics

Thesis Supervisor: Kamal Youcef-Toumi
Title: Professor of Mechanical Engineering

Acknowledgments

My time with LIGO at MIT has been an absolutely invaluable experience. It was during this time that I learned what it means to be an engineer and a scientist. I must first thank Sam Finn who introduced me to LIGO by allowing me to attend his weekly lab meetings while I was just a freshman at Penn State. David Shoemaker deserves a special thanks for connecting me to the LIGO group at MIT when I began my graduate work and for being a constant source of support throughout.

Each of my four committee members deserve a special and personal thanks. First, my advisors Nergis and Kamal. Nergis's unwavering support and her extensive knowledge of the science behind LIGO contributed immensely to this work. Nergis also has an awe-inspiring ability to maintain the larger scope of a project and find direction among infinite possibilities. This is a quality I will strive to replicate throughout my career. Kamal's expansive expertise in the field of controls kept me grounded in the engineering world amidst a sea of scientists. His guidance in searching for controls contributions and his steady encouragement to justify, explore, and explain all directions taught me to turn my ideas into work worthy of publication. From my supervisor Rich, I learned first hand how to setup, run, and debug the experiments in the lab. Rich showed me the excitement of conducting experiments, and the patience it often requires. Professor Slotine encouraged me to take this work to the next level. With this encouragement, the contributions established in this thesis are greatly enhanced.

Professor Scott Hughes contributed immensely to the science goals here. His class on relativity that I sat in on during the fall of 2011 infinitely expanded my understanding of the subject matter (or at least my realization of how little I in fact knew). Additionally, our few but dense conversations directly led to the identification of further science benefits to LIGO from this work.

Thanks to Myron for all the help in the lab. Myron has an ability I can only equate with magic (seriously) to just make things happen in a seemingly effortless way. Without Myron's help I would never have gotten out of the lab. To Marie, I

could never offer enough thanks. All the travel help, purchases, cookies, fax machine rescuing, advice, etc...I would have been lost without. Thanks to Fred for all the computer help and advice. He directly supported the computation required by the simulations in this thesis. Not to mention he often kept my computers running when they needed help.

To Alex, Rolf, and Jay all the way at Caltech for taking many 'please help' calls regarding LIGO electronics and software. I'll never forget the time Jay flew cross country on a single day's notice just to repair one of those pesky ESD cables. You are missed, my heart goes out to you and your family.

Matt, Lisa, Sam, and Peter. I am immensely grateful for all your cavity help and all the discussions we had about how my work fits in with the greater scope of LIGO.

To the entire suspensions group. I wish I could thank everyone of them here. I would like to call out some of them personally. Norna, Janeen, Mark, and Ken Strain for many helpful conversations regarding all things suspensions. Thanks to Jeff for becoming the only other full time suspension's team member at MIT and driving the implementation of much of my work at the sites. The guys at RAL: Joe, Ian, and Justin for showing me that a few masses on springs can provide years worth of fascination.

Thanks to everyone else at MIT. To call out a few: Jonathon Soto for all those deep conversations we had in the brief time we shared an office. I always valued bouncing ideas back and forth with you. Thanks to John Miller, Sharon, and Rutu for all your help in the BSC; Laurent for introducing me to the concept of modal damping; Fabrice for providing another engineering perspective; Ruslan for our data analysis conversations; and Rai for your input and passion.

A special thanks to Nevan for sharing an apartment in Tang for 3 years and for sharing some special drinks. To Kyle for your many visits from Hatfield that gave me some distracting sanity. And of course for all those awesome board games.

Finally, my family. Mom, dad, and Nina, and Darren, I wouldn't have gotten this far without you. To Jamie for your constant and unwavering support, often from very far away. Without you this thesis would not exist.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The remaining years of this decade promise to be an exciting time for gravitational wave astronomy. Current upgrades to existing gravitational wave observatories such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) are likely to make direct observations of these waves from massive compact events in space possible for the first time. For some of these events, such as black hole mergers and early universe processes, gravitational waves may be the only detectable radiation emitted. For others such as supernovae, these waves can augment electromagnetic observations, providing ever deeper astrophysical understanding.

Gravitational waves have not yet been directly observed because they interact so weakly with any possible detector design. Consequently, the science available to current (and foreseeable future) observatories is limited by many sources of noise. One class of these noise sources for ground based detectors such as LIGO consists of non-stationary disturbances coming through the ground. These disturbances are constantly evolving in time depending on weather, tides, and man-made activities. Many active control loops are involved to compensate for these disturbances to maintain the detectors at a set operating point. The main contributions in this thesis target the development of an adaptation algorithm to optimize certain control loops for the observation of astrophysical sources of gravitational waves.

1.1 Behavior of Gravitational Waves

Gravitational waves (GWs) are currently undetected phenomena predicted to exist by Einstein's Theory of General Relativity. In this theory accelerating mass produces GWs in a manner analogous to the way accelerating charges produce electromagnetic waves. Additionally, GW's interaction with surrounding space is incredibly weak. Only the most massive and violent events in the universe such as supernovae and collisions of black holes or neutron stars produce GWs strong enough to consider detecting [1].

The effect of GWs is quite different from electromagnetic waves. As they propagate they compress and stretch space. Consider the ring of particles in Figure 1-1 below. If the axis of propagation is perpendicular to the page then the shape of space is altered in the plane of the page, perpendicular to propagation. Two perpendicular axes in this plane simultaneously have opposite effects; one is stretched while the other is compressed.

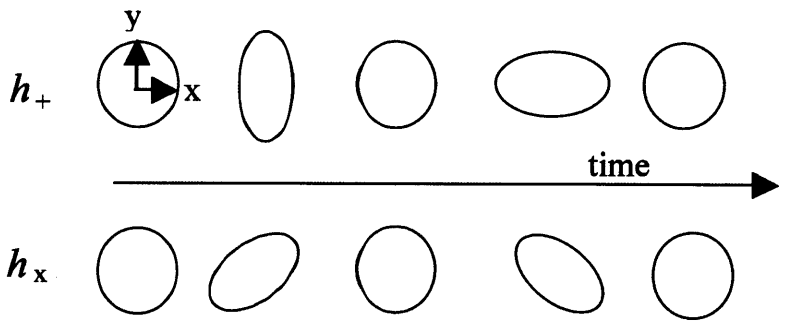

Figure 1-1: Effect of a GW on a ring of particles. The top half shows the effect of a wave polarized with axes, x and y, parallel to the vertical and horizontal axes of the page. The lower half shows the effect of a wave polarized at 45° relative to the upper wave. Space is alternately compressed and stretched along the wave's x and y axes as it propagates through the plane [1].

The figure shows two different polarizations of a GW [14, 15]. Polarization is defined as the tilt of the plane wave's x and y axis relative to the observer. The upper half of the figure, h+, is a case when the polarization is parallel with the vertical and horizontal axes of the page. The lower one is a case when the polarization is tilted 45°. At the observational distances we are considering here, all waves are assumed to

be planar.

The stretching of space is similar to the concept of mechanical strain in that it is a proportional effect. If two objects are far apart, the distance between them will change more than for two objects that are close together. The amount of strain produced by a GW, h, is proportional to the strength of the wave. For a GW of strength h, the specified distance L is altered by $\Delta L = hL$. For more details on GWs and associated general relativity see [14, 15].

1.2 Sources of Gravitational Waves

Sources of GWs that are expected to be strong enough for detection by current and upcoming observatories include coalescing compact binaries, pulsars, supernovae, and a stochastic background. These sources encode in their waves rich information about general relativity and astrophysics. A non-exhaustive summary of these sources and the scientific knowledge they encode in their waves is presented in this section. More information can be found in [16, 17, 18].

1.2.1 Binary Inspirals

A binary inspiral is a coalescing pair of compact, massive objects such as black holes and neutron stars. Such events begin with the objects orbiting around their common center of mass. The orbital distance and period gradually decay due to energy loss through the emission of gravitational radiation. The pair begins spiraling into the center of mass and the GW signal increases both in frequency and amplitude creating a chirp-like signal. Finally both objects crash into each other emitting a final burst of gravitational radiation. Figure 1-2 plots the waveform for last 0.25 seconds of two 1.4 solar mass inspiraling neutron stars.

Inspirals are a potentially rich source of information. Their GWs encode information about their population in the universe and mass and spin properties. The tidal disruption of a neutron star merging into a black hole produces information on the neutron star matter's equation of state of which there is still much uncertainty. The

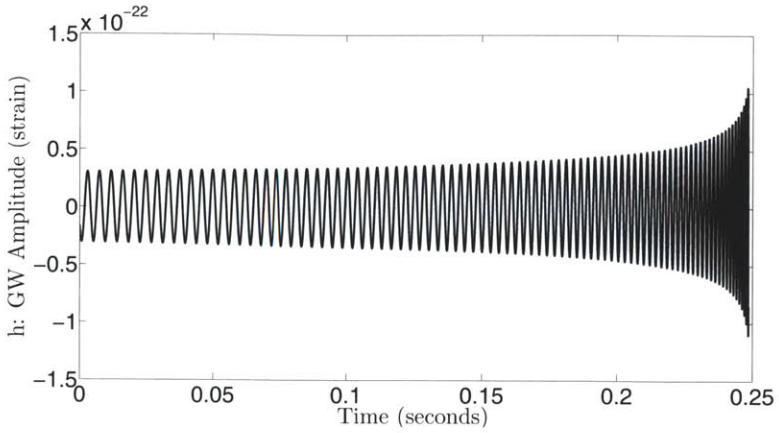

Figure 1-2: The last quarter second of an inspiral waveform for two 1.4 solar mass neutron stars 100 Mpc from Earth.

merging of two black holes creates the most extreme warping of space time, providing detailed information on the as yet untested nonlinear strong field dynamics of general relativity.

1.2.2 Pulsars

Pulsars are rotating neutron stars that emit GW radiation due to their rotation. This is another mechanism for neutron stars to emit observable gravitational radiation. Spherical symmetry prevents wave emission, but if the star is slightly asymmetric (its ellipticity) and rotating quickly enough, then waves strong enough to observe from Earth may be produced at twice the pulsar's rotational frequency.

Long term observations of these waves as the Earth revolves around the sun reveal to high precision the location in the sky of the pulsar, which enables the blending of electromagnetic observations. Comparison's between GW and electromagnetic observations can divulge inhomogeneities in the density of the pulsar and its ellipticity. GWs can also carry more detailed information about the star's structure such as the properties of its crust, crust-core interactions, viscosity, etc. There is also the possibility of observing a so-called r-mode oscillation in fast spinning stars, which is an unstable oscillation in the GW induced flow of material within the star. This oscillation contributes to the strength of the waves, thus enhancing the oscillation further until dissipative forces bring it to an equilibrium.

1.2.3 Bursts

Bursts represent a class of cosmic events characterized by high energy output on short time scales. One example is the type-II supernova, which is the explosive death of a massive star and the subsequent collapse of the core to a neutron star or black hole. The exact evolution of a supernova is not well known, so this source provides an excellent opportunity to expand scientific knowledge of such phenomena. Many of these events have already been observed through neutrino and electromagnetic emissions. Merging these observations with concurrent GWs would produce a more complete picture of these events than any that has been achieved to date. GW emission requires the collapse of the stellar core to have spherical asymmetry, however it is believed this is usually the case. For example, if the core is spinning its collapse will not be symmetric and strong waves can be produced. Since these events lead to neutron stars and black holes there is a considerable opportunity to study these compact objects just as they are born in ways that are not achievable otherwise.

Gamma ray bursts include another set of possible sources in this class. These events emit strong bursts of gamma rays over a period of a fraction of a second to 100 seconds. They are observed approximately once per day by dedicated satellites, in particular by the Swift mission [19]. The specific triggers for gamma ray bursts are not well formalized, but it is believed they come from a variety of sources such as supernovae, accretion around a black hole, and the merger of compact binary objects.

1.2.4 Stochastic Background

Stochastic background radiation encompasses a final class of GWs. The background could consist of either an incoherent superposition of many discrete weak sources or primordial radiation from very early times following the Big Bang. The stochastic signals are expected to be broadband and extremely weak. They are generally modeled as being isotropic, stationary, and Gaussian, though these assumptions are not necessarily true. There are many ways to break the anisotropy assumption. A simple example is for a background dominated by discrete sources in the Milky Way, such

as binary white dwarfs. Since the galaxy is not spherical the strongest signals would point towards the bulk of the Milky Way in the sky. The stationary assumption is reasonable given that the universe evolves on cosmological timescales that are much greater than the anticipated lifetime of GW observation. The Gaussian assumption is reasonable in particular for a background dominated by early universe radiation.

Current understanding of the nature of this background is rather poor. It is unknown for instance which of the two stochastic components will dominate, the primordial radiation or the superposition of discrete sources. Conversely however, there is room for large gains in knowledge. The primordial radiation, if visible, will contain information about the universe as early as 10^{-22} s after the Big Bang. In comparison, our current earliest view of the universe comes from Cosmic Microwave Background Radiation (CMBR) originating at 10^5 years after the big bang [20]. Observations of the stochastic background will help answer many of these unknowns.

1.3 LIGO

For the first time the direct detection of GWs may now be possible due to the work of observatories such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) [3, 21, 22]. To date, no direct detections of GWs have been made due to the weak interactions of gravitational radiation. A typical wave produced from such powerful sources as described in section 1.2 is expected to induce a strain of only 10^{-22}, about 10^4 times smaller than the diameter of a proton for the 4 km observation paths used in LIGO. Earth-based detectors, such as LIGO, have the added complication of natural and man-made noise traveling through the ground. Nearly all the work of these detectors up to this point is focused on developing ways to amplify the signal and reduce sources of noise that will wash out these infinitesimally small signals.

In order to bring observations of gravitational waves into the realm of regular astronomy, the second generation of LIGO known as Advanced LIGO [3] is currently under construction. Advanced LIGO increases the sensitivity of the first generation of LIGO, known as Initial LIGO by tenfold over a broad frequency band. This increases

the volume of space visible to LIGO by a factor of 1000, which brings the expected GW detection rate from about once is a handful of years to as much as once a day.

1.3.1 Advanced LIGO Layout

Advanced LIGO installation is currently underway within the existing LIGO vacuum envelopes at the two LIGO sites in Livingston, Louisiana and Hanford, Washington. See Figures 1-3 and 1-4 for photos of these observatories. To improve upon the Initial LIGO sensitivity nearly all hardware with the exception of the vacuum system itself will be upgraded. These systems include seismic isolation, optics, and lasers [3].

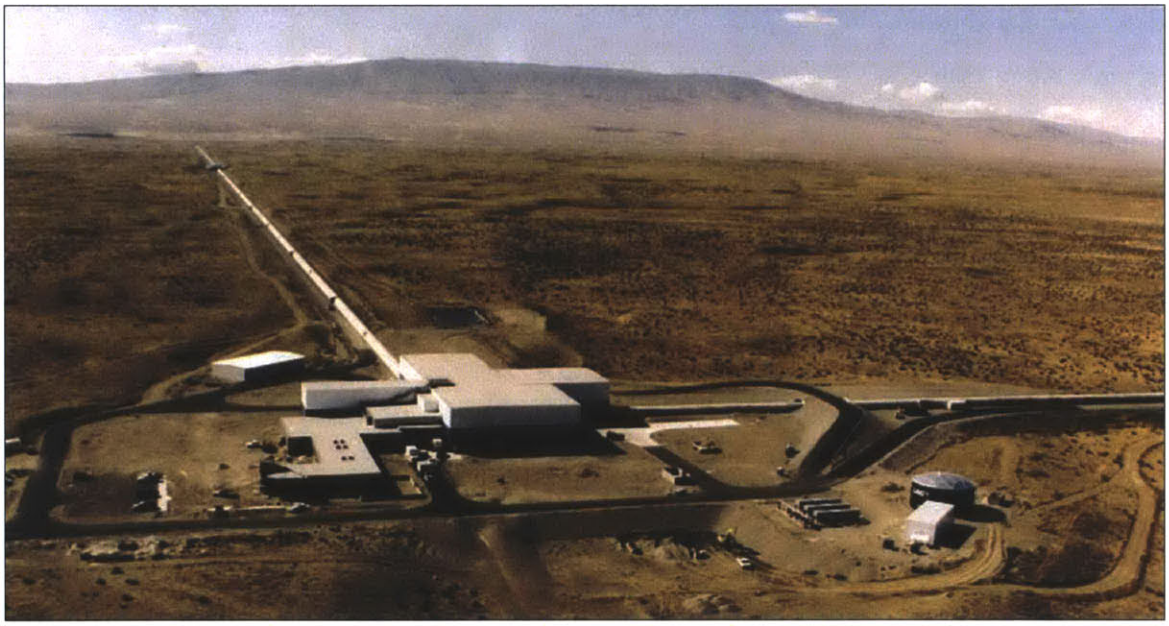

Figure 1-3: A photograph of the LIGO Hanford Observatory in Washington State [2].

Each Advanced LIGO detector consists of a Michelson interferometer with 4 km long Fabry-Perot cavities in each perpendicular arm, used to measure the differential arm length caused by the passing wave. A long arm length was chosen in order to maximize the length change induced by the strain of the wave. The interferometer is housed in a vacuum envelope evacuated to about 10^{-9} Torr to prevent interference from gas particles.

Figure 1-5 shows the schematic diagram of these interferometers. The exact details of how these detectors work are rather complicated, but the essential principles can

Figure 1-4: A photograph of the LIGO Livingston Observatory in Louisiana. Courtesy of Aero Data.

be described as follows. A laser injects 1064 nm (infrared) light into a beam splitter that splits the light into the two orthogonal arms. A Fabry-Perot cavity in each arm of the Michelson, comprising optics referred to as an input test mass (ITM) and an end test mass (ETM), stores the laser light to increase the phase sensitivity of the interferometer. The test masses are suspended as pendulums in order to isolate them from ground motion and act as free test particles, such as those in the rings of Figure 1-1. A photon will make approximately 100 round trips between these optics in order to amplify the length measurement of each arm. The light then recombines at the beam splitter and continues on to a photodetector. Measuring the intensity of the light at this photodetector provides a measure of the phase difference of the light in each arm, and thus the differential length [1].

The purpose of multiple observatories is motivated by several factors. One is to locate which part of the sky the GW came from and determine its polarization. Another is to create veto scenarios so that GWs can be distinguished from local noise sources such as a heavy truck driving down a nearby road. A GW of cosmic origin should be present coincidentally at both sites whereas the truck will not.

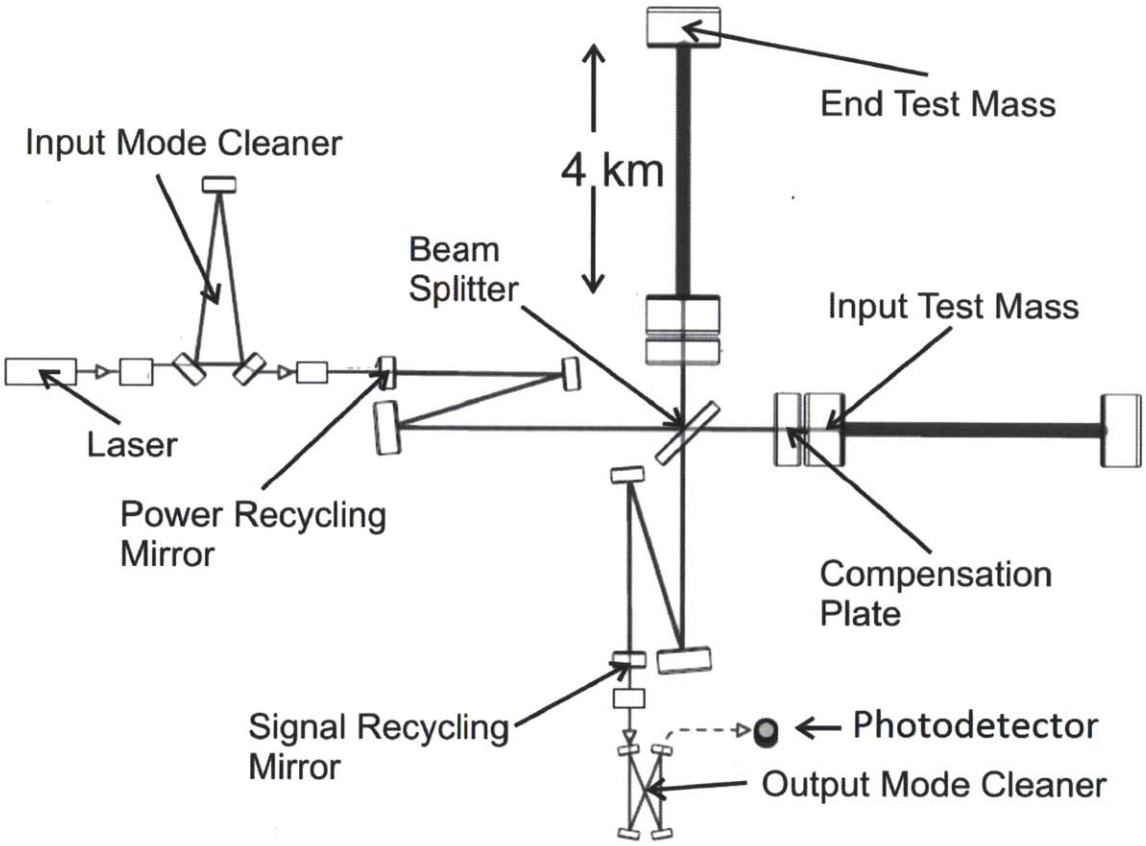

Figure 1-5: A top view of the optical layout of an Advanced LIGO observatory (not to scale). A propagating GW passing through the observatory alters the differential length of the two 4 km perpendicular light storage arms. This differential length is measured with the amplitude of the light on the photodetector. Adapted from [3].

1.3.2 Sensitivity Limiting Noise Sources

In order to make the detectors sensitive enough to see GWs, many different sources of noise and disturbances in the LIGO detection band (10 Hz to 8 kHz) have to be eliminated or reduced. The dominating sources are categorized into seismic, gravity gradients, thermal, and quantum noise. Quantum noise consists of both radiation pressure and shot noise. Each source contributes to specific parts of the spectrum [3]. Figure 1-6 shows each one relative to the Advanced LIGO design sensitivity. The measured Initial LIGO sensitivity from the Livingston Observatory is included for reference.

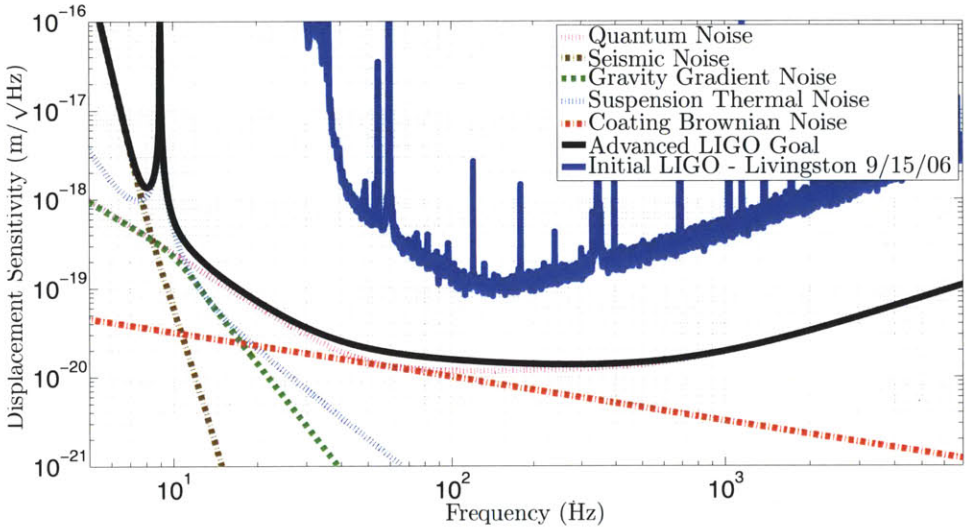

Figure 1-6: Projected displacement spectral density for the Advanced LIGO design with contributing noise sources. The strain spectral density sensitivity is found by scaling the displacement down by 4000 m, the length of the interferometer arms. The measured Initial LIGO displacement is plotted for reference. Adapted from [3, 4].

Seismic disturbances dominate at frequencies below about 10 Hz. The term seismic is used to collectively refer to all sources coming through the ground. These sources include actual seismic motion from tectonic plates and volcanic activity. However, it also includes noise from ocean waves crashing onto the shores as well as man-made activity. A nine order of magnitude attenuation is needed to meet the Advanced LIGO sensitivity requirement at the 10 Hz start of the GW detection band.

The 10 Hz to 40 Hz region is dominated by radiation pressure noise, thermal noise,

and a possible contribution from gravity gradient noise. Radiation pressure noise is the disturbance from the momentum transfer of random numbers of photons from the interferometer laser reflecting off the test masses. It scales with the square root of the laser power [23]. The thermal noise in this band is associated with the mechanical dissipation of the test mass suspension systems. Thermal energy in the materials making up the suspensions (particularly the silica fibers supporting the test masses) couple to the motion of the test mass along the interferometer axis. Gravity gradient noise refers to classical (as opposed to relativistic) fluctuations in the local gravitation field from variations in the distribution of nearby mass, particularly density fluctuations in the local ground.

From about 40 Hz to 200 Hz radiation pressure combines with shot noise and thermal noise in the test mass optical coatings to limit the sensitivity. Radiation pressure and shot noise are both quantum noises associated with the random distribution of discrete photons in the laser light. Radiation pressure physically moves the test masses as described above, whereas shot noise is a sensing problem. The random number of photons at any given moment falling on the photodetector measuring the interferometer output generates the shot noise contribution. It scales with the inverse square root of the laser power. Consequently, there is a fundamental design trade-off when choosing the laser power since radiation pressure scales with the square root of the laser power. Frequency dependent squeezing technology is currently under development to improve upon this quantum noise trade-off [24]. The optical coating thermal noise has the same process as the suspension thermal noise except that here the source is in the coating on the test mass surface. Beyond 200 Hz shot noise will be the sole dominating noise.

1.3.3 Seismic Isolation Systems

Advanced LIGO requires the installation of sophisticated seismic isolation systems to reach design sensitivity in the seismic and thermal noise dominated band of the spectrum. This isolation is achieved with three cascading systems. For the test masses, where the noise contributions are most critical, these systems encompass a

total of seven stages of both passive and active isolation. Figure 1-7 is an illustration of how the three systems fit together in the vacuum chambers that house the test masses. This section focuses on these particular test mass systems. Similar systems with fewer stages exist for many of the auxiliary optics within the interferometer [25].

The first system, known as the Hydraulic External Pre-Isolator (HEPI), functions outside the vacuum chambers that house the test masses and internal isolation systems. HEPI is a single stage that senses and actively removes ground motion by about an order of magnitude between the frequencies of 0.1 Hz to 10 Hz. Its specifications call for actuation in all six degrees of freedom (DOFs), force generation up to 2000 N, a range of ± 1 mm in translation, ± 1 mrad in rotation, and a noise floor no greater than $10^{-9}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 1 Hz. The actuation is applied to the four corners of the support tubes that carry the internal seismic isolation systems. Bellows allow the motion to be transmitted to the inside of the vacuum chamber. HEPI has inertial and displacement sensors which feedback to collocated laminar flow hydraulic actuators. Seismometers placed nearby on the ground further enhance isolation with feedforward control [3, 5].

The second system, known as the Internal Seismic Isolation (ISI) system is shown in a detailed drawing by Figure 1-8. It supports the suspensions inside the chambers and is directly supported by HEPI. It is a two stage system providing both active and passive isolation. The ISI has inertial and displacement sensors that feedback to electromagnetic actuators to provide active isolation from about 0.2 Hz to 30 Hz. The two stages are suspended by cantilever springs which provide passive isolation beyond the bandwidth of the active isolation control. Overall, the ISI isolates by about a factor of 300 at 1 Hz and 3000 at 10 Hz [3, 5].

The second stage of the ISI is equipped with an optical table from which a passively isolating test mass suspension hangs. These suspensions are chains of four stages hanging from each other, where the bottom stage is an interferometer optic serving as a test mass. The test mass suspensions are known as the quadruple pendulum. The quadruple pendulum is describe in detail in Chapter 2.

Advanced LIGO's non-test mass optics receive seismic isolation as well. Their

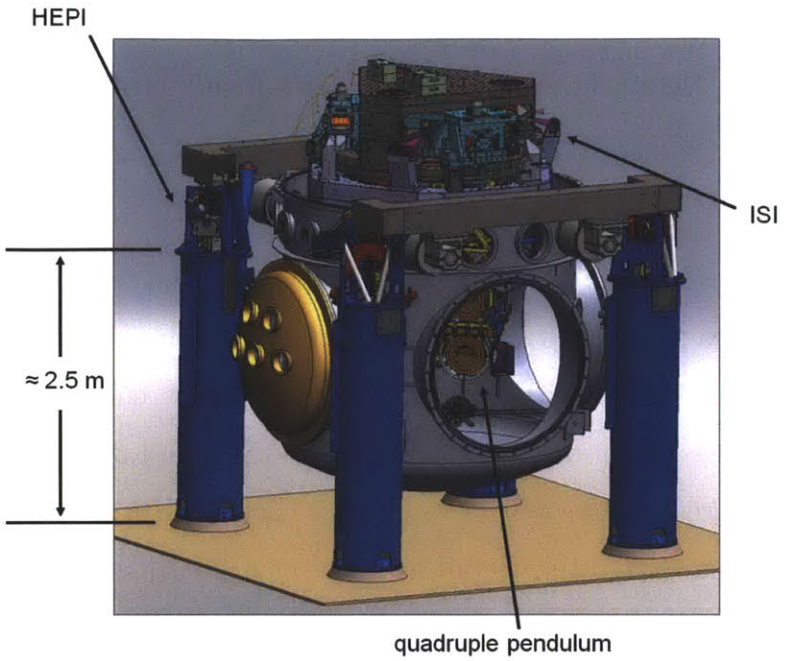

Figure 1-7: Layout of the three cascading systems of seismic attenuation for Advanced LIGO's test masses: HEPI, a single stage active isolation system external to the vacuum chamber; The ISI, a two stage active and passive isolation system inside the vacuum chamber; and a quadruple pendulum, a passively isolating system whose bottom stage is a test mass. Adapted from [5, 6].

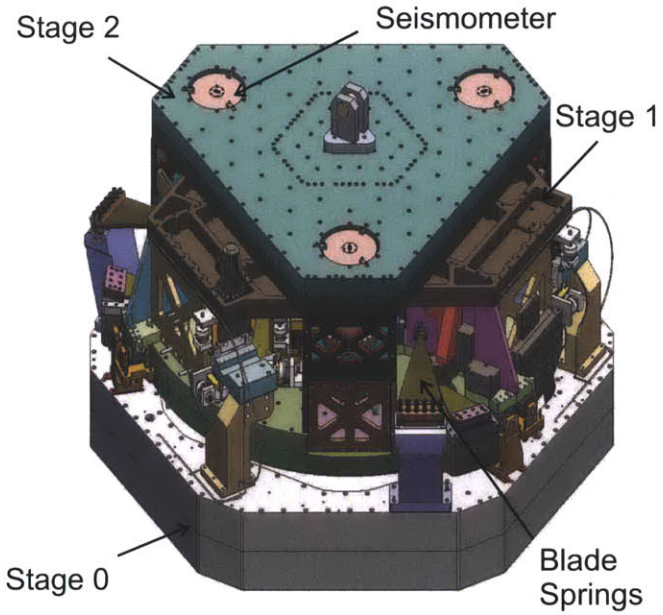

Figure 1-8: The two stage active-passive Internal Isolation System (ISI). The ISI is supported inside the vacuum system by HEPI. The ISI supports the quadruple pendulum that suspends the ETM and ITM optics [3].

isolation is achieved with simpler versions of the aforementioned systems, generally with fewer stages. Details of these systems can be found in reference [25].

1.3.4 Problems and Challenges

To measure the weakly interacting gravitational waves, LIGO's test masses must have a minimum relative displacement spectral density less than $10^{-19}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ between 10 Hz and 10 kHz. This sensitivity is realized at the low end of LIGO's spectrum through the advanced isolation systems described in Section 1.3.3. These systems have many coupled degrees of freedom that must be controlled to properly steer the test masses. LIGO requires that the overall relative root mean square (RMS) displacement and orientation between the test masses be less than 10^{-15} m and 10^{-9} rad respectively [26] as a result of the fact that the interferometer's sensitivity to higher order nonlinear terms increases with large motions of the test masses and isolation systems. These nonlinearities cause the generally low frequency ground vibrations to upconvert to frequencies where LIGO hopes to measure GWs. Many of the nonlinearity sources are either not measured or poorly understood, preventing their electronic subtraction from the interferometer's output. Some of the known or suspected sources result from laser light scattering off the walls of the interior vacuum system, laser light clipping from falling off the optics, the approximately quadratic nature of the interferometer output, higher order dynamics in the isolation systems, nonlinear actuator responses, and creak (sliding dislocations) in the materials of the isolation systems.

Additionally, the seismic disturbances the LIGO observatories experience is not constant in time. The magnitude of the noise changes depending on weather, earthquakes, and highly variable man-made (anthropogenic) disturbances. Ideally, all the control loops would be designed with sufficiently large gains to keep the test masses within an RMS of 10^{-15} m and 10^{-9} rad during the worst case seismic events. In practice however, non-negligible sensor noise exists in many of the feedback control loops, limiting these gains. Therefore it is likely that purely linear control of the interferometer will not provide optimal GW sensitivity for all seismic disturbances.

Indeed, Initial LIGO's sensitivity evolved with the seismic disturbances. Figure

1-9 illustrates an example of moderately large low frequency seismic disturbance upconversion at the Hanford Observatory during February of 2006. The top plot shows different cases of disturbances between 0.1 Hz and 4 Hz. The lower plot shows the resulting interferometer displacement sensitivity between 20 Hz and 200 Hz. In this case upconversion was most apparent for the disturbance at 1.2 Hz. Historically, the largest seismic disturbances experienced by Initial LIGO would move the interferometers so far beyond their linear range that they would go off-line until the period passed.

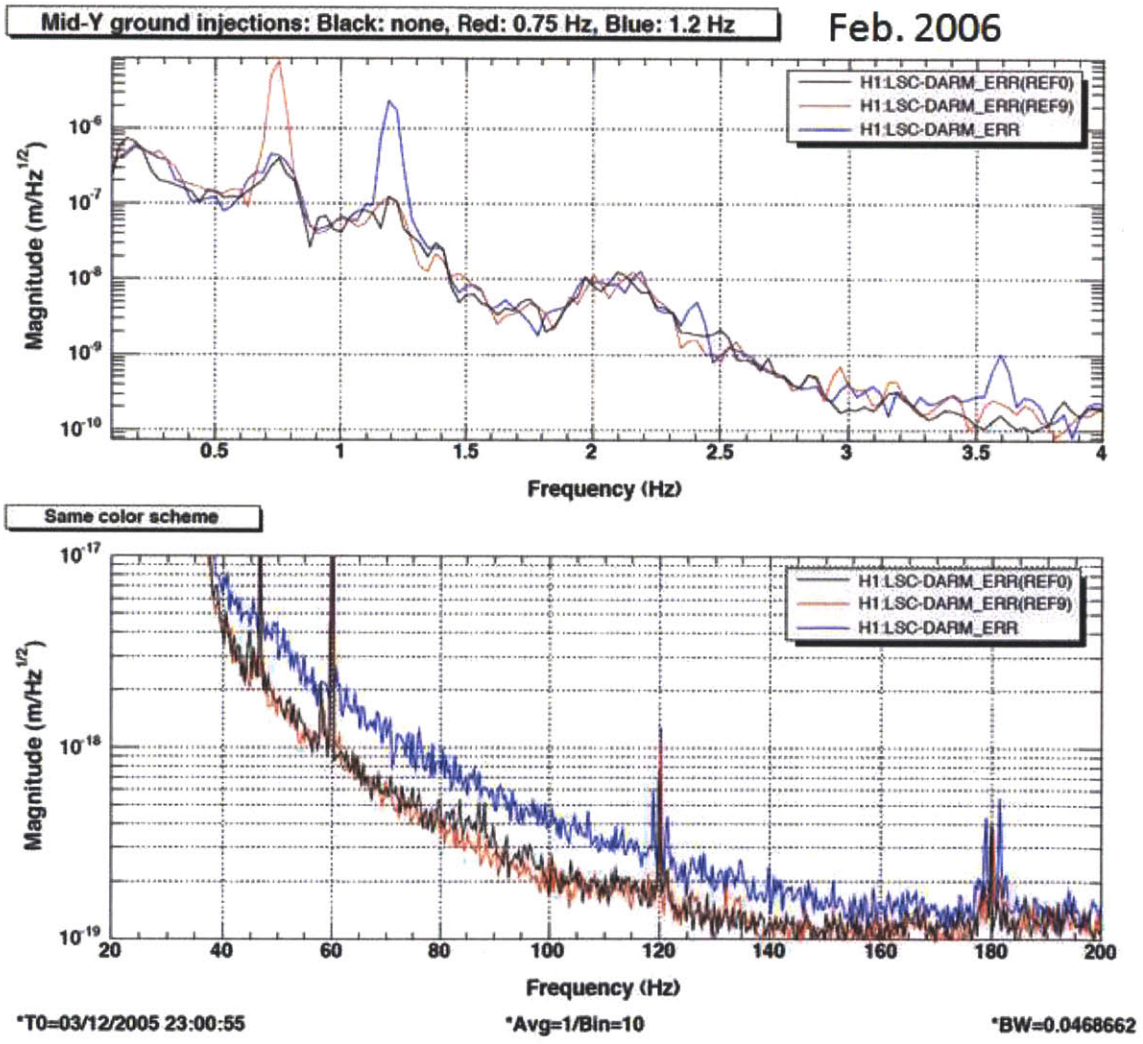

Figure 1-9: Upconversion of seismic disturbances in Initial LIGO at the Hanford, WA observatory during February 2006. Courtesy of Samuel Waldman.

To date, LIGO has devoted much effort to optimize the design of control loops that

position the interferometer's many coupled degrees of freedom while rejecting seismic disturbances with the minimal amount of process noise. These control loops were designed with many types of methods including traditional loop shaping techniques, feedforward control, and modal damping [27, 4, 28, 29]. Certain adaptation techniques have been explored as well. Driggers et al [27] explored adaptive feedforward cancellation of seismic disturbances and Zuo [30, 31, 32] explored adaptive feedback control of isolation systems with unknown or time-varying plant parameters. Reference [31] develops the concept of model reaching adaptive control. It improves upon model reference adaptive control by defining a goal for the system's dynamics, rather than its output. This change eliminates the need to measure ground disturbances. Reference [32] employs a frequency shaped sliding control method to robustly handle unknown or time varying system parameters.

1.4 Proposed Adaptive Approach for LIGO Suspensions

Previous adaptive feedback work focused on unknown or time varying isolation table parameters. It did not go as far as considering the pendulums or show how the control or seismic disturbances influenced the output of the interferometer. At this point in Advanced LIGO's development, the pendulums are well understood and their parameters are static. Adaptive feedforward work did consider the interferometer output under changing seismic conditions, but the adaptation was limited to time scales longer than 10 minutes and it can be very computationally intensive for good performance.

Additional gains could be made with adaptive feedback control on the pendulums that takes directly into account the interferometer performance and unique aspects of these pendulums. Such a controller would adapt its parameters in real-time in some optimal way to balance the trade-off between rejecting non-stationary seismic disturbances and the interferometer's sensitivity to upconversion. A need also remains

to define the optimal trade-off directly in terms of LIGO's science goals. Adaptive modal damping is developed to implement this real-time automatic tuning for the case of the local damping of the quadruple pendulums, which is one component of the interferometer's overall feedback control architecture. However, similar adaptive methods are applicable to other control components such as test mass angular control, control of auxiliary optical cavities, and the isolation loops of the seismic systems.

Local damping refers to any control strategy used to reduce or 'damp' the large amplitude displacements of the pendulums induced by the residual seismic disturbances amplified by high quality factor mechanical resonances. The proposed strategy, adaptive modal damping, incorporates a modal damping scheme enveloped by an adaption algorithm that adjusts the amount of damping applied to each mode in response to changing seismic disturbances. Modal damping is a convenient way to decouple the modes of vibration so they can be damped independently. Since each mode responds as a simple second order system, the design of each mode's compensator is minimally complex. As a result, it is relatively easy to find an optimal trade-off between damping and noise amplification for each mode. This trade-off is defined in terms of LIGO's sensitivity to the astrophysical sources of GWs listed in Section 1.2. The adaptive part of the control monitors the amplitude of each mode's response and scales the feedback gains of the compensators in real-time. More gain is chosen when a modal signal increases in amplitude and less gain when it decreases. The control also switches its adaptation rates for quick, coarse responses to sudden disturbances and slow, precise responses to steady disturbances.

1.5 Thesis Overview

The remaining chapters of this thesis support the main contributions of applying adaptive modal damping to the Advanced LIGO quadruple pendulums. Chapter 2 describes the design and requirements of the quadruple pendulum used to isolate and support the Advanced LIGO test masses. Chapter 3 describes the derivation of a model of the quadruple pendulum's equations of motion and the system identifica-

tion methods used to test the pendulum's performance against that model. Chapter 4 describes the methods used to adjust the model's parameters to match the measured performance. Chapter 5 uses this model to analyze the minimum actuator size needed to meet the quadruple pendulum's control requirements. It is shown that the current actuator designs are sufficient to meet these requirements. The concept of modal damping is developed in Chapter 6. The adaptive algorithm that automatically adjusts the modal damping parameters is described in detail in Chapter 7. Experimental results are presented for the control of a triple pendulum. Chapter 8 applies adaptive modal damping to optimize the performance of a quadruple pendulum for the detection of GWs from astrophysical sources. Simulated results are presented.

1.6 List of Contributions

1. **Modeling the quadruple pendulum** (Chapter 4): I propose a procedure to select which quadruple pendulum model parameter uncertainties and system identification measurements contribute the most to the pendulum's high dimensional dynamics (Section 4.1). The procedure is given a known bounded uncertainty for each of the parameters and a list of possible measurements relevant to the pendulum's performance. Those parameters whose uncertainty has a large influence on the measurements are selected. Those measurements that are sufficiently decoupled from the selected uncertainties are rejected. The result is a reduced system identification task and a simplified parameter estimation problem. The results also suggest what physical components of the pendulum can be improved to minimize the influence of parameter uncertainty on the measurements.

2. **Actuator Sizing** (Chapter 5): I develop a method for determining the minimum required actuator size for a system. This method uses singular value decomposition in the frequency domain to estimate how much drive each actuator requires and at what frequencies those drives are most efficient. The analysis confirmed the current quadruple pendulum design has enough dynamic range.

3. **Modal Damping** (Chapter 6): Modal damping is a method that permits the damping of each mode in a resonant system. The method simplifies the pendulum's active damping control by decomposing the problem of a high dimensional system with many coupled degrees of freedom into a set of simple decoupled single degree of freedom problems with similar dynamics. Modal damping of LIGO's multi-stage pendulum's is limited due to an incomplete set of sensors and actuators. Within these limitations I refined the state estimator optimization with a novel cost function that reflects LIGO's unique noise and damping requirements (Section 6.3). Additionally, I established the maximum closed loop damping ratios achievable for each mode. This calculation becomes very complex as the number of modes increases (Section 6.4).

4. **Adaptive Modal Damping** (Chapter 7): The adaptive component to modal damping allows the control system to automatically tune its own loop gains in real-time in response to changing environmental conditions. This adaptation allows the system to maintain an optimal trade-off between disturbance rejection and process noise amplification. My contributions from this experimental work include the development of a novel adaptive optimization scheme that responds to non-stationary stochastic variables to maintain a system, such as the LIGO interferometers, at an optimal operating point. The adaptation has the ability to switch its step size and step rate to either respond quickly to sudden large increases in the disturbance or converge slowly and carefully to the most optimal solution.

5. **Adaptive Modal Damping Applied to Advanced LIGO** (Chapter 8): I develop cost functions for the adaptive modal damping scheme that are optimal to the detection of GWs using the Advanced LIGO interferometers. These cost functions govern the way the modal damping adapts to time varying seismic disturbances. The impact on Advanced LIGO's sensitivity to GWs from binary inspirals and the stochastic background is considered.

Chapter 2

The Quadruple Pendulum

The quadruple pendulum is an extension of the three stage, triple pendulum designed for use in the German-British Gravitational Wave Detector (GEO600). Reference [33] details the design of the triple pendulum and is useful for reviewing the thought process used in extending the mathematics behind a single pendulum to an n-stage pendulum. References [34, 35, 36] describe the quadruple pendulum design.

2.1 Performance Requirements

2.1.1 Spectral Requirements

The purpose of the quadruple pendulum is to limit the contribution of seismically induced motions of the ITM and ETM optics within the GW detection band of Advanced LIGO starting at 10 Hz. Due to the strict nature of this requirement, care must be taken in the design of the pendulum such that it does not introduce other noise contributions that overpower and negate its seismic filtering property. Figure 2-1 shows the expected seismic motion of the table from which the pendulum hangs (blue line) and the required motion of the test mass (green line). To achieve the requirement given the input motion, the pendulum must provide more than six orders of magnitude isolation within the GW detection band.

The dominant two noise contributions stemming from the pendulum itself are

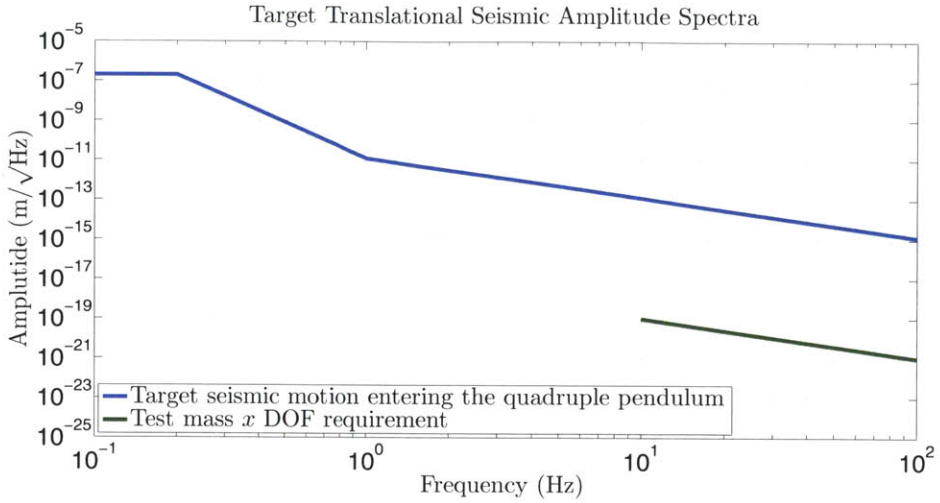

Figure 2-1: Blue line: expected motion along the translational DOFs of the table from which the quadruple pendulum hangs. The rotational DOFs have similar numerical values, but in units of radians rather than meters [4, 7]. Green line: the test mass requirement along the x DOF within the GW detection band. The pendulum must achieve more than six orders of magnitude isolation within this band.

thermal noise and process noise. The amount of thermal energy converted to kinetic energy in the test mass depends on the temperature of the surrounding environment. The only way to reduce this energy is to cool the system, which is not in the scope of Advanced LIGO. However, the spectral contributions of this thermal energy are optimized by carefully selecting the materials from which the pendulum is made. Process noise, or technical noise in LIGO terminology, comes from the feedback control systems that position the test masses within the interferometer. It includes sensor noise, actuator noise, and any other noise originating from the electronics of these controls. Technical noise is limited by careful design of electronics, sensors, actuators, and feedback loops.

The limits on the spectral motion of the test masses are summarized in Table 2.1. The maximum contribution of process noise to this limit is listed in the right hand column. Each test mass axis, or DOF, has a unique requirement depending on how strongly it couples to the interferometer output. These axes are defined by the coordinate system in Figure 2-2a. x lies along the interferometer axis, y is the perpendicular horizontal axis, and z is the vertical axis. Roll is rotation about x, pitch

is about y, and yaw is about z. The most stringent in-band (GW band) requirement on test mass motion is along the interferometer axis, x, i.e. the axis sensitive to the GW. Vertical test mass motion along the z axis is expected to directly contribute about one part in a thousand simply because the local direction of earth's gravity points in slightly different directions at the ends of the arms due to the curvature of the earth. Motion along the y axis is assumed, in the worst case, to contribute mechanically about a part in a thousand. The technical noise for these translational DOFs is restricted to 1/10 the overall limit, so that it may be considered a negligible component. Pitch and yaw motion are limited by the centering of the beam on the test masses. The value stated in Table 2.1 assumes the static test mass position is within 1 mm of the desired position. The spectral limit on pitch and yaw can be relaxed however if the static position of the test mass is known to be better. There is no specific rotational technical noise requirement provided the overall motion stays within the limit. A roll requirement was originally stated for sapphire test masses. Now that they are silica, this requirement needs reevaluation. In principal roll is limited only by the optical properties of the test mass, e.g. wedges and birefringence.

Table 2.1: The spectral limits on the motion of the test mass within the Advanced LIGO GW detection band. The technical noise is restricted to contribute less than one part in ten to the overall motion of the test mass along the translational DOFs. Motion along the rotational DOFs is merely specified as an upper limit [13].

DOF	**Overall Motion at 10 Hz rolling off at f^{-2}**	**Technical Noise**
X	$10^{-19}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$	1/10
Y	$10^{-17}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$	1/10
Z	$10^{-16}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$	1/10
Yaw	$10^{-17}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$	NA
Pitch	$10^{-17}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$	NA

2.1.2 Non-spectral Requirements

These non-spectral requirements are control related items that put restrictions on test mass RMS motion and suspension mode damping performance. Due to the thermal noise restrictions, the quadruple pendulum was designed with materials with very

low natural damping. Consequently, the resonant frequencies have very high quality factors (Qs). Many of these resonances are actively damped, such as the suspension modes (body modes of the suspension stages). The damping control is required to ring down the suspension modes within 10 seconds by at least e^{-1}, where e is Euler's number [13]. It is not necessarily required that both the damping and technical noise requirements are met simultaneously. For example, if a mode rings up the damping can be increased for 10 seconds, possibly breaking the technical noise requirement, until it rings down.

Active feedback loops are also employed to control the differential arm length in the interferometer. These loops require the test masses to follow each other within an RMS displacement tolerance of 10^{-15} m. Similarly, the pitch and yaw alignment of the test masses is to be controlled to an RMS angular tolerance of 10^{-9} rad [26].

2.2 Mechanical Design

The quadruple pendulum was designed to realize the ETM and ITM requirements specified in Section 2.1. Fig. 2-2 provides an illustration and a photograph of a prototype. It is a stable pendulum consisting of two hanging vertical chains of four stages each. The stages in each chain are numbered top down 1 through 4, where the fourth stage in the main (front) chain is a test mass consisting of a highly reflective optic. The reaction chain is used to provide a quiet actuation platform to filter any disturbance or noise that might couple through the actuators used to position the test mass. In each chain stages 1 and 2 are approximately 22 kg and stages 3 and 4 are a combined 80 kg. The main and reaction (back) chains are about 2 m from top to bottom. At stage 4 they hang 5 mm apart for the ETMs and 20 mm for the ITMs. Maraging steel cantilever springs support stages 1 through 3 and provide isolation along the vertical axis. The lowest two stages of the main chain are made of fused silica glass and are connected by welded fused silica fibers 400 μm in diameter. In this way, stages 3 and 4 consist of a single monolithic piece of fused silica. This material was chosen because its low loss property has advantageous thermal noise

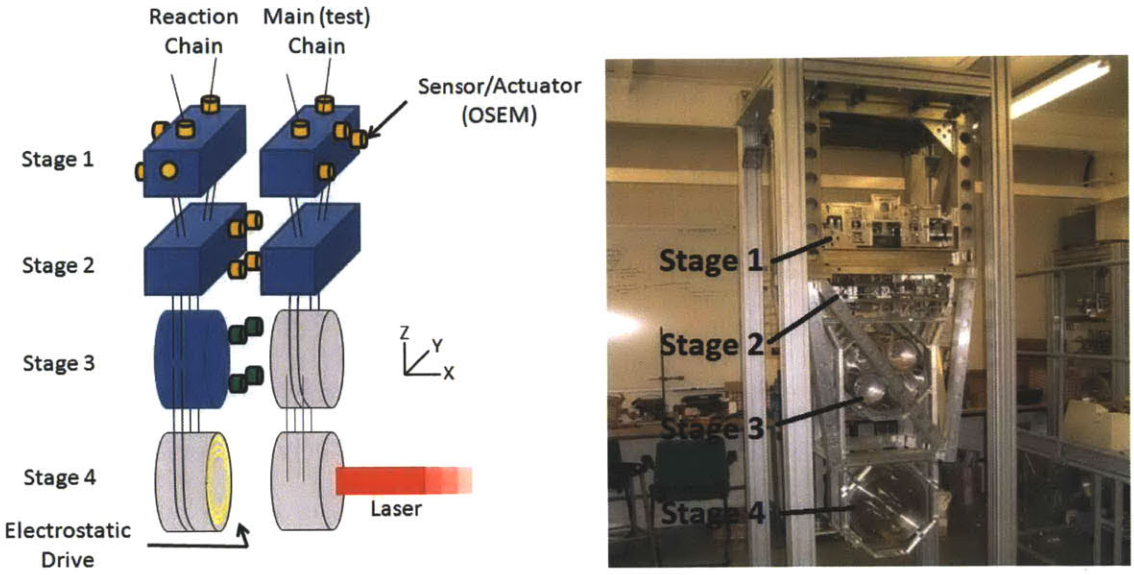

(a) Quadruple pendulum diagram

(b) Quadruple pendulum photo

Figure 2-2: (a): An illustration of the quadruple pendulum. It consists of two chains, a main chain (right) and a reaction chain (left). Stage 4 of the main chain is the interferometer optic. Three stages of cantilever springs provide vertical isolation. Sensor actuator devices (OSEMs) provide active damping and control in conjunction with an electrostatic drive (ESD). The reaction chain is used as a seismically isolated actuation surface. (b): A photograph of a prototype quadruple pendulum at the Rutherford Appleton Lab in the UK. Stages 3 and 4 are stainless steel dummy stages whereas the production versions are fused silica glass. Stages 1 and 2 are almost entirely covered by the surrounding cage. The cage's purpose is to mount sensors, actuators, and to catch the stages. Copyright Science and Technology Facilities Council.

characteristics. Since the reaction chain has less stringent noise requirements its stage 4 optic is supported by a loop of steel wire, which is mechanically more robust than glass. The glass stages of the pendulum are indicated in Figure 2-2a by a shade of gray and the metal stages are indicated by blue.

Each stage of the pendulum is modeled as a rigid body connected elastically by very lightly damped springs to the neighboring stages. Consequently, a second order differential equation is associated with each stage providing f^{-2} isolation above the pendulum's resonant frequencies, where f represents frequency. Thus, by using four stages a performance of f^{-8} is achieved. In this way the pendulum realizes six to seven orders of magnitude of seismic isolation in the single decade between its mechanical resonances and the low frequency end (10 Hz) of Advanced LIGO's GW

band. Figure 2-3 plots the modeled isolation performance of pendulums consisting of 1 through 4 stages, where the lowest stage is a reflective optic. The curves are the transfer functions between the pendulum's ground displacement and the displacement of the lowest stage along the interferometer axis. The quadruple and triple pendulum curves represent actual Advanced LIGO pendulum designs. The single and double are illustrative examples.

All active control is achieved with a network of sensors and actuators. Collocated sensor/actuator devices called OSEMs (Optical Sensor Electro-Magnet) are placed on Stages 1, 2 and 3, indicated by the yellow and green cylinders in Figure 2-2a. Fig. 2-4 includes an illustration of this device. Stage 4 contains actuators known as the electrostatic drive (ESD). Section 2.3 provides more detail on these devices.

The six OSEMs placed around each stage 1 are referenced to the pendulum's ground. They are used to damp the high quality factor mechanical resonances of the pendulum and provide low frequency test mass positioning control. Damping control is permitted only at stage 1 since the OSEM sensor noise of $10 \times 10^{-10}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ beyond 10 Hz [37] is non-negligible compared to the required Advanced LIGO sensitivity. As a result, the pendulum mechanically attenuates the sensor noise amplification through the pendulum chain below. A total of 22 out of 24 modes receive damping from these OSEMs. These modes are between 0.5 Hz and 5 Hz. They are designed to couple strongly to stage 1 to ensure observability and controllability for the damping control [34]. The remaining two modes occur at about 9 Hz and 13 Hz and couple poorly to stage 1. They largely represent vertical displacement and roll between stages 3 and 4 (extension of the glass fibers). Consequently, they are naturally well isolated from ground displacements and remain undamped.

The four actuators placed between the main and reaction chains at the second, third, and fourth stages control the x displacement and yaw and pitch rotation of the test mass. The pendulum is designed to split the test mass control between the various stages so that larger and noisier actuators are applied to the higher stages where there is better mechanical attenuation to the test mass. The primary concern is the coupling of OSEM magnets to environmental magnetic fields. The second and

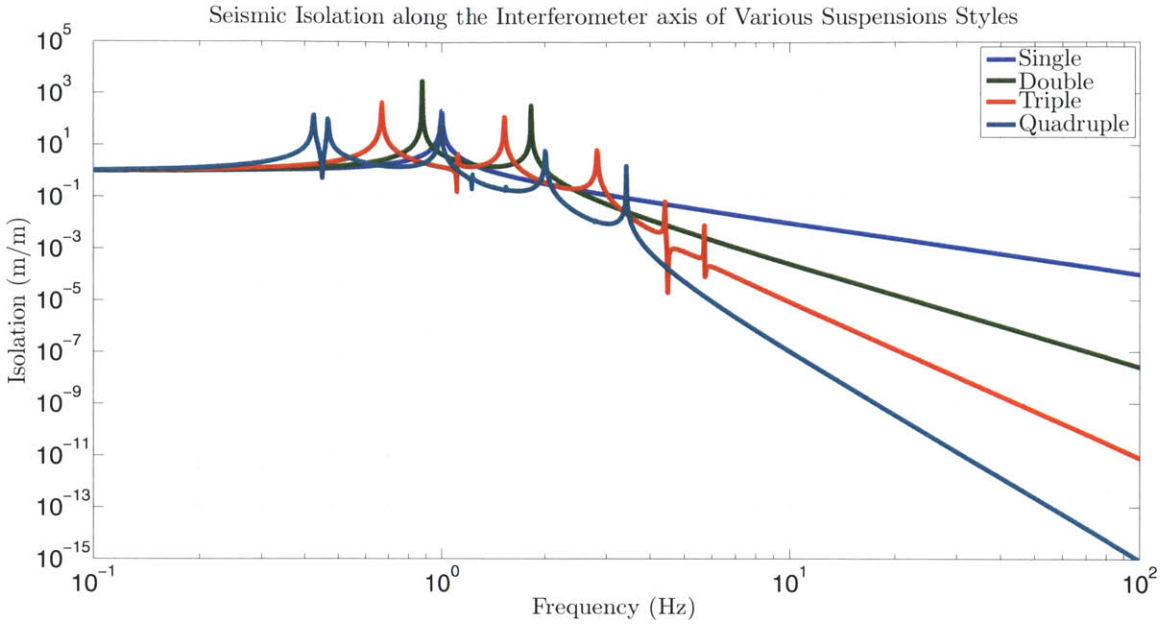

Figure 2-3: The isolation performance for pendulums consisting of 1 through 4 stages, where the lowest stage is an optic. Each curve is a simulated transfer function between ground motion and the displacement of the lowest stage of the pendulum along the interferometer axis. The triple and quadruple pendulum curves are simulations of actual Advanced LIGO pendulum designs. The single and double pendulum curves are illustrative examples.

third stages have four OSEMs each, and at stage 4 there are four ESD actuators. The error signal sent to each of these stages is the position of only the test mass measured directly from interferometric signals. The OSEM sensors here are used only for alignment purposes during the mechanical assembly of the pendulum.

2.3 Sensors, Actuators, and Electronics

The actuators discussed here are specifically those for prototype quadruple pendulum at the Massachusetts Institute of Technology (MIT). Their designs are similar to those used at the LIGO sites, but are not exactly the same.

2.3.1 Optical Sensor Electro-Magnet (OSEM)

The OSEMs on stage 1 provide sensing and actuation for the active damping control of the mechanical resonances. They are also used for low frequency control of the test mass. A detailed description of these OSEMs is found in [38].

These OSEM actuators contain an 8 mm long, 800 turn coil of wire that actuates on 10 mm long by 10 mm diameter NdBFe, nickel plated magnets attached to the top stage. This coil-magnet pair produces 2.05 N/A. The coil current is driven by a voltage source that drives current through the 30 Ω coil. This current driver includes a filter to reduce electronic noise from the digital-to-analogue converter (DAC) in the sensitive region of the LIGO interferometer at 10 Hz and above [9] (see Fig. 2-5). The current driver receives an input signal between ±10 V from the DAC, allowing the actuator to apply a maximum of 0.205 N at zero frequency.

The OSEMs on stage 2 are the same model as those on stage 1 and use the same magnets. However, the coil current driver has a different response due to the stricter noise requirements being closer to the test mass. The DC gain is about 60 times lower than the stage 1 driver. Also, this second stage filter has a steeper roll off [10] (see Fig. 2-6). The maximum DC output of these OSEMs is 3.4 mN.

The OSEMs on stage 3 have an even more strict noise limit being only one stage away from the test mass. Consequently, these OSEMs are a different model, with their own style of magnet. They have a 5.25 mm, 400 turn coil of wire, and actuate on a 6 mm long, 2 mm diameter magnet. The smaller coil and magnet pair makes these actuators much weaker than those on the higher stages, putting out only 16 mN/A, at the benefit of much lower noise. Due to the actuator's relative weakness, the stage 3 current driver filter does not need as steep a slope as those at stage 2. See Fig. 2-7.

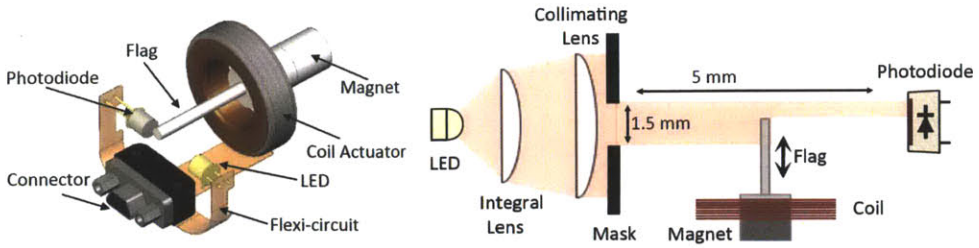

Figure 2-4: A drawing of the working parts of an OSEM. The basic OSEM components consist of an LED, photodiode, and coil of wire. A flag mounted to a stage on the quadruple pendulum blocks part of the LED light and produces a position dependent signal from the photodiode. When a current is run through the coil an actuation force is produced on a permanent magnet mounted under the flag. Adapted from [8].

These OSEMs will put out a maximum DC force of $43\,\mu$N [11].

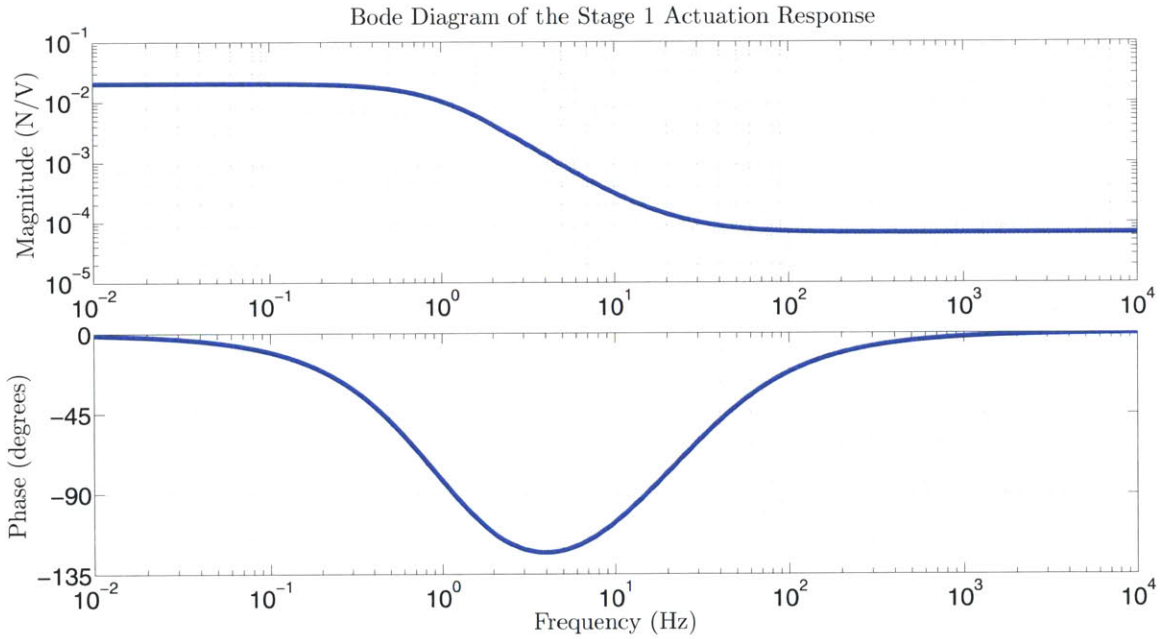

Figure 2-5: The modeled frequency response of the stage 1 actuation including the current driver and a coil-magnet pair. The current driver has two poles at 1 Hz, and zeros at 10 Hz and 31 Hz [9].

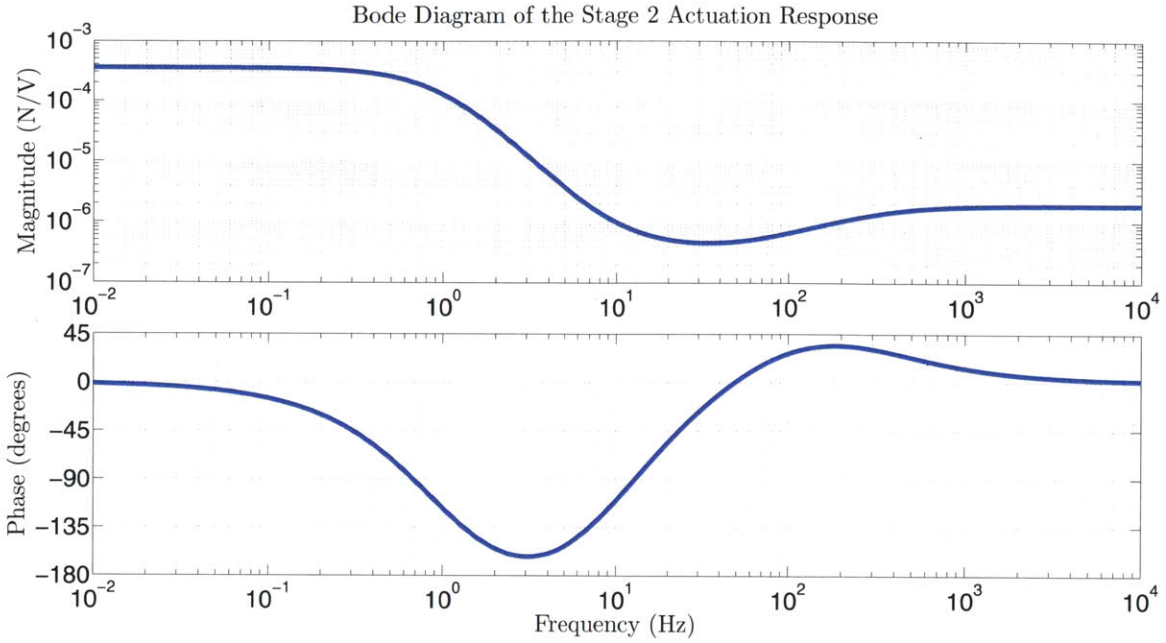

Figure 2-6: The modeled frequency response of the stage 2 actuation including the current driver and a coil-magnet pair. The current driver has three poles at 1 Hz, one pole at 325 Hz, three zeros at 10 Hz, and one zero at 60 Hz [10].

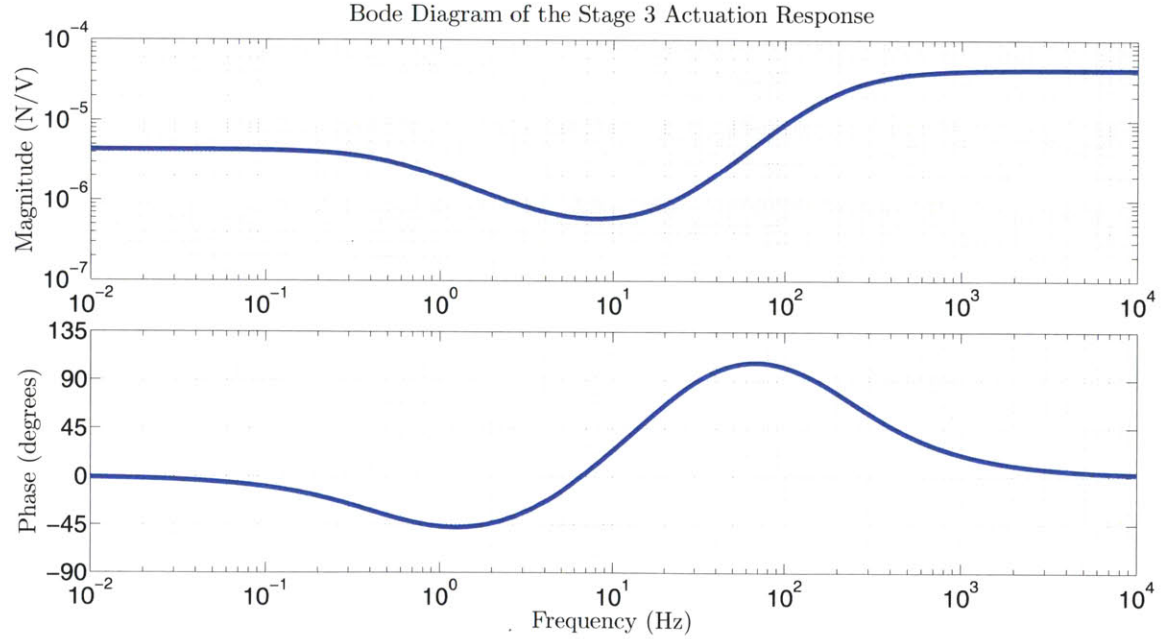

Figure 2-7: The modeled frequency response of the stage 3 actuation including the current driver and a coil-magnet pair. The current driver has poles at 0.5 Hz, two at 200 Hz and zeros at 5 Hz and two at 20 Hz [11].

2.3.2 Electrostatic Drive (ESD)

The strictest noise requirements are placed on the actuator at the test mass itself. Consequently, no magnets are permitted on it to avoid all possible coupling with stray magnetic fields. Physical contact of actuator parts with the test mass also increases mechanical loss and therefore creates suboptimal thermally excited fluctuations. These issues motivate the decision to use the ESD to actuate the test mass.

The ESD is described in detail in [12]. Fig. 2-8 shows a photograph of the reaction chain's stage 4 highlighting the ESD. It has four quadrants, each with a pair of interlaced conductive gold traces on the reaction chain. When a voltage difference is applied between these electrode pairs the resulting electric field attracts the nearby parallel dielectric surface of the glass test mass in front of it. Figure 2-9 illustrates this principle.

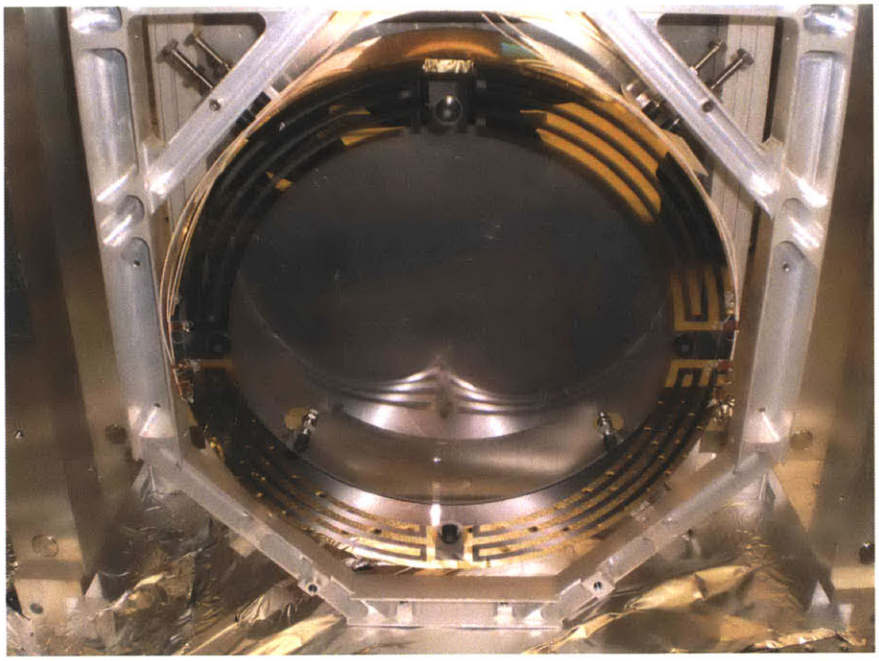

Figure 2-8: A photograph of a prototype quadruple pendulum electrostatic drive (ESD) on the reaction chain bottom stage. Each quadrant has two interlaced gold traces. A potential difference between the traces will apply a force on the dielectric surface of the nearby test mass.

The force each quadrant produces is nonlinear, given by

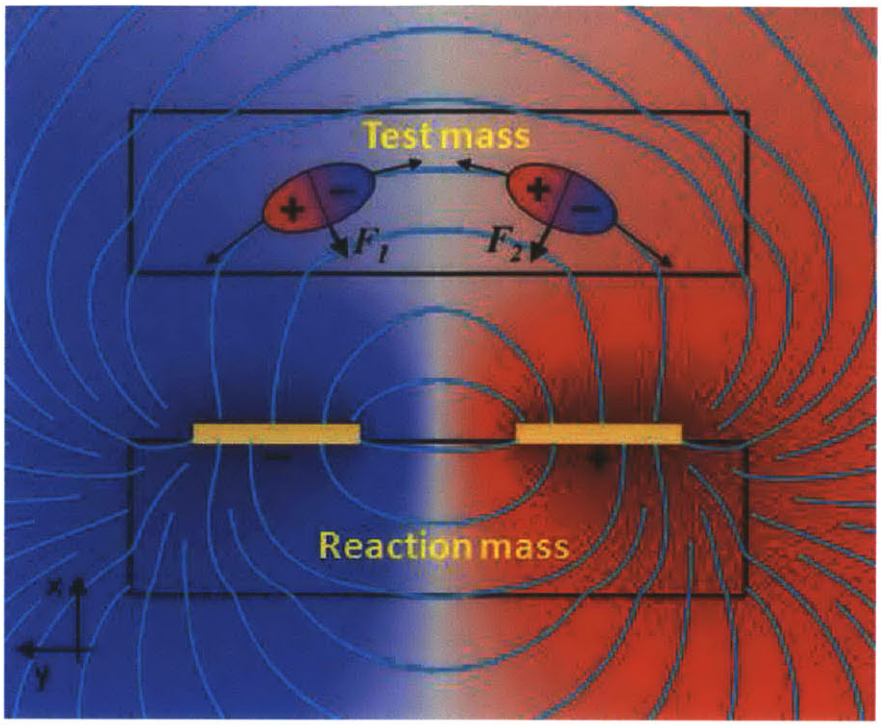

Figure 2-9: A diagram illustrating the working principle of the ESD. The upper rectangle represents the test mass containing two polarized molecules; the lower rectangle represents the reaction mass bearing two electrodes. The electric field lines are shown in cyan [12].

$$F = \alpha(\Delta V)^2. \tag{2.1}$$

The force is proportional to the square of the voltage difference between the two gold traces and a coefficient α, known as the coupling coefficient. α is dependent not just on the geometry of the gold pattern but also on the gap size between each stage 4. This dependence is modeled in Figure 2-10.

Eq. (2.1) must be linearized if linear control analysis is to be used. The voltage is linearized digitally by providing the square root of the desired control force to the DAC. The dependence of α on the gap size cannot be removed, however, since the gap is not measured in this system. Nonetheless, the gap size may be considered constant since the relative motion between the test mass and ESD is small compared to the gap size, on the order of 10^{-7} m RMS.

The speed of the actuator's response, many kHz, is much greater than the bandwidth of any reasonable design of an active control loop. As a result no time delays are included in the model of (2.1).

The maximum ΔV for this ESD is 800V and α is $2.9 \times 10^{-10}\,\mathrm{NV}^{-2}$ for a 5 mm gap. These values provide a maximum applied force of 190 μN with all four quadrants engaged. However, since the applied force is always attractive, bipolar actuation is achieved by implementing a static force offset of half the maximum. This necessary offset limits our effective actuation to 95 μN.

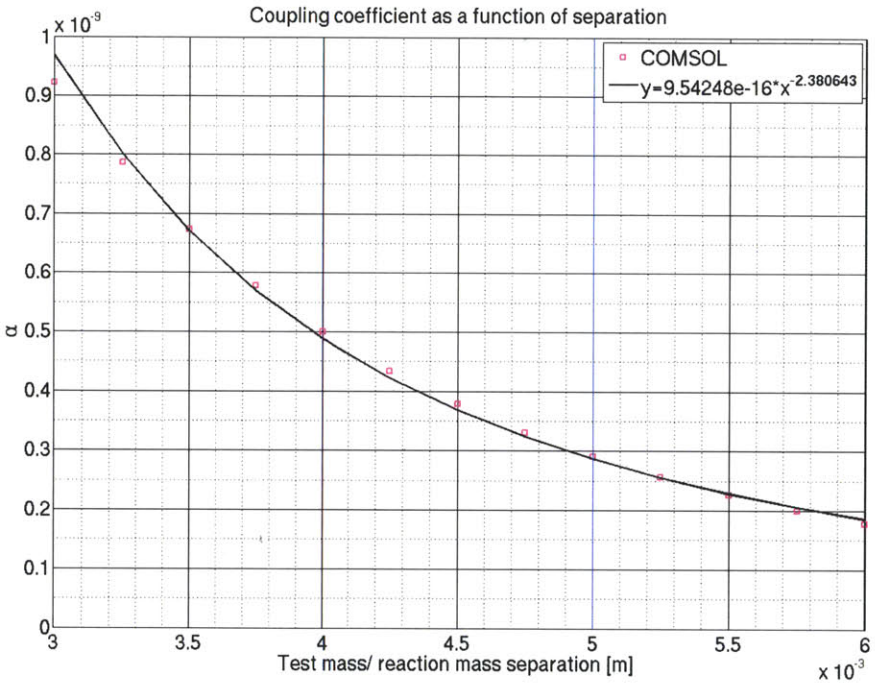

Figure 2-10: The modeled dependence of the ESD coupling coefficient α on the gap between the test mass and stage 4 of the reaction chain [12].

Chapter 3

Quadruple Pendulum Model and System Identification

A good system model is essential to performance predictions and good control design. Such models must be verified with appropriate system identification. In this chapter, Section 3.1 describes the derivation of the quadruple pendulum model this thesis employs. Section 3.2 describes the system identification methods used to characterize the prototype quadruple pendulum at the Massachusetts Institute of Technology (MIT) and verify the model.

3.1 Model

The version of the quadruple pendulum model used here was written in Mathematica® by Mark Barton and builds off previous work done by others in the LIGO group including Calum Torrie and Ken Strain [39]. A summary of the evolutionary history of the model is given in [40].

The calculations done in Mathematica® are based on the method of normal mode analysis as it is described in "Classical Mechanics" by Herbert Goldstein [41]. For a detailed description of how the model was developed by Barton see [39]. This section provides a summary of the theory and implementation.

The normal mode analysis invokes energy methods to derive the system of dynamic

equations. For small perturbations around the equilibrium the system is adequately assumed to be linear. The process is as follows:

1. Express the potential energy, E_P, in terms of the system's coordinates, $\mathbf{x}$.

$$E_P = E_P(x_1, ..., x_n) = E_P(\mathbf{x}) \tag{3.1}$$

2. Express the kinetic energy, E_K in terms of the system coordinates and coordinate velocities.

$$E_K = E_K(x_1, ..., x_n, \dot{x_1}, ..., \dot{x_n}) = E_K(\mathbf{x}, \dot{\mathbf{x}}) \tag{3.2}$$

   For a linear system the kinetic energy depends only on the coordinate velocities.

3. Minimize the potential energy to solve for the equilibrium coordinate vector $\mathbf{x}_{eq}$.

$$\mathbf{x}_{eq} = [x_{1(eq)}, ..., x_{n(eq)}]^T \tag{3.3}$$

4. Differentiate the potential energy twice to find its Hessian matrix. Since the dynamics are linear, this Hessian matrix is the symmetric stiffness matrix $\mathbf{K}$, with matrix elements K_{ij}.

$$K_{ij} = \frac{\partial^2 E_P}{\partial x_i \partial x_j} \tag{3.4}$$

   In fact, for a general nonlinear system the potential energy can be expanded with a Taylor series as

$$E_P(\mathbf{x}) \approx E_P(\mathbf{x}_{eq}) + \left[\frac{\partial E_P}{\partial x_i}\right]_{\mathbf{x}_{eq}} (\mathbf{x} - \mathbf{x}_{eq}) \tag{3.5}$$

$$+\frac{1}{2}(\mathbf{x} - \mathbf{x}_{eq})^T \left[\frac{\partial^2 E_P}{\partial x_i \partial x_j}\right]_{\mathbf{x}_{eq}} (\mathbf{x} - \mathbf{x}_{eq}) \tag{3.6}$$

   The first term represents the potential energy of the equilibrium position. By

shifting this arbitrary zero point, this term can be eliminated. The second term represents the gradient, which is also zero at the equilibrium position. For a linear system there are no terms beyond the third. Consequently, the potential energy can more simply be written as

$$E_P = \frac{1}{2}(\mathbf{x} - \mathbf{x}_{eq})^T \mathbf{K}(\mathbf{x} - \mathbf{x}_{eq}) \tag{3.7}$$

5. A similar expansion can be written for the kinetic energy. As a result, the mass matrix $\mathbf{M}$, with matrix elements M_{ij} is expressed as the Hessian matrix of the kinetic energy.

$$M_{ij} = \frac{\partial^2 E_K}{\partial \dot{x}_i \partial \dot{x}_j} \tag{3.8}$$

6. The resulting equations of motion are

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{K}(\mathbf{x} - \mathbf{x}_g) = \mathbf{f} \tag{3.9}$$

$\mathbf{x}_g$ represents the 'ground' surface to which the pendulum is attached. $\mathbf{f}$ represents external forces.

Historically, a MATLAB® model of a simplified version of the same system was developed and remains convenient for some purposes, especially for use with Simulink®. The Mathematica® model was used to generate new, more general core code for the MATLAB® model by exporting symbolic expressions for key matrices in MATLAB® syntax. The Mathematica® model also exports a parameter file, again in MATLAB® syntax (or one can easily be created by hand). This file gives numeric values for all the physical parameters, e.g., masses, stiffnesses, gravity, numbers of wires, and others. A wrapper script in MATLAB® then creates a MATLAB® numerical state-space system from the numerical parameters and symbolic matrices.

3.2 System Identification

The goal of system identification here is to measure as much about the pendulum's dynamics as possible in order to predict control performance, mechanical cross-couplings between DOFs and the values of the as-built parameters. Due to the relatively large number of physical parameters with uncertainty, the more data collected the better the chance of finding realistic values for these parameters. The data comes in the form of measured resonant frequencies and transfer functions.

3.2.1 Resonant Frequency Measurements

Resonant frequency measurements have the advantage that good sensor alignment and calibration is not needed. Further, the resonant frequencies contain the vast majority of the information needed to identify the model parameters. The disadvantage is that they provide little information about cross-couplings. In general they can be measured for four different, undamped, pendulum configurations to collect as much information as possible (see the illustration in Figure 3-1):

1. Free quadruple pendulum
2. Triple pendulum hanging from a locked stage 1
3. Double pendulum hanging from a locked stage 2
4. Single pendulum hanging from a locked stage 3

Locking various stages down changes the resonant frequencies of the free stages and localizes the effect of the uncertain parameters relating to those stages. These four configurations expand the number of resonances from 24 free pendulum resonances to a total of 60 (24 free + 18 triple + 12 double + 6 single). In theory even more configurations could be collected from free stages above locked stages, e.g. stage 1 free with stage 2 locked. These other configuration are not practical however since resonances are extremely sensitive to wire tension and locking a stage down undoubtedly changes wire tension above due to shifts in the position of the masses.

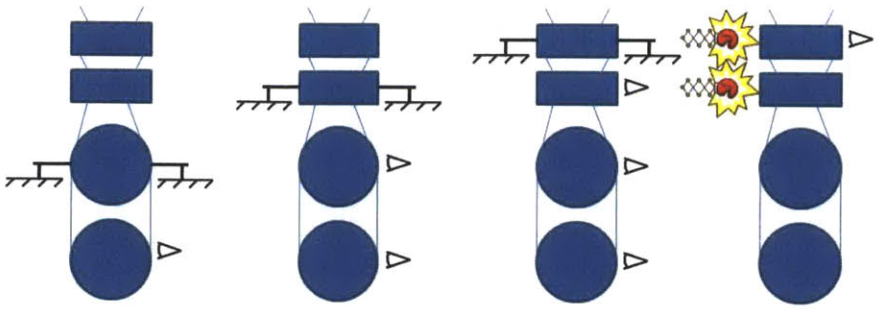

Figure 3-1: The four system identification configurations of the quadruple pendulum from which measurements are taken. The configuration on the left shows measurements being taken on the single pendulum of stage 4 while stage 3 is locked. The next shows measurements on the double pendulum of the lower two stages, then the triple pendulum of the lower three stages, and finally the full pendulum. The sketched in eye balls indicate where measurements are taken. The fists indicate where excitations are applied with the OSEMs to measure transfer functions.

The free pendulum resonant frequencies are extracted from stage 1 to stage 1 transfer functions (these serve Section 3.2.2 as well). Measurements on the other three configurations are taken with the OSEMs on stage 2 and stage 3 and an optical lever on stage 4. These triple, double, and single pendulum resonant frequencies are extracted from power spectra where the excitation is simply from background air currents and seismic noise (more than enough).

Nominally, the OSEMs on stages 2 and 3 measure only x, yaw, and pitch. However, intentional temporary misalignments of the OSEMs allow cross-coupling to y, z, and roll due to OSEM flags that only partially occlude the LED beams. In this way resonances from all six DOFs are visible to these OSEMs. The reaction chain is locked down when using stage 2 and stage 3 OSEMs to ensure that all measurements are of the main chain.

The optical lever measuring stage 4 is limited to measurements of only pitch and yaw. The x resonances were captured as well through mechanical x-pitch coupling, but the others were not visible.

For the data collected in this thesis a total of 50 out of the possible 60 resonant frequencies were measured on the quadruple pendulum with a resolution at least as fine as 2 mHz. The highest frequency roll mode of the triple suspension configuration was not visible to either the OSEMs or the optical lever. In the double configuration

only the optical lever was used since it was not realized that the stage 3 OSEMs could provide off-axis measurements until it was too late. Hence only x, pitch, and yaw modes were obtained missing 6 of the possible 12. For the single pendulum case, the optical lever was the only option, so again x, pitch, and yaw were the only visible DOFs, obtaining 3 out of a possible 6. At a later time it could not be proven that the single pendulum test was not influenced by the compliance of the Viton® rubber padding in the stops locking the glass stage 3. Thus only 47 out of a possible 60 frequencies were retained for the parameter estimation problem in Chapter 4.

3.2.2 Transfer Function Measurements

Transfer functions have the advantage that they contain cross-coupling information. The disadvantage is that the OSEMs must be well aligned and have sufficient dynamic range.

As stated earlier, transfer functions are measured from stage 1 to stage 1 for the free pendulum configuration. This yields a six-by-six transfer function matrix which contains the drive from every stage 1 DOF to every stage 1 DOF between 0.1 Hz and 10 Hz. These stage 1 transfer functions were measured at MIT with a resolution no more coarse than 2 mHz to accurately resolve the high Q factor resonance peaks.

Transfer functions were also measured from the stage 2 x, yaw, and pitch DOFs to all six stage 1 DOFs for the free pendulum configuration. Simultaneous measurements to stage 2 and stage 3 would be ideal, but the alignment of the sensors was not trusted. Since the free pendulum resonances were already localized from the stage 1 transfer functions, a more coarse 5 mHz resolution was used here. No transfer functions were measured with excitations from stage 3 or stage 4. Both stage 3 and stage 4 have insufficient dynamic range.

Chapter 4

Parameter Estimation Problem

The quadruple pendulums require precise models to adequately predict their performance and commission their control systems. Current models of these quadruple pendulums contain 67 unique physical parameters to fully characterize the system with the desired complexity. These parameters contain information related to the inertia, stiffness, and geometry of the pendulum's components. All these parameters contain some bounded a priori uncertainty, of which some may be significant while others are not.

This chapter outlines an approach to fit the model to the measured data for the quadruple pendulum prototype at MIT. Section 4.1 develops a procedure to identify the parameters with significant a priori uncertainties. The measurements important to estimating the true values of these parameters are also identified with this procedure. Section 4.2 describes the Gauss-Newton algorithm chosen to perform this estimation with the selected parameters and measurements. Section 4.3 presents the experimental results of this model fitting approach.

4.1 Sensitivity Selection Procedure

Before a model can be fit to experimental data the relevant parameters and measurements must be identified. This section develops a sensitivity analysis procedure to identify and rank those parameters with a priori uncertainties that are most critical to

the system's performance. The procedure then selects the measurements that couple most strongly to this parameter set. Through this procedure classes of uncertainties and measurements are identified that dominate the observed error between the model and the physical system. The procedure benefits the problem of off-line parameter estimation in two ways. First, by minimizing the number of parameters and measurements required and second, by maximizing the conditioning of the estimation problem. This latter point increases the likelihood of finding a global minimum.

A key tool for this procedure is the sensitivity matrix. A system's sensitivity matrix is rich with information. Fundamentally, it tells us how the model parameters influence the measurements, which is the interpretation followed in this section. However, it is also the slope at a local point in the model's parameter error space, which means it can calculate descent directions for parameter estimation algorithms. Another use is calculating the Fisher Information Matrix or its inverse, the Cramér-Rao lower bound. These matrices put lower limits on how well unknown parameters can be estimated in the presence of noisy measurements. This chapter eventually exhausts all of these uses. The sensitivity matrix is derived below in a way specifically useful for the purposes of this chapter.

4.1.1 Analysis

A general system has the form

$$\mathbf{y} = f(\boldsymbol{\theta}, \boldsymbol{\mu}) + \boldsymbol{\nu} \tag{4.1}$$

$$\boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\epsilon} \tag{4.2}$$

$$\boldsymbol{\epsilon} \in [-\bar{\boldsymbol{\theta}}; \bar{\boldsymbol{\theta}}] \tag{4.3}$$

$\mathbf{y}$ is the m length column vector of measured system outputs, $\boldsymbol{\theta}$ is the n length vector of system parameters, $\boldsymbol{\mu}$ is the vector of system inputs, and $\boldsymbol{\nu}$ is the measurement noise. The system parameters are assumed to be known within some deviation, $\boldsymbol{\epsilon}$, of the designed values $\boldsymbol{\theta}_0$. The unknown value $\boldsymbol{\epsilon}$ is known to be bounded by $\pm\bar{\boldsymbol{\theta}}$.

The system is modeled by Eq. (4.4) where $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\epsilon}}$ are the estimated parameter values and parameter deviations respectively. $\hat{\mathbf{y}}$ is the output of the model $\hat{f}$.

$$\hat{\mathbf{y}} = \hat{f}(\hat{\boldsymbol{\theta}}, \boldsymbol{\mu}) \tag{4.4}$$

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}_0 + \hat{\boldsymbol{\epsilon}} \tag{4.5}$$

Generally, the sensitivity matrix is derived from the squared model-observation error in Eq. (4.6)

$$V(\hat{\boldsymbol{\epsilon}}) = \frac{1}{2}(\mathbf{y} - \hat{\mathbf{y}}(\hat{\boldsymbol{\epsilon}}))^T(\mathbf{y} - \hat{\mathbf{y}}(\hat{\boldsymbol{\epsilon}})) \tag{4.6}$$

$$\mathbf{g}(\hat{\boldsymbol{\epsilon}}) = \mathbf{J}(\hat{\boldsymbol{\epsilon}})^T(\mathbf{y} - \hat{\mathbf{y}}(\hat{\boldsymbol{\epsilon}})) \tag{4.7}$$

where V is the scalar valued square error. Deriving with respect to $\hat{\boldsymbol{\epsilon}}$ yields Eq. (4.7) where $\mathbf{g}$ is the n length column gradient vector and $\mathbf{J}$ is the $m \times n$ Jacobian matrix. The Jacobian matrix is defined in terms of partial derivatives in Eq. (4.8).

$$J_{ij} = -\frac{\delta \hat{y}_i}{\delta \hat{\epsilon}_j} \tag{4.8}$$

$$1 < i < m \tag{4.9}$$

$$1 < j < n \tag{4.10}$$

A common way to make the Jacobian matrix dimensionless is to normalize each element with the values of the parameters and measurements. In this way, all sensitivities are comparable [42, 43]. This dimensionless matrix $\mathbf{S}$ is referred to here as the sensitivity matrix.

$$S_{ij} = \frac{\hat{\theta}_j}{y_i} J_{ij} \tag{4.11}$$

There are two reasons this form of the sensitivity matrix does not suit our needs. First, certain parameters such as geometric values can be defined from arbitrary reference points. For example, if θ_j refers to the horizontal position of one end of a suspension wire, the value could arbitrarily be shifted to zero simply by changing the reference point. This shift would not influence the pendulum's dynamics and would unfairly yield a sensitivity value of zero. Second, our goal is to characterize the importance of parameter uncertainty. The absolute value of a parameter is not important.

Modifying Eq. (4.11) by scaling the Jacobian with the a priori parameter uncertainties $\bar{\boldsymbol{\theta}}$ rather than the parameter values themselves yields a new dimensionless sensitivity matrix $\bar{\mathbf{S}}$ that scales with this uncertainty.

$$\bar{S}_{ij} = \frac{\bar{\theta}_j}{y_i} J_{ij} \tag{4.12}$$

Using singular value decomposition, this sensitivity matrix is decomposed into three multiplicative matrices.

$$\bar{\mathbf{S}} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T \tag{4.13}$$

$\mathbf{U}$ is the $m \times m$ matrix of orthogonal, unit length, left singular vectors. This matrix characterizes the relative importance of the system measurements. $\mathbf{V}$ is the $n \times n$ matrix of right orthogonal, unit length, singular vectors, and characterizes the relative importance of the parameter uncertainties. In general, the transpose on $\mathbf{V}$ is a complex conjugate transpose, however for this application the matrices are real. $\boldsymbol{\Sigma}$ is the $m \times n$ diagonal matrix of singular values σ_k listed in descending order where k runs from 1 to n. The singular values quantify how much small parameter changes in the directions of the right singular vectors influence the measurements in the directions of the left singular vectors.

These matrices can be used to select the parameters and measurements that are

most sensitive to the system's dynamics. A similar analysis is performed using principal component analysis (pca) in [42, 43]. However, this pca analysis is limited to $\mathbf{\Sigma}$ and $\mathbf{V}$, ignoring the information in $\mathbf{U}$. As a consequence, the pca analysis contains limited information about the measurements. Here we introduce a method to rank the importance of both the model parameters and the measurements.

The matrix norm $\| \bar{\mathbf{S}} \|$ is defined as the length of all the singular values such that

$$\| \bar{\mathbf{S}} \| = \sqrt{\sum_{k=1}^{n} \sigma_k^2} \tag{4.14}$$

The relative sensitivity magnitude of the uncertainty for each physical model parameter $\hat{\theta}_j$ can be quantified as p_j in the following way

$$p_j = \| \bar{\mathbf{S}} \|^{-1} \sqrt{\mathbf{v}_j \mathbf{\Sigma}^T \mathbf{\Sigma} \mathbf{v}_j^T} \tag{4.15}$$

$$1 < j < n \tag{4.16}$$

where $\mathbf{v}_j$ is the j^{th} row vector of $\mathbf{V}$. The vector $\mathbf{p}$, of which p_j are the components, is unit length. The closer p_j is to 1, the more influence $\bar{\theta}_j$ has on the measurements relative to the other parameter uncertainties.

Equivalently, p_j can be determined by calculating the 2-norm of the j^{th} column of $\bar{\mathbf{S}}$ and scaling it by $\| \bar{\mathbf{S}} \|^{-1}$ as given by Eq. (4.17).

$$p_j = \| \bar{\mathbf{S}} \|^{-1} \sqrt{\sum_{a=1}^{m} \bar{S}_{aj}^2} \tag{4.17}$$

In a similar way, the sensitivity magnitude for each measurement y_i can be quantified as q_i

$$q_i = \| \bar{\mathbf{S}} \|^{-1} \sqrt{\mathbf{u}_i \boldsymbol{\Sigma} \boldsymbol{\Sigma}^T \mathbf{u}_i^T} \tag{4.18}$$

$$1 < i < m \tag{4.19}$$

where $\mathbf{u}_i$ is the i^{th} row vector of $\mathbf{U}$. The vector $\mathbf{q}$, of which q_i are the components, is unit length. The closer q_i is to 1, the more information y_i has on the parameter set $\boldsymbol{\theta}$ relative to the other measurements.

Equivalently, q_i can be determined by calculating the 2-norm of the i^{th} row of $\bar{\mathbf{S}}$ and scaling it by $\| \bar{\mathbf{S}} \|^{-1}$ as given by Eq. (4.20).

$$q_i = \| \bar{\mathbf{S}} \|^{-1} \sqrt{\sum_{b=1}^{n} \bar{S}_{ib}^2} \tag{4.20}$$

4.1.2 Procedure

The procedure for using the analysis in Section 4.1.1 to identify the important uncertainties and associated measurements is described here and then summarized in the following enumerated list.

First the sensitivity matrix $\bar{\mathbf{S}}$ is calculated. The parameter uncertainty vector $\mathbf{p}$ is then found. This vector is sorted in descending order and renamed $\tilde{\mathbf{p}}$.

Selecting the $n^* \leq n$ most sensitive parameters is now a straight forward procedure. A desired sensitivity level η_p is chosen where $0 < \eta_p \leq 1$. The n^* parameters are chosen such that n^* is just large enough that the squared length of the first n^* elements of $\tilde{\mathbf{p}}$ is greater than or equal to η_p. The vector of the n^* most sensitive parameters is referred to as $\tilde{\mathbf{p}}^*$, where $\eta_p \leq \| \tilde{\mathbf{p}}^* \|^2 \leq 1$. The sensitivity matrix $\bar{\mathbf{S}}$ is then recalculated with the reduced parameter set. This recalculation can be as simple as removing the columns associated with the rejected parameters.

Next, the measurement sensitivity vector $\mathbf{q}$ is calculated. This vector is then sorted in descending order and renamed $\tilde{\mathbf{q}}$. The procedure for identifying the m^* most

sensitive measurements is largely the same, where $m^* \leq m$. A desired sensitivity level η_m is chosen where $0 < \eta_m \leq 1$. The m^* measurements are chosen such that m^* is just large enough that the squared length of the first m^* elements of $\tilde{\mathbf{q}}$ is greater than or equal to η_q. The vector of the m^* most sensitive measurements is referred to as $\tilde{\mathbf{q}}^*$, where $\eta_q \leq \| \tilde{\mathbf{q}}^* \|^2 \leq 1$. Ideally one would choose η_q such that $m^* \geq n^*$. This choice would satisfy a necessary, though not sufficient, condition needed for a parameter estimation algorithm to converge to the true parameters. In the general case however, this necessary condition may not be achievable.

The following list summarizes the steps of the selection procedure.

1. Identify the a priori uncertainty bounds on all model parameters.
2. Establish the list of feasible measurements.
3. Calculate the sensitivity matrix $\bar{\mathbf{S}}$ between the measurements and parameters.
4. Determine the parameter uncertainty weights $\mathbf{p}$ from $\bar{\mathbf{S}}$.
5. Generate $\tilde{\mathbf{p}}$ by sorting $\mathbf{p}$ in descending order.
6. Determine the desired quantity of sensitivity to preserve by choosing η_p and select the n^* most sensitive parameters such that $\eta_p \leq \| \tilde{\mathbf{p}}^* \|^2 \leq 1$.
7. Recalculate $\bar{\mathbf{S}}$ with the reduced parameter set.
8. Determine the measurement weights $\mathbf{q}$.
9. Generate $\tilde{\mathbf{q}}$ by sorting $\mathbf{q}$ in descending order.
10. Determine the desired quantity of sensitivity to preserve by choosing η_q and select the m^* most sensitive measurements such that $\eta_q \leq \| \tilde{\mathbf{q}}^* \|^2 \leq 1$.

4.2 Parameter Estimation Method

This section presents the parameter estimation routine written around Mark Barton's MATLAB® model that varies the selected physical parameters in order to minimize

the error between the model and the selected data. The chosen routine is Gauss-Newton, a common approximation to Newton's Method for least squares applications [44]. The MATLAB® code is included in reference [45].

4.2.1 Newton's Method

Newton's method is a second order method that uses both the first and second derivatives to quickly find the minimum of a function. A function $f(\vec{x})$ is approximated with a second order Taylor series as

$$V(\boldsymbol{\theta}_{k+1}) \approx f(\boldsymbol{\theta}_k) + \mathbf{g}_k^T(\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k) + \frac{1}{2}(\boldsymbol{\theta}_{\theta+1} - \boldsymbol{\theta}_k)^T\mathbf{H}_k(\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k) \tag{4.21}$$

In this application the scalar V would be the sum of the squared error between the model and the measured data. The dependent variable $\boldsymbol{\theta}$ is a column vector of length n, and represents the parameters we wish to fit. $\mathbf{g}$ is the column vector gradient, length n, of V with respect to $\boldsymbol{\theta}$. $\mathbf{H}$ is the n by n Hessian matrix of V. The Hessian is the symmetric matrix of second derivatives with respect to $\boldsymbol{\theta}$. The subscript k represents the current iteration number of the algorithm.

Similarly, the gradient can be approximated with a first order Taylor series.

$$\nabla V_{k+1} = \mathbf{g}_{k+1}^T \approx \mathbf{g}_k^T + (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k)^T\mathbf{H}_k \tag{4.22}$$

An approximate solution for the minimum can be found at $\boldsymbol{\theta}_{k+1}$ by setting the gradient in Eq. 4.22 to zero.

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \mathbf{H}_k^{-1}\mathbf{g}_k \tag{4.23}$$

Iterating over this calculation will typically converge to a minimum in relatively few steps. It does not guarantee that the minimum is global.

4.2.2 Gauss-Newton Algorithm

One of the major difficulties with Newton's method is the calculation of $\mathbf{H}^{-1}$. If a single calculation is needed to find the error, than n calculations are needed to find the gradient, and n^2 calculations are needed to find the Hessian. Then the Hessian needs to be inverted. Thus, computation quickly gets very slow with the square of the number of parameters being solved. In this application each calculation includes compiling the model. Thus, if we have ten parameters we want to fit, we need to compile the model at least 100 times just to get the Hessian matrix for a single iteration of the routine. The Gauss-Newton method is an approximation to this method, among others, that sacrifices some accuracy on the Hessian matrix for an improvement on computation time.

The algorithm depends on knowledge of the gradient vector and an approximate Hessian matrix to run. One way to easily estimate the gradient vector is to use the finite difference method. Each single parameter is stepped by a small amount, the model recompiled, and the change in error calculated. The list generated for each parameter is the gradient vector. Care must be taken to use an appropriate step size. Too small a step and roundoff errors dominate, too large and the estimate is not a faithful representation of the gradient at the current point.

The Hessian approximation can be derived by writing the cost V in (4.21) exactly as a function of the model residuals, the measured difference between the physical pendulum and the modeled pendulum. In (4.24) $\mathbf{r}$ represents the length m column vector of residuals ($\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$).

$$V(\boldsymbol{\theta}) = \frac{1}{2}\mathbf{r}(\boldsymbol{\theta})^T\mathbf{r}(\boldsymbol{\theta}) \tag{4.24}$$

Taking the derivative with respect to the model parameters, $\boldsymbol{\theta}$, gives the gradient vector $\mathbf{g}$.

$$\mathbf{g}(\boldsymbol{\theta}) = \mathbf{J}(\boldsymbol{\theta})^T\mathbf{r}(\boldsymbol{\theta}) \tag{4.25}$$

where $\mathbf{J}$ is the Jacobian matrix of the residuals with respect to the parameters.

$$J_{ij} = \frac{\delta r_i}{\delta \theta_j} \tag{4.26}$$

$$1 < i < m \tag{4.27}$$

$$1 < j < n \tag{4.28}$$

Note that we have derived a second way to calculate the gradient vector. The Jacobian is calculated using finite differences in a manner similar to that described for the first gradient estimation method. In practice whichever gradient calculation method is the most convenient at a particular time is the one used.

The Hessian matrix is derived by taking another derivative with respect to the parameters.

$$\mathbf{H}(\boldsymbol{\theta}) = \mathbf{J}(\boldsymbol{\theta})^T \mathbf{J}(\boldsymbol{\theta}) + \mathbf{L}(\boldsymbol{\theta})\mathbf{r} \tag{4.29}$$

$\mathbf{L}$ is some matrix that includes many second derivatives. Since this matrix is multiplied by $\mathbf{r}$, in the neighborhood of the optimal solution, where $\mathbf{r} = 0$, the Hessian can be approximated using only first order derivatives from the Jacobian.

$$\mathbf{H}(\boldsymbol{\theta}) \approx \mathbf{J}(\boldsymbol{\theta})^T \mathbf{J}(\boldsymbol{\theta}) \tag{4.30}$$

The extent of the 'neighborhood' is not well defined a priori. It is possible for Gauss-Newton to fail to converge if the starting model parameters are too far from the true parameters. Based on experience the algorithm has no trouble converging to the quadruple pendulum model's true solution given the a priori parameter uncertainties.

A sufficient condition for convergence to a unique global solution is that $\mathbf{J}$ have full column rank everywhere in the region of exploration inside the parameter space. A necessary, but not sufficient, condition for full column rank is that $m \geq n$, i.e. at least as many measurements as there are parameters being fit. It is difficult to numerically guarantee that $\mathbf{J}$ has full column rank everywhere, and not obvious how

to show it analytically. However, confidence that this requirement is met can be greatly enhanced by using as many measurements as possible, particularly those that couple well to the parameters being tuned by the algorithm. The rank, or better the condition number (ratio of maximum and minimum singular values), of $\mathbf{J}$ can also be tested for various points in the parameter space, including those the algorithm passes through.

The Gauss-Newton routine used here runs through the following steps:

1. Start with an initial guess of the model parameters.
2. Calculate the Jacobian matrix $\mathbf{J}_k$.
3. Calculate the search direction $\mathbf{d}_k = -\mathbf{J}_k^{\dagger}\mathbf{r}_k$. For full column rank $\mathbf{J}_k$, $\mathbf{d}_k = -(\mathbf{J}_k^T\mathbf{J}_k)^{-1}\mathbf{g}_k$. '$\dagger$' indicates the Moore-Penrose pseudoinverse.
4. Do a line search in the direction of $\mathbf{d}_k$ to find an appropriate step size α_k so that $\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k + \alpha_k\mathbf{d}_k$.
5. If $|\mathbf{g}_k^T\mathbf{g}_k|$ is greater than some user chosen scalar value, return to step 2. Otherwise, stop.

4.2.3 Implementing Gauss-Newton on the Data

Residual Calculation

One key aspect of the implementation of any fitting routine is how the error is defined. There are many ways to define the error between the model and the data, each of which will yield different results; unless the uncertainties are small enough that the routine converges to exactly the as built parameters, but this is not realistic.

Since most of the data involves measured resonant frequencies, an obvious choice is to compare them to the resonances of the model. Then one still needs to decide how to compare the resonances. Insight comes from the fact that percent errors on physical parameters tend to cause percent errors on mode frequencies. For instance quadrupling the length of a single pendulum halves its frequency. Thus, a shift of

10 mHz at 0.5 Hz represents much more parameter error than 10 mHz at 5 Hz. So it makes sense to use some sort of percent error on the resonant frequencies, rather than absolute error.

Consequently, resonance residuals are calculated with the following equation, where ω represents frequency in radians/second. These units were chosen instead of Hz because these are the units MATLAB® returns, and dividing by 2π requires unnecessary time.

$$r_i = \sum_{i=1}^{m} \frac{\omega_{i,measured} - \omega_{i,model}}{\omega_{i,measured}} \tag{4.31}$$

The remaining data fed into the model comes from magnitudes of the measured transfer functions from 7 Hz to 10 Hz. The residual terms for this data are calculated as the average percent difference between the modeled and measured magnitudes.

Estimating the Expected Parameter Solution Error

Given imperfect measurements we can expect the algorithm to converge to a solution with some error. The minimum size of this error can be estimated using a mathematical device called the Fisher Information Matrix (FIM). A book entitled *Optimal Measurement Methods for Distributed Parameter System Identification* by Dariusz Uciński gives a good description of the FIM [46].

Conveniently, the FIM is largely dependent on the system Jacobian matrix $\mathbf{J}$, which is already calculated by the Gauss-Newton algorithm as described in Section 4.2.2. The only other information needed is the variance of the measurement. If $\mathbf{C}_\epsilon$ is the covariance matrix of the measurement error, then the covariance matrix of the parameter estimate error, $\mathbf{C}_\theta$, is given by the inverse of the Fisher Information Matrix $\mathbf{M}$,

$$\mathbf{C}_\theta = \mathbf{M}^{-1} \tag{4.32}$$

$$\mathbf{M} = \mathbf{J}^T \mathbf{C}_\epsilon^{-1} \mathbf{J} \tag{4.33}$$

or more directly,

$$\mathbf{C}_\theta = \mathbf{J}^\dagger \mathbf{C}_\epsilon \mathbf{J}^{\dagger T} \tag{4.34}$$

If it is desired to apply weighting factors to certain measurements, for example if some measurements are trusted more than others, then $\mathbf{C}_\theta$ is adjusted by

$$\mathbf{C}_\theta = \mathbf{J}^\dagger \mathbf{W} \mathbf{C}_\epsilon \mathbf{W} \mathbf{J}^{\dagger T} \tag{4.35}$$

Where $\mathbf{W}$ is a diagonal matrix that weights the relative importance of each measurement. For this application the transfer function data were scaled down by a factor of 10 relative to the resonances, otherwise they tend to dominate the total error while containing the greatest uncertainty. The standard deviation of the resonance measurement error was taken to be the measurement resolution, 2 mHz. The transfer function magnitudes were assumed to be accurate to 3%, but the error is possibly greater.

The minimum standard deviation of the parameter error is the square root of the diagonal elements of $\mathbf{C}_\theta$. It is important to realize that the FIM yields only the *minimum* parameter error. The actual error can be larger. Further, unknown errors in the model (e.g. those parameters that are assumed correct but not, or incorrect measurements) will increase the error. Further still, the FIM depends on the calculation of $\mathbf{J}$, which itself is an estimation based on the assumption of a good model.

Reducing the Model for the Various System Identification Configurations

It was necessary to adjust the model to reflect the various sub-pendulum system identification configurations. This is a very simple, almost trivial procedure. Locking a mass has the effect of forcing its position and velocity to zero. An easy way to reflect this in the model is to compile the model as usual and simply delete the relevant rows and columns in the resulting state space matrices. For example, when

stage 4 is free by itself all the rows and columns representing the top three stages are deleted. In this way the model could be compared to the measured resonances of these sub-pendulums as part of the fitting routine's error calculation.

4.3 Experimental Results

4.3.1 Selection Procedure

In this section the sensitivity selection procedure developed in Section 4.1 is applied to the prototype Advanced LIGO quadruple pendulum at MIT. The values of the selected parameter set are estimated with the selected measurements using the Gauss-Newton algorithm detailed in Section 4.2.

Following step 1, some of the pendulum's parameters are well known a priori because they are either directly measured or are confined by well known constraints. For example, the mass of each stage is measured to either 22 kg or 40 kg within about 5 g. The lengths of the suspension wires, about 0.5 m, are cut to a precision better than 1 mm. The location of the ends of these wires in the pendulum are assumed to have an uncertainty of 10 μm in the horizontal plane. Other parameters have more uncertainty. The ends of the same wires in the vertical plane are only known within about 1 mm because they are fixed to adjustable springs. The lowest two stages are connected directly by glass fibers rather than steel wires hung from springs. The effective bending point of these fibers cannot be directly measured. Consequently, a finite element model estimates this value with an assumed uncertainty of 2 mm. Rotational inertias associated with each stage are more difficult to determine as well. The initial guess of those values is typically calculated from CAD models. The actual value is assumed to be within 2% here. Spring stiffness uncertainty is also assumed at 2%. These are characteristic examples of a total of 67 unique parameters.

For step 2, the pendulum's performance is determined by its resonant frequencies. The transfer function magnitude measurements are left out here for simplicity. The resonances lie between 0.5 Hz and 10 Hz and by design have very low damping.

Consequently, they can easily be measured down to mHz precision or better given sufficient time. In general, there are 60 resonant frequencies available as described in Section 3.2.1.

In practice, all possible measurements are taken to check for unexpected errors. However, the ones where the most precision (and time and effort) is needed are the ones reported by the selection procedure. The prototype setup at MIT was limited in that 13 of these resonance measurements were not obtainable as described in Section 3.2.1. The production article pendulums at the Hanford and Livingston observatories will have the option to measure all 60 resonances.

Considering these measurements, the parameter uncertainty described above, and the model described in Section 3.1, the quadruple pendulum's 47×67 sensitivity matrix $\bar{\mathbf{S}}$ is constructed according to step 3. For steps 4 and 5 the 67 uncertainty weights $\mathbf{p}$ are calculated and ranked in descending order to generate the sorted list $\tilde{\mathbf{p}}$.

For step 6, $\eta_p = 0.97$ was chosen. That is, by this measure we choose to preserve 97% of the uncertainty sensitivity in the selected parameters. A trade-off exists in deciding the best value of η_p. The closer to 1, the more parameter sensitivity is preserved, but the more poorly scaled the parameter estimation is likely to be. 97% was chosen here because it represents the largest number of parameters with a well scaled Jacobian matrix. The Jacobian matrix determines the parameter estimation descent direction and must be sufficiently well scaled to converge to the global minimum of the model error. With an $\eta_p = 0.97$, 59 parameters are rejected leaving $\tilde{\mathbf{p}}^*$ with $n^* = 8$ components. The vector elements of $\tilde{\mathbf{p}}$ are plotted in Figure 4-1 as the solid line. The squared norm of this vector up to the j^{th} element is also plotted as the dashed line. After the first 8 most sensitive parameters (dotted line) this norm is greater than η_p.

The first two parameters in $\tilde{\mathbf{p}}^*$ represent the uncertainty in bending points of the glass fibers supporting stage 4. The second two parameters represent the uncertainty in the position of the ends of the vertical springs in stages 1 and 2. The fifth parameter represents the vertical stiffness along the length of the glass fibers. The next two represent the stiffness of the vertical springs in stages 1 and 2. The last parameter

represents the rotational inertia about the stage 2 y (pitch) axis (refer to Figure 2-2).

In steps 7, 8, and 9 the columns of $\bar{\mathbf{S}}$ relating to the 59 rejected parameters are removed and the descending measurement weights $\tilde{\mathbf{q}}$ are calculated.

At step 10 it is decided to target η_q at 0.99. The pendulum's control requirements demand a rather good match with the model so η_q was chosen to be close to 1. This value results in a $\tilde{\mathbf{q}}^*$ with $m^* = 20$. This value satisfies the necessary condition $m^* \geq n^*$ required to estimate unique parameter values. Figure 4-2 plots the elements of $\tilde{\mathbf{q}}$ as the solid line and the squared norm of the vector up to the i^{th} component as the dashed line. The squared norm passes η_q at the 20^{th} most sensitive measurement (dotted line).

It is observed that the worst error in mode frequencies between the model and measurements is concentrated at resonances that couple strongly to pitch motion. It is thus not surprising that half the selected measurements are resonances that couple strongly to pitch motion including the first 4 elements of $\tilde{\mathbf{q}}^*$. All of the remaining selected measurements are resonances representing modes of vibration along the z axis or roll. These frequencies depend strongly on the selected stiffness parameters.

4.3.2 Parameter Estimation

The Gauss-Newton routine described in Section 4.2 was used to fit the chosen model parameters to the selected measured data. The MATLAB® code used to implement the routine is found in reference [45]. All resonances were measured to a resolution of 2 mHz. This value was chosen to maximize the precision of the measurements within the limited time available.

Since the worst error in mode frequencies was initially concentrated at the pitch modes, Figure 4-3 compares the measured stage 1 pitch to pitch transfer function to the model before the fit. The blue curve is the model and the red curve is the measurement. The top plot presents the transfer function magnitude and the bottom plot the transfer function phase. The worst error in mode frequency is about 16%.

The Gauss-Newton model fitting routine was run twice to estimate the reduced parameter set under two conditions. The first run included all 47 measurements, the

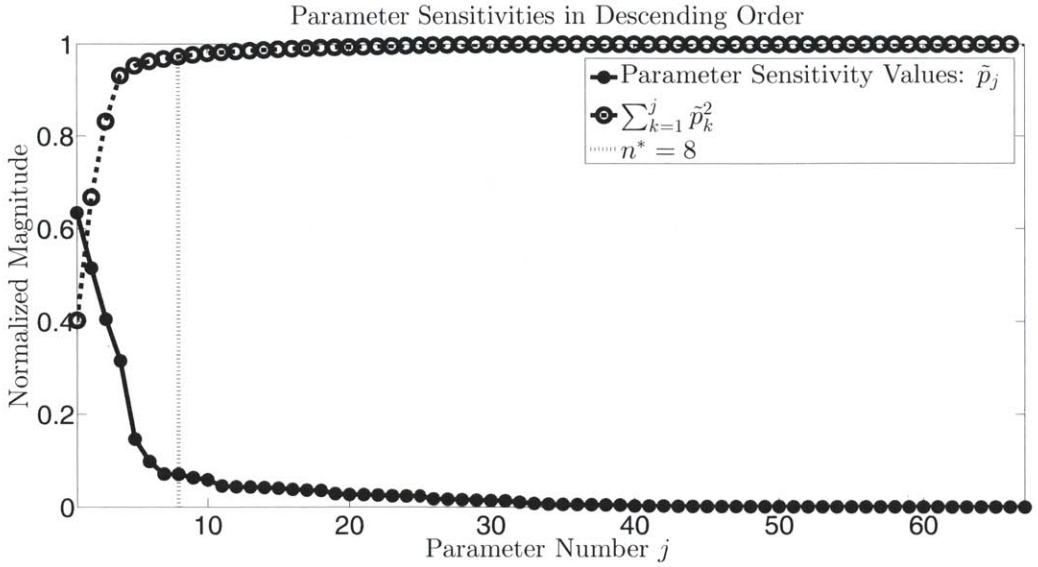

Figure 4-1: The solid line shows the values of the parameter uncertainty sensitivities in descending order ($\tilde{p}_j$). The dashed line shows the squared length of $\tilde{\mathbf{p}}$ up to its j^{th} component. The dashed line can be interpreted as the integral of the squared elements of the solid line. The dotted line indicates that all parameters beyond $j = n^* = 8$ in $\tilde{\mathbf{p}}$ are rejected.

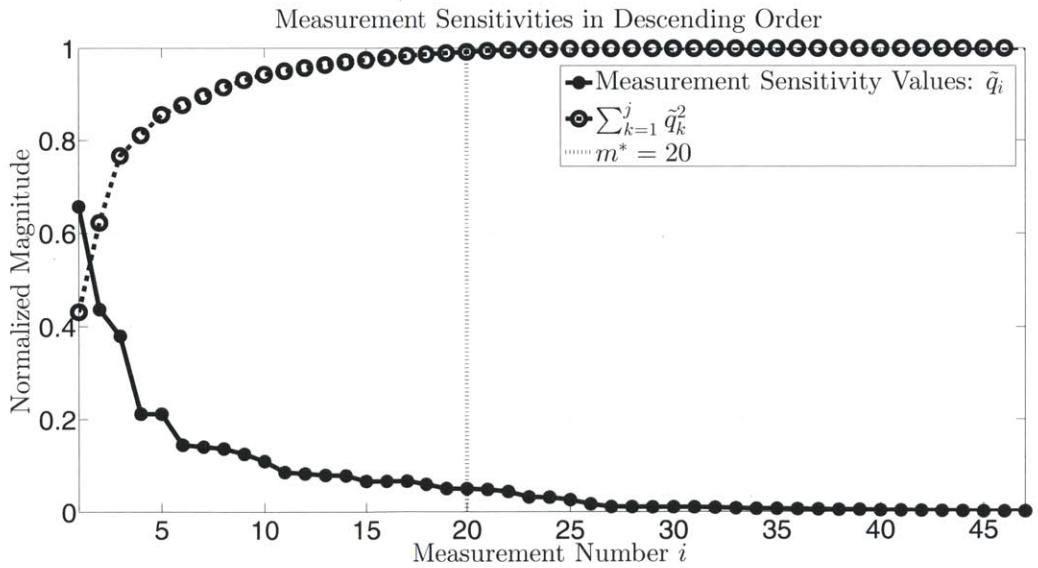

Figure 4-2: The solid line shows the values of the measurement sensitivities in descending order ($\tilde{q}_i$). The dashed line shows the squared length of $\tilde{\mathbf{q}}$ up to its i^{th} component. The dashed line can be interpreted as the integral of the squared elements of the solid line. The dotted line indicates that all parameters beyond $i = m^* = 20$ in $\tilde{\mathbf{q}}$ are rejected.

second run included only the reduced set of 20 measurements. A comparison of the results is presented below. The model fitting routine created a model that matches the measurements quite well in both cases.

Figure 4-4 presents the same information as Figure 4-3, but after running the fitting routine to estimate the selected 8 parameters using all 47 measurements. Figure 4-5 shows the same information again but after running the fitting routine using the reduced set of 20 measurements. In both cases the worst error in mode frequency is reduced to about 1%.

Figure 4-6 plots the error evolution between the model and the measurements for each iteration of the model fitting routine in both cases. The solid line represents the full measurement case and the dashed line the reduced case. The error is the sum of the squared percent difference between the model and each measurement. Notice that the error in both fitting cases is nearly identical. This match in the error indicates that the measurements ignored in the reduced set contribute very little information to the parameter estimation problem.

The FIM analysis from Section 4.2.3 estimates how much error we should expect in the estimated parameters after running the fit. The glass fiber bending points are identified to within about 0.1 mm, the ends of the springs 0.5 mm, the fiber vertical stiffness 0.1%, and the spring stiffnesses and the stage 2 pitch inertia 0.5%. These uncertainties are well below the a priori uncertainties in all cases.

4.4 Conclusion

The novel sensitivity selection procedure presented in this chapter reduces the parameter estimation problem from one that includes 67 parameters and 47 measurements to one that includes 8 parameters and 20 measurements. The parameter selection component to the procedure quantifies the importance of the a priori parameter uncertainty to determine which parameters require estimation to refine the model. The measurement selection component quantifies how much information each measurement contributes to the problem of estimating the selected parameters. Those

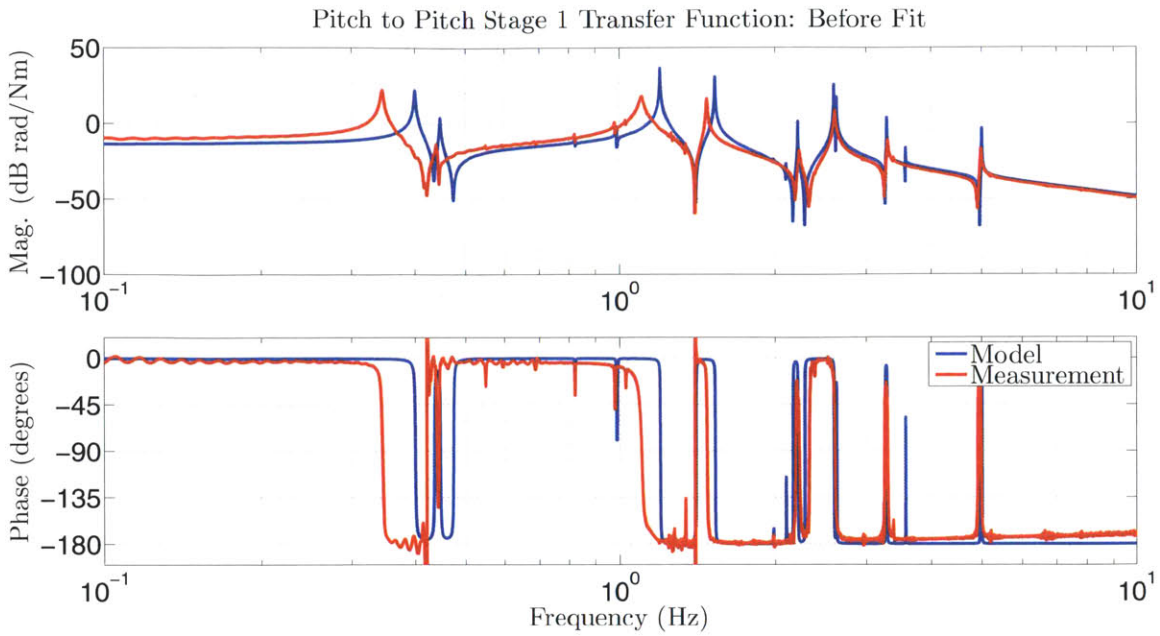

Figure 4-3: Original quadruple pendulum model against the measured stage 1 pitch to pitch transfer function. The blue curve is the model, the red the measurement.

measurements that contribute negligible information are not required.

The parameter estimation results show a significant improvement in the match between the model and the measurements. After the fit, the maximum percent error in mode frequencies is reduced from 16% to 1%. Thus, the selection procedure is effective in choosing those parameters that require estimation. The 27 least sensitive measurements contribute less than 1% of the parameter uncertainty information contained in the full set of 47 measurements. Discarding these measurements has negligible impact on the parameter estimation results. This measurement selection procedure is useful in determining how to focus limited resources to those measurements that are most important. For the quadruple pendulum parameter estimation problem, there would be no need to collect the least sensitive 27 measurements at all. The time required to measure them could be better spent refining the top 20.

This sensitivity analysis also provides a lot of useful information about the quadruple pendulum itself. It states that the most important uncertainty is focused on effective wire and fiber bending points. Spring stiffness follows in importance. A distant third is rotational inertia. Some further investigation reveals that the bending points

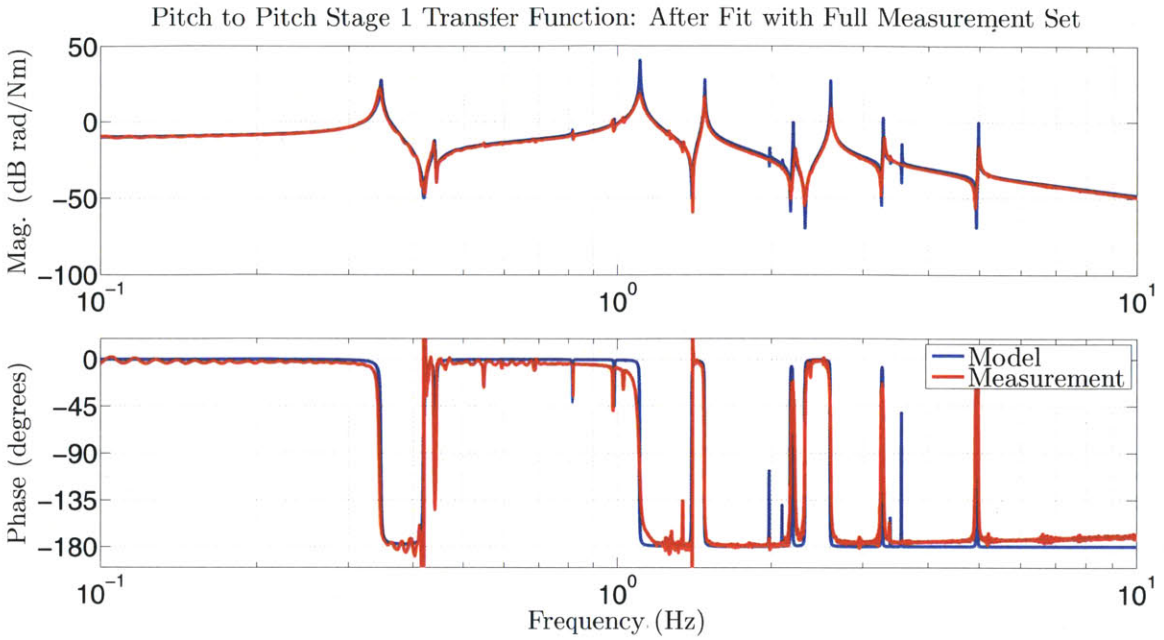

Figure 4-4: Model fit result, using all 47 measurements, against the measured stage 1 pitch to pitch transfer function. The blue curve is the model, the red the measurement.

determine the rotational pitch stiffness. These bending points are nominally placed only a few millimeters from the centers of mass. It is this distance to the center of mass that determines pitch stiffness. In this example, the uncertainty is the same order of magnitude as the nominal distance itself, which is why the largest error appeared at pitch resonant frequencies. We can thus conclude that on future quadruple pendulums more effort should be placed on measuring or constraining these values more precisely.

The limitation with this sensitivity analysis is that one needs to know the a priori parameter uncertainty bounds. Choosing unfair bounds will skew which parameters are selected. Additionally, this analysis says nothing on how precise each measurement should be made. It assumes each measurement is exact. The Fisher Information Matrix must be used to determine how much precision the selected measurements require to reduce the uncertainty on the selected parameters.

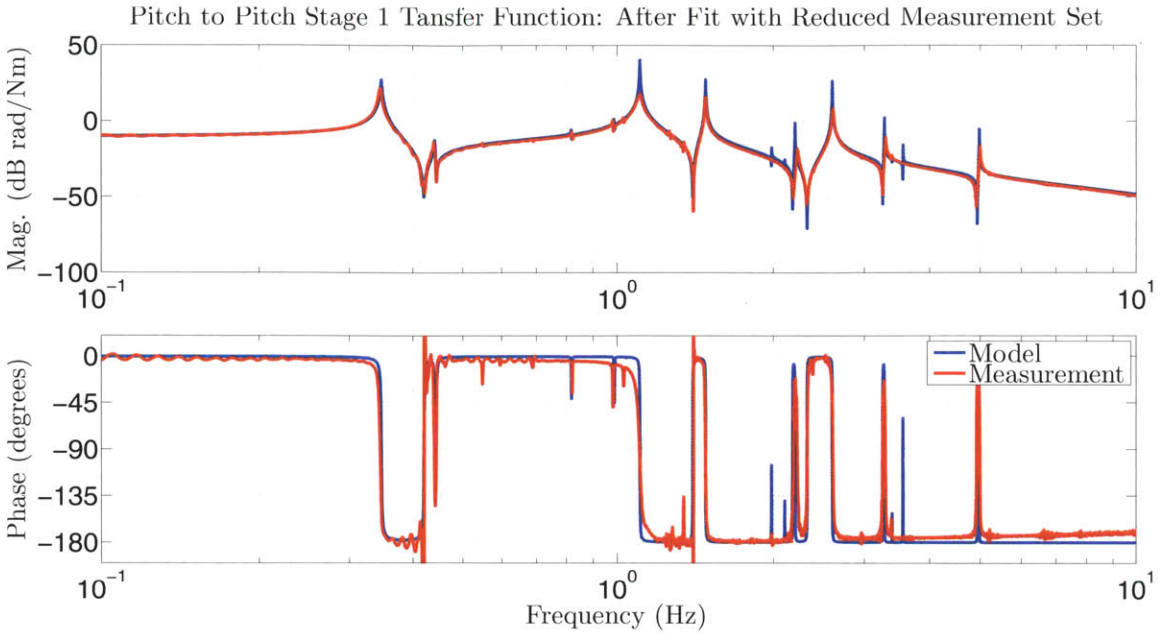

Figure 4-5: Model fit result, using the selected 20 measurements, against the measured stage 1 pitch to pitch transfer function. The blue curve is the model, the red the measurement.

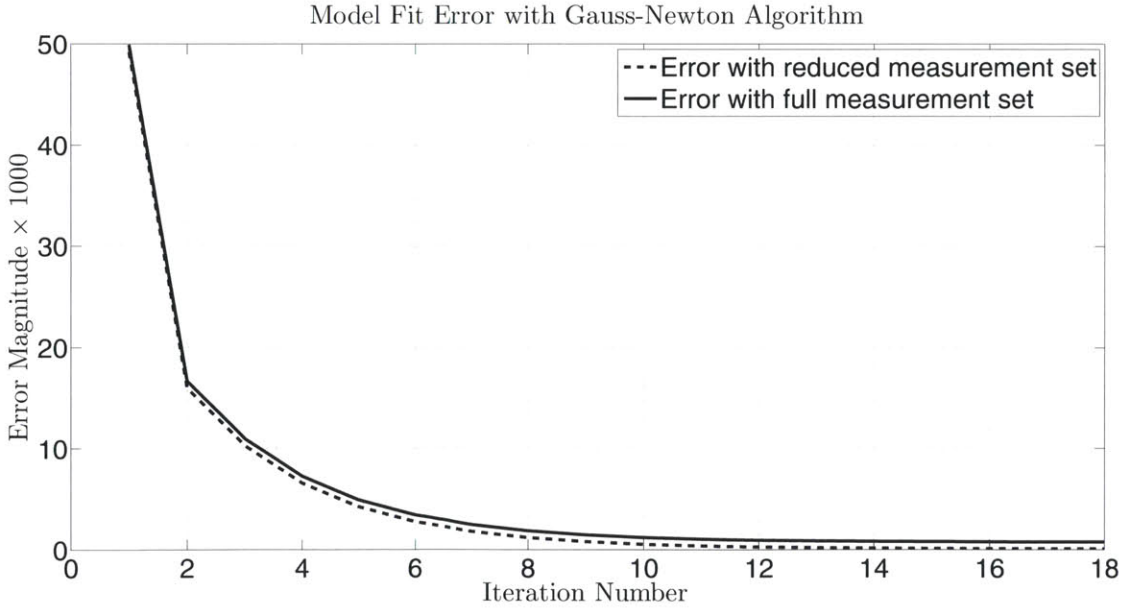

Figure 4-6: The solid line is the sum of the squared error at each iteration of the fitting routine when all 47 measurements are used. The dashed line is the sum of the squared error when only the selected 20 measurements are used.

Chapter 5

Actuator Sizing

This chapter applies the pseudoinverse of the plant transfer function matrix to validate that the actuators provide sufficient dynamic range to fully reject the stochastic disturbances. This method is independent of feedback design, but does suggest characteristics of the optimal design [47].

The model of the pendulum used in this analysis is discussed in detail in Chapters 3 and 4 and references [39, 4, 48]. A reduced order version of the state space matrices is listed in Appendix A. The model was reduced by considering only motion of the stages of the pendulum along the x axis, or the interferometer axis. The reduction is possible because the symmetry of the pendulum sufficiently decouples motion along this axis from all the others.

As discussed in Section 2.1.2, the relative displacement between the test masses must be suppressed with active control to under 10^{-15} m RMS. This suppression involves rejecting the black 'Advanced LIGO Goal' line in Figure 1-6. Figure 5-1 extends this line down to 0.1 Hz to make the amount of suppression more obvious. Without this control, the black line has an overall displacement of about 3×10^{-7} m RMS. Thus, at least a factor 3×10^{8} suppression is required. The control task can be thought of as a tracking problem since the test masses are actuated to follow as closely as possible this black line (to suppress relative displacement). The 'Advanced LIGO Goal' will be considered as a reference tracking signal for this chapter.

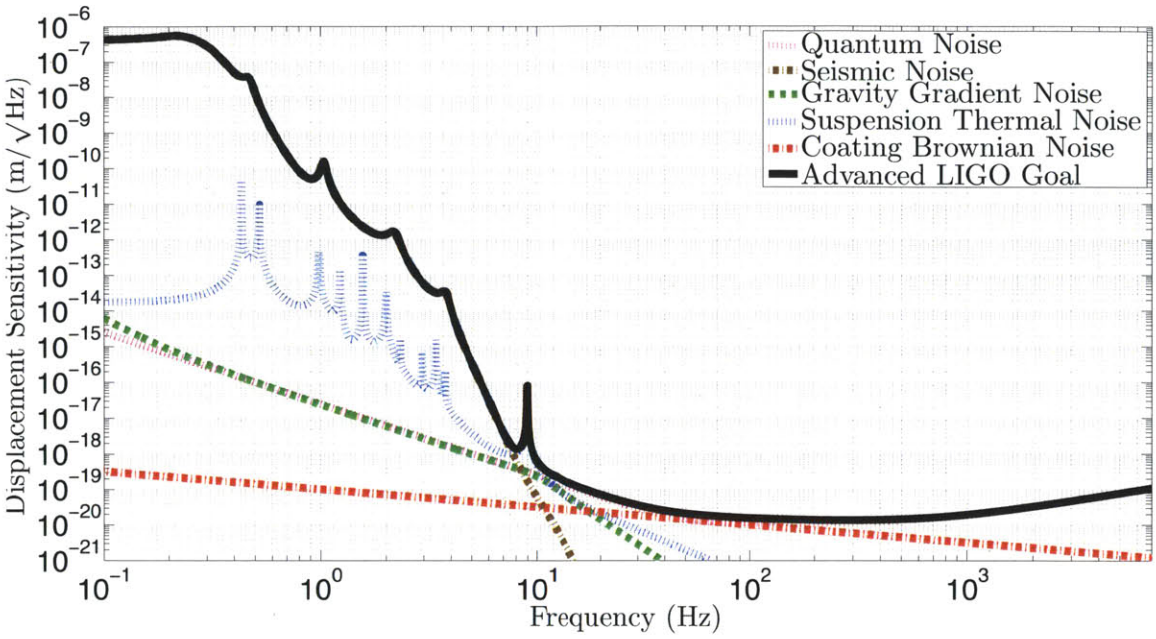

Figure 5-1: The Advanced LIGO displacement sensitivity from Figure 1-6 extended to low frequencies.

5.1 Quadruple Pendulum Actuator Sizing

The schematic in Figure 5-2 and the block diagram in Fig. 5-3 summarize the signal flow for this tracking control. In reality the control will be achieved by actuating on both ETMs. Since they are conceptually identical, this loop is simplified by assuming there is only one, with double the dynamic range. The test mass position, y is measured interferometrically. This signal is compared to the desired position, R, to generate the error signal. This error signal is passed to feedback filters that drive each of the four pendulum stages along the axis of the interferometer.

The frequency response of the sensor is ignored in this case to focus on the actuators. The DAC saturates at $\pm 10\,\mathrm{V}$, so the actuators must have enough dynamic range to maintain the DAC voltages, $\mathbf{u}$, within this limit.

To model the least mean square DAC voltages required to follow the desired trajectory R, the models of the actuators given in Section 2.3 (plotted in Figs. 2-5, 2-6, and 2-7) are merged with the pendulum model in Appendix A to create a 1×4 plant $\mathbf{G}$. Only the final output of the pendulum model is considered here. The goal of the control loop is to apply control forces to solve the transfer function (5.1). We

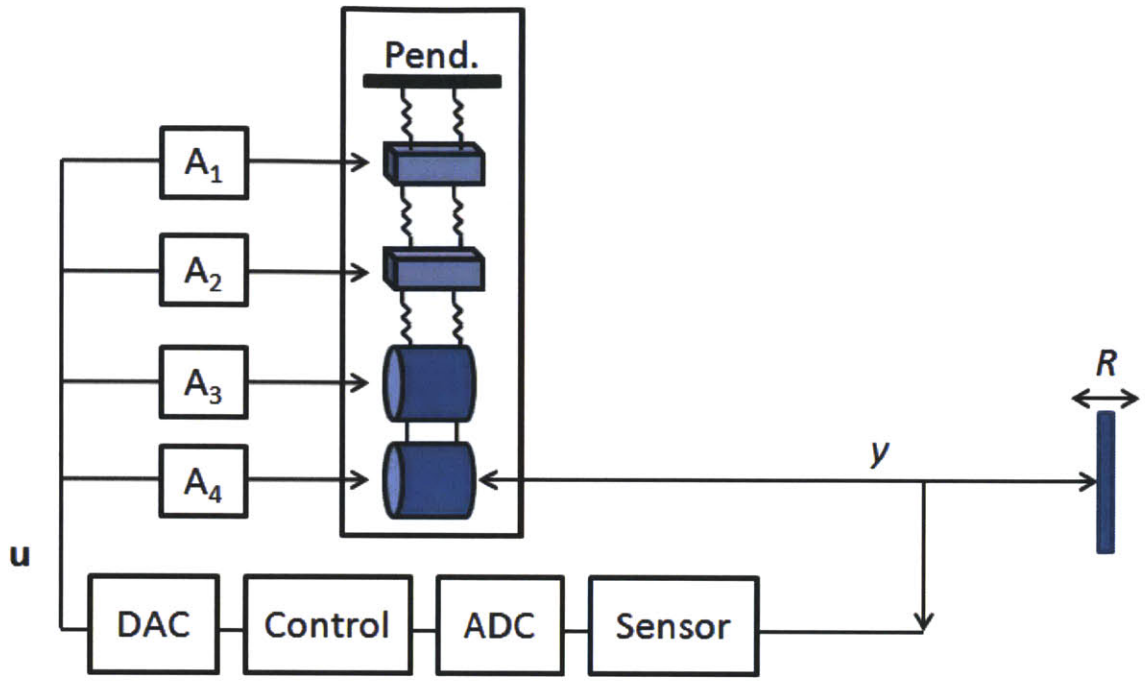

Figure 5-2: A schematic diagram of the test mass (stage 4) positioning control loop. Actuation is applied on all four stages of the pendulum to control the position of the test mass, y, along the interferometer axis. The signals $\mathbf{u}$ must be maintained within the $\pm 10\,\mathrm{V}$ limit of the DAC. A_1 to A_4 represent the actuator dynamics.

assume the feedback loops have enough gain to approximate this simplification within the band of interest (up to 10 kHz).

$$R = \mathbf{G}\mathbf{u} \tag{5.1}$$

R is taken to be the 'Advanced LIGO Goal' amplitude spectrum from Fig. 5-1.

The least squares DAC voltages are given by the 4×1 pseudoinverse, $\mathbf{G}^{\dagger}$. However, since R is a random Gaussian signal, we do not need the phase information from $\mathbf{G}^{\dagger}$. The amplitude spectrum of the least squares actuation, $\mathbf{u}_{ls}$, is

$$\mathbf{u}_{ls} = |\mathbf{G}^{\dagger}|R \tag{5.2}$$

where $|\mathbf{G}^{\dagger}|$ is the real column vector containing the magnitudes of each complex element in $\mathbf{G}^{\dagger}$.

Fig. 5-4 shows the resulting amplitude spectra of the DAC voltage sent to each

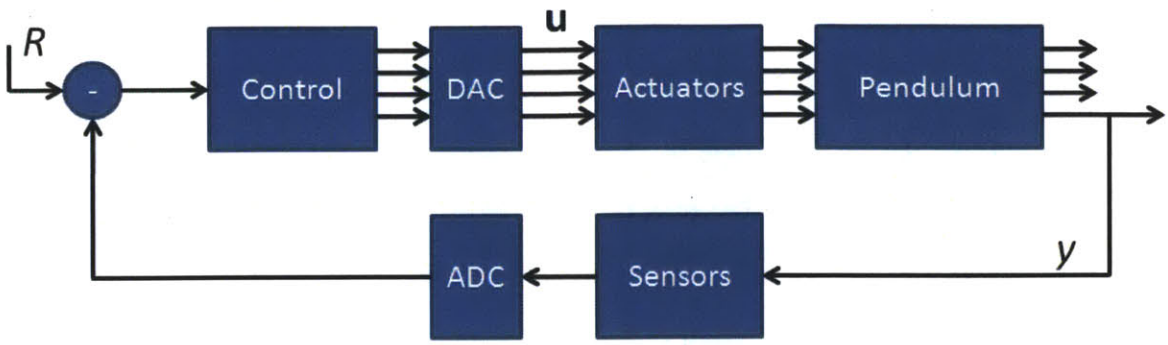

Figure 5-3: The block diagram of the test mass positioning control loop. Actuation is applied on all four stages of the pendulum to control the position of the test mass, y, along the interferometer axis. The signals $\mathbf{u}$ must be maintained within the $\pm 10\,\mathrm{V}$ limit of the DAC.

actuator. The relative distance between values at each frequency represents the relative importance of each actuator. Larger values state higher importance. We see then that it is most advantageous to use the stage 4 actuator above 4 Hz while using the state 1 actuator below. From a least squares point of view the stage 2 and 3 actuators gives us little.

Note that $\mathbf{G}^{\dagger}$ (in general non-causal) does not represent a feedback control law, but rather the mapping of R to the minimum open loop force that would be required to move the test mass from rest to a trajectory with the amplitude spectrum of R. Thus, this least squares result sets a lower limit on an optimal feedback control design.

The least squares RMS voltages for the actuators, $\bar{\mathbf{u}}_{ls}$, are found in Eq. (5.3) by integrating the $\mathbf{u}_{ls}$ power spectra and taking the square root. The probabilities, $\mathbf{p}$, of saturating the DAC at a given time step are calculated in Eq. (5.4) using the error function erf.

$$\bar{\mathbf{u}}_{ls} = \sqrt{\int_0^{\infty} \mathbf{u}_{ls}^2 df} \tag{5.3}$$

$$\mathbf{p} = 1 - \mathrm{erf}\left(\frac{10}{\sqrt{2}\bar{\mathbf{u}}_{ls}}\right) \tag{5.4}$$

The factor 10 is a scaling factor from the DAC saturation limit.

Table 5.1 summarizes the RMS DAC voltages and probabilities of saturation for the solution of the least squares actuation required at each stage. All actuators have

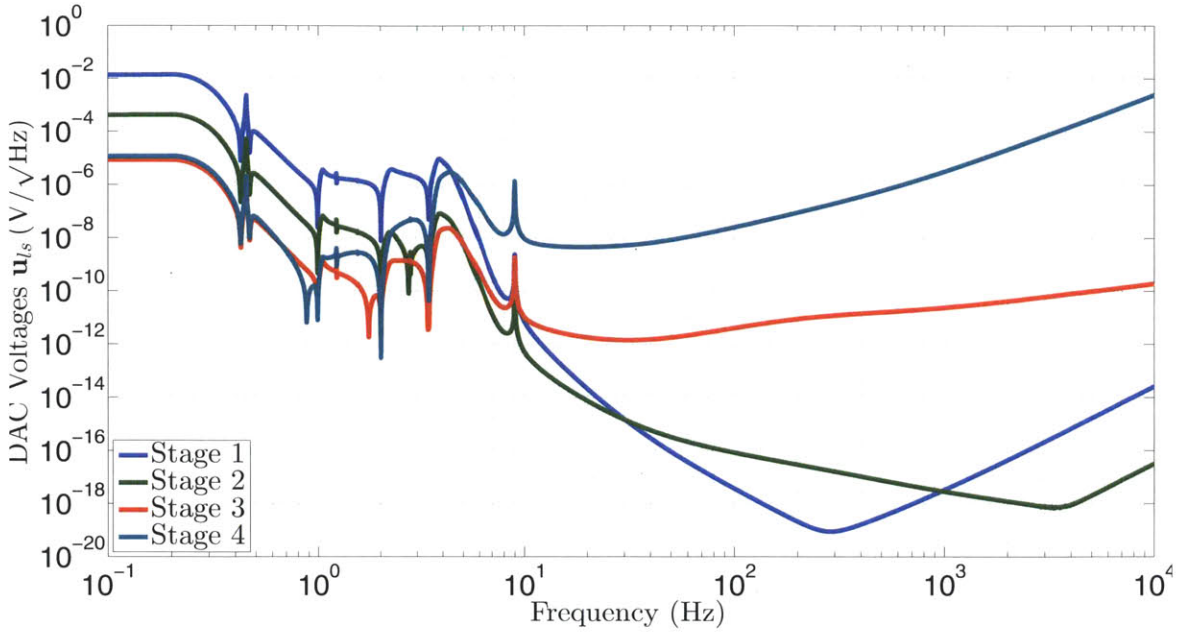

Figure 5-4: Amplitude spectra of the DAC voltages $\mathbf{u}_{ls}$.

virtually no risk of saturation. The stage 4 actuator comes the closest to saturation. If the spectrum of R is stationary then the probability of saturation will remain virtually non-existent.

However, on a stormy day R will increase and the control loops will demand more drive, possibly causing saturation. Fortunately, the information in the table and in Fig. 5-4 provide hints as to how we might tune the control loops by telling us which actuators are most advantageous at which frequencies. This knowledge allows the actuation to be efficiently offloaded to other stages. Further, the finite bandwidth of the control loops, within 1000 Hz, reduces the force estimate needed at stage 4. The analysis in this chapter assumes a feedback loop with broadband infinite gain, while we only need enough to limit test mass motion to 10^{-15}m RMS.

Additionally, we expect stormy weather to impact the low frequency band below 0.2 Hz. It turns out that scaling the reference signal R in this band scales the voltages of the upper three stages given by Table 5.1 by a similar amount. The stage 4 voltage is largely unaffected. Thus, we can use these analysis techniques to conclude that from a least squares point of view we can tolerate an increase in R from stormy weather by a factor of almost 100, a reasonable margin.

Table 5.1: The least squares DAC voltages $\bar{\mathbf{u}}_{ls}$ of each actuator and the probabilities $\mathbf{p}$ of saturating at the $\pm 10\,\mathrm{V}$ limit. The probabilities refer to each DAC sample time.

	$\bar{\mathbf{u}}_{ls}$ (V RMS)	$\mathbf{p}$
Stage 1	0.01	0.0
Stage 2	3.2×10^{-4}	0.0
Stage 3	6.5×10^{-6}	0.0
Stage 4	0.097	0.0

Further, as indicated by Figure 5-4, the stage 2 and 3 actuators are doing very little here. Consequently, part of the stage 4 actuation can be offloaded to these stages to reduce the load at other stages at the expense of an increase in the total amount of drive. In fact, there are a number of other benefits for offloading the drive to higher stages. Perhaps the most important reason is stability. Crossing the stage 1 and 4 actuation as steeply as the 4 Hz crossover shown in Fig. 5-4 will undoubtedly lead to a seriously unstable phase margin. Bridging the actuation bands at these two stages with even a small amount of drive at the stages in between will greatly improve the phase margin.

5.2 Conclusion

Using the pseudoinverse of the plant transfer function and the modeled system noise it was shown that the minimum mean square actuation required for the Advanced LIGO quadruple pendulum falls within the design constraints with reasonable margins. This analysis is possible without feedback loop design. Further, the pseudoinverse quantifies which actuator is most effective at each frequency providing critical information for later feedback design.

In general, additional analysis is needed during the design of any feedback loop since stability requirements will prevent the minimum mean square actuation from actually being achieved. The need to maintain reasonable phase margin prevents any given feedback filter from rolling off in magnitude too steeply. This limitation causes any error signal or noise outside the bandwidth of a particular actuator's loop to leak through and contribute to the RMS drive of that actuator.

Chapter 6

Modal Damping

This chapter develops the method of modal damping. Modal damping is a key component of the method of adaptive modal damping described in Chapter 7. Section 6.1 outlines the basic working parts of modal damping. Section 6.2 describes the design of the modal feedback filters. The state estimation used to reconstruct the incomplete sensor signals is designed in Section 6.3. Finally, the damping limitations imposed by the incomplete set of actuators are determined by Section 6.4.

The work in this chapter expands upon the modal damping work in references [4, 28] in a number of ways. First, the estimator design has been streamlined, particularly for MIMO modal damping. Second, the estimator outputs the modal velocities rather than the modal displacements, which improves sensor noise filtering at virtually no cost. Third, a theoretical maximum damping limit has been established. Finally, the modal feedback filter design has been simplified and given greater phase margin. This new filter design has less noise filtering, but permits a larger range of damping gains with the assumption that the damping will be adjusted as needed.

6.1 Overview

In order to best meet the challenges stated in Section 1.3.4, we would like to have an easily tunable damping design with high gain at the resonances, and very low gain at 0 Hz and above 10 Hz. Low gain at 0 Hz can be achieved simply by feeding back

velocity signals, but low gain at 10 Hz is difficult to achieve in a stable way.

A design technique that contributes well to these goals and to the adaptation developed in Chapter 7 is known as modal damping. Modal damping is a method that permits the damping of each mode in a resonant system. The method simplifies the pendulum's active damping by decomposing the problem of a high dimensional system with many coupled degrees of freedom into a set of simple decoupled single degree of freedom problems with similar dynamics. It is not adaptive on its own but is easily worked into an adaptation scheme due to its decoupling property.

Figure 6-1 provides a block diagram of the modal damping signal flow for the four quadruple pendulum modes along the x axis which couple most strongly to the interferometer output. These x modes are already sufficiently decoupled from the other 20 by the symmetry of the system, allowing x displacement of the 4 stages to be handled independently. The diagram would have a similar structure if all modes were considered.

The modal decomposition that makes modal damping possible is listed in Eqs. (6.1) to (6.3). Eq. (6.1) represents the Cartesian equations of motion of the pendulum where $\mathbf{M}$ is the mass matrix, $\mathbf{K}$ the stiffness matrix, $\mathbf{f}$ the vector of control forces and $\mathbf{x}$ the displacement vector. The eigenvector basis $\mathbf{\Phi}$ of the matrix $\mathbf{M}^{-1}\mathbf{K}$ transforms between $\mathbf{x}$ and the modal displacement coordinates $\mathbf{q}$. Substituting Eq. (6.2) into Eq. (6.1) and multiplying on the left by $\mathbf{\Phi}^T$ we get the modal equations (6.3). The subscript m indicates that the mass matrix, stiffness matrix, and the force vector are now in the modal domain. Note that these modal matrices are diagonal, which is a mathematical result of the fact that the dynamics of each modal DOF are decoupled.

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{f} \tag{6.1}$$

$$\mathbf{x} = \mathbf{\Phi}\mathbf{q} \tag{6.2}$$

$$\mathbf{M_m}\ddot{\mathbf{q}} + \mathbf{K_m}\mathbf{q} = \mathbf{f_m} \tag{6.3}$$

For the application of damping a real system there are two important transfor-

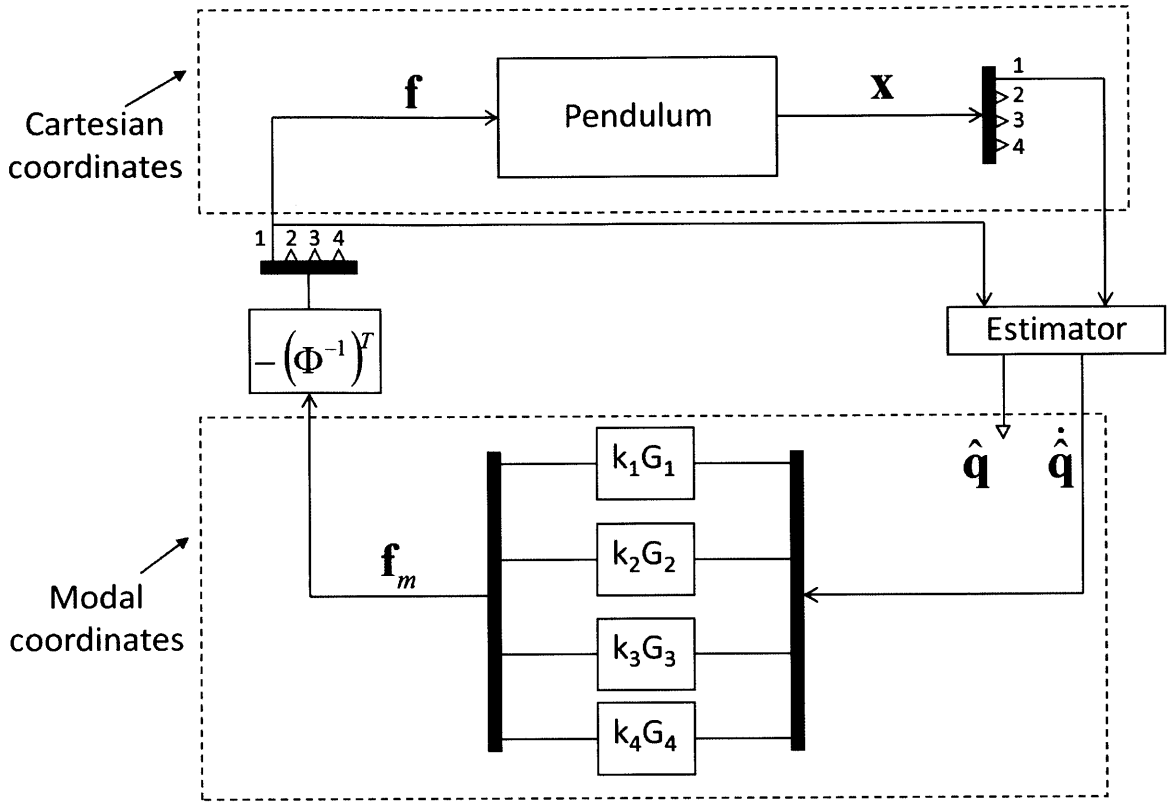

Figure 6-1: A block diagram of a modal damping scheme for the four x modes. An estimator converts the incomplete sensor information into modal signals. The modal signals are then sent to damping filters, one for each mode. The resulting modal forces are brought back into the Cartesian coordinate system through the transposed inverse of the eigenvector matrix $\mathbf{\Phi}$. Only stage 1 forces are applied for enhanced sensor noise filtering to stage 4.

mations required.

$$\mathbf{q} = \mathbf{\Phi}^{-1}\mathbf{x} \tag{6.4}$$

$$\mathbf{f} = (\mathbf{\Phi}^{-1})^T\mathbf{f_m} \tag{6.5}$$

Eq. (6.4) transforms the sensor signals into the modal domain. Control can then be applied where the dynamics are decoupled. This transformation requires the full measured state. As Figure 6-1 shows, we only measure stage 1 of the pendulum, so an estimator is used to perform the change of coordinates. The design of this estimator is detailed in Section 6.3.

The second transformation, (6.5), occurs after the modal damping feedback filters, labeled as k_1G_1 to k_4G_4 in Figure 6-1. It converts the modal damping forces $\mathbf{f_m}$ generated by these filters into Cartesian forces that can be applied to the pendulum by the actuators. Note that actuation is only available at stage 1 for this damping control. Although actuators exist at the lower stages, employing them would allow part of the sensor noise to bypass the mechanical filtering of the stages above.

Consequently, the applied damping forces are a projection of the modal damping forces to stage 1. This limitation means that the damping is not truly modal, resulting in cross coupling between modal feedback loops. Fortunately, the analysis in Section 6.4 states that this coupling is minimal for closed loop damping ratios less than 0.2. This upper limit is more than enough to meet the damping requirements.

6.2 Feedback Loop Design

Note, that Figure 6-1 indicates we choose to feedback the modal velocities $\hat{\dot{\mathbf{q}}}$ rather than the displacements $\hat{\mathbf{q}}$. If we were to feedback the displacements we would need an AC coupling zero at 0 Hz to ensure the filters meet the AC coupling requirement. One advantage of an estimator is that we can simply take the modal velocity outputs directly instead. Importantly, feedback with the estimated velocities also has less noise amplification.

The nominal design of the feedback filters is relatively simple since each modal plant is identical except shifted in frequency and magnitude. The filter design has a total of 2 zeros and 4 poles. A complex pole pair is placed at 2.5 times the frequency of the mode for lowpass filtering. These poles are placed 10° off the imaginary axis to achieve slight enhancements in the filtering and phase margin. The factor of 2.5 was chosen to achieve the most aggressive filtering possible while leaving enough phase margin to allow sufficiently high damping ratios. The remaining complex pole and zero pairs create a notch in the loop gain at 10 Hz to create a more aggressive cutoff in the lowpass filtering near the beginning of the GW detection band. This notch is designed to have a depth of a factor of 10 and a quality factor of 10. It is realized

with the zeros and poles at 10 Hz, 2.866° and 30° off the imaginary axis respectively.

Eq. (6.6) is the transfer function form of these damping filters, with arbitrary scaling, where s is the Laplace variable and ω_n is the modal frequency in radians/s. Eq. (6.7) is an example filter for a 1 Hz mode. Figure 6-2 illustrates an example loop gain transfer function and root locus diagram with this example filter.

$$G_{\omega_n} = \frac{s^2 + 2\pi s + (20\pi)^2}{(s^2 + 5\sin(10°)\omega_n s + (2.5\omega_n)^2)(s^2 + 20\pi s + (20\pi)^2)} \tag{6.6}$$

$$G_{1\,\mathrm{Hz}} = \frac{s^2 + 6.283s + 3948}{(s^2 + 5.455s + 246.7)(s^2 + 62.83s + 3948)} \tag{6.7}$$

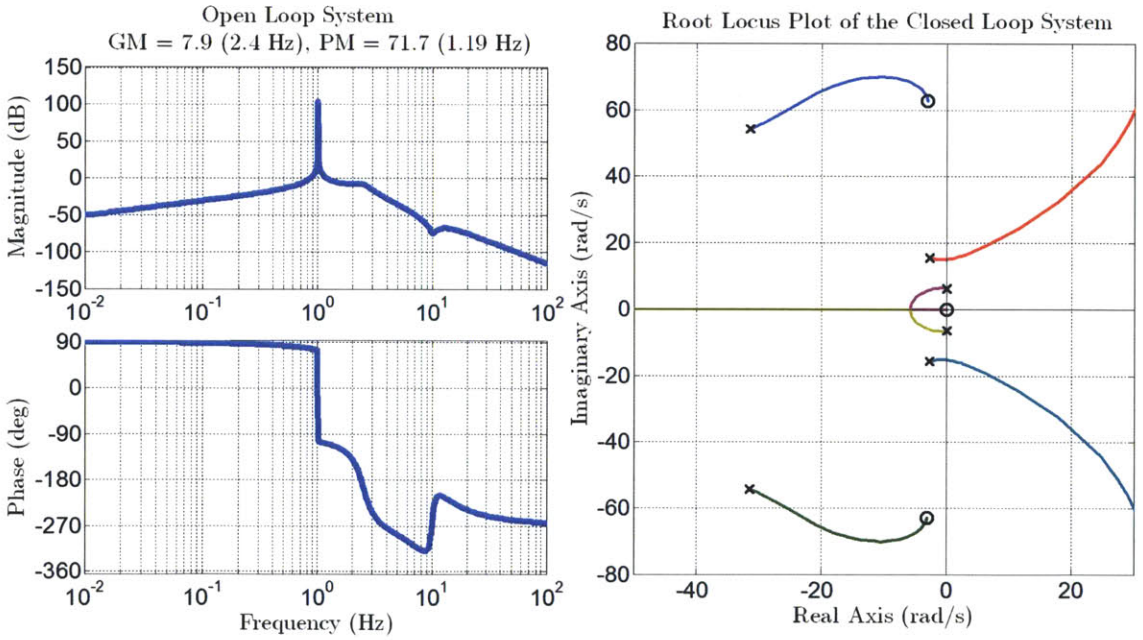

Figure 6-2: Left, the loop gain transfer function of an example 1 Hz modal oscillator with its damping filter. Right, the root locus plot of the modal damping loop. The plant contributes the large resonant peak and 0 Hz zero, the damping filter contributes the remaining poles and zeros. All the damping loops have the same basic shape but are shifted in frequency and gain (except the fixed frequency 10 Hz notch).

6.3 State Estimation

The mathematics of modal damping requires the positions of all four stages to be measured. However, only stage 1 is directly observed, as Figure 6-1 illustrates. The stages below have effectively no sensors since any measurement would reference a moving platform with its own dynamics. Consequently an estimator, as in (6.8), must be employed to reconstruct the full dynamics.

$$\begin{bmatrix} \dot{\hat{\mathbf{q}}} \\ \ddot{\hat{\mathbf{q}}} \end{bmatrix} = \mathbf{A_m} \begin{bmatrix} \hat{\mathbf{q}} \\ \dot{\hat{\mathbf{q}}} \end{bmatrix} + \mathbf{B_m u} - \mathbf{L_m}(\mathbf{C_m} \begin{bmatrix} \hat{\mathbf{q}} \\ \dot{\hat{\mathbf{q}}} \end{bmatrix} - \mathbf{y}) \tag{6.8}$$

The estimator equation in (6.8) receives the control input $\mathbf{u}$ and the stage 1 sensor signals $\mathbf{y}$ in the Cartesian frame and estimates the corresponding modal signals $\dot{\hat{\mathbf{q}}}$ using the modal state space matrices $\mathbf{A_m}$, $\mathbf{B_m}$, $\mathbf{C_m}$ obtained from the system model in Appendix A. The subscript m indicates the matrices refer to the modal frame. The goal is now to design the estimator feedback gain $\mathbf{L_m}$.

The observer separation principle [49] allows us to keep the already established control scheme, provided that the estimator has an accurate model of the pendulum, which is the assumption here. Note that the estimator replaces the direct signal transformation in Eq. (6.4).

The estimator is designed using the Linear-Quadratic Regulator (LQR) technique [49] which solves the cost function in (6.9) and (6.10).

$$J = \int_r^\infty (\begin{bmatrix} \tilde{\mathbf{q}}^T & \dot{\tilde{\mathbf{q}}}^T \end{bmatrix} \mathbf{Q} \begin{bmatrix} \tilde{\mathbf{q}} \\ \dot{\tilde{\mathbf{q}}} \end{bmatrix} + {\mathbf{z_m}}^T \mathbf{R} \mathbf{z_m}) \, dt \tag{6.9}$$

$$\mathbf{L_m} = \underset{\mathbf{L_m}}{\operatorname{argmin}}(J) \tag{6.10}$$

The $\mathbf{Q}$ matrix weights the accuracy of the modal state estimation while the $\mathbf{R}$ matrix weights the cost of using a noisy measurement. $\mathbf{z_m}$ is defined as

$$\mathbf{z_m} = -\mathbf{L_m}^T \begin{bmatrix} \tilde{\mathbf{q}} \\ \dot{\tilde{\mathbf{q}}} \end{bmatrix} \tag{6.11}$$

$\mathbf{L_m}$ is the estimator feedback matrix determined by (6.10).

An estimator could be generated in the form of a Kalman filter for this application. However, a Kalman filter is concerned with optimal sensor data recovery whereas our goal is to optimize the noise reaching the pendulum's test mass. Sacrificing sensor signal accuracy and the amount of damping is acceptable to a certain degree if the noise performance is improved. Consequently, a more direct approach is taken to obtain the values of $\mathbf{Q}$ and $\mathbf{R}$.

In this chapter we develop modal damping for just the x axis dynamics. However, up to six stage 1 DOFs can be considered simultaneously. The following estimator optimization technique is constructed for this more general case, and then applied to simpler case of x dynamics.

First, $\mathbf{Q}$ is set by placing along the diagonal the square of the modal velocity impulse response amplitudes. Each amplitude is given by $M_{m,ii}^{-1}$ where $M_{m,ii}$ is the modal mass of mode i. In this way the modes with the greatest responses will be damped more efficiently. For $\mathbf{R}$, we restrict the matrix to be diagonal, a common assumption. We are then left with up to six parameters to simultaneously optimize relating to six Cartesian stage 1 measurement signals. To choose the best value of $\mathbf{R}$ the control design is first set so that the damping time requirement is met assuming full state information. An optimization routine then simulates the performance of the closed loop system for estimators designed using many values of $\mathbf{R}$ over a sufficiently large space. A cost function, (6.12), dependent on the performance criteria is calculated for each value of $\mathbf{R}$. The value that minimizes this cost function in the chosen interval is optimal.

The two competing performance criteria are the 10 s settling time and the sensor noise amplification as stated in Section 2.1. The sensor noise is a known measured quantity approximately $10^{-10}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ beyond 10 Hz for each sensor. Since the con-

tribution to test mass motion drops off quickly in the frequency domain, we will only consider the contribution at the start of the gravitational wave detection band, 10 Hz. Note that these criteria are not required to be met simultaneously. However, in this section we considered the stronger case where they are met simultaneously.

These two performance criteria are represented in the optimization routine with the cost function in (6.12).

$$J_R(\mathbf{R}) = \max_{\text{DOF}}(T_s^2) + \max_{\text{DOF}}(N^2) \tag{6.12}$$

$$\mathbf{R} = \underset{\mathbf{R}}{\text{argmin}}(J_R) \tag{6.13}$$

T_s is the normalized stage 4 settling time; chosen as the maximum settling time of all the DOFs normalized by the requirement, 10 s. N is the maximum sensor noise contribution to the test mass relative to the respective requirement at 10 Hz for all DOFs considered (from Table 2.1).

As an example, the results of (6.12) obtained from a MATLAB® script (Appendix B) are plotted in Figure 6-3 for the x DOF where $\mathbf{R}$ is scalar valued. The optimal value occurs at $\mathbf{R} = 9 \times 10^{-8}$. Since the total cost is less than 1, both design requirements are met at this optimal point. The spectral density of the test mass displacement under this optimized modal damping control is plotted in Figure 6-4. Note that the sensor noise in this example is reduced from the OSEM noise level by $\sqrt{2}$ because two sensors are used to measure x displacement.

6.4 Top Mass Limited Damping

Projecting the modal damping forces to the subset of stage 1 actuation causes the decoupled modal degrees of freedom to couple back together again. However, the modes may be considered decoupled for sufficiently small damping.

To quantify these limits, we make two assumptions. First, we assume perfect knowledge of the entire system state by invoking the observer separation principle. This principle allows us to consider the influence of the estimator independent of the

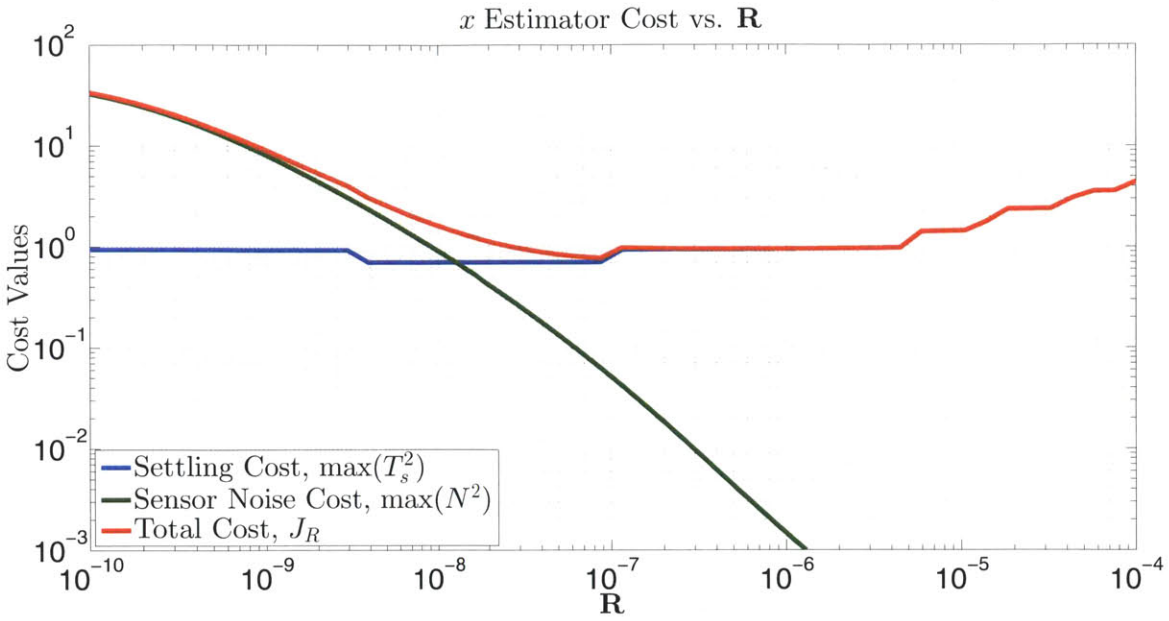

Figure 6-3: The components of the cost function (6.12) for the x DOF as a function of $\mathbf{R}$ calculated by the optimization routine. At each value of $\mathbf{R}$ the closed loop system performance is simulated using the estimator design based on the LQR solution with that particular $\mathbf{R}$ value.

control. In Figure 6-1 this assumption replaces the estimator with the time derivative of the matrix transformation in Eq. (6.4).

The modal damping forces then influence the plant according to

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{\Gamma}(\mathbf{\Phi}^{-1})^T\mathbf{f_m} \tag{6.14}$$

The matrix $\mathbf{\Gamma}$ is the projection matrix that restricts actuation to stage 1. It is diagonal, where the elements corresponding to actuated DOFs are 1, and those corresponding to unactuated DOFs are 0.

Second, the modal damping filters G_1 to G_4 are assumed to be unity gains, $G_i = 1$. The gains k_i in Figure 6-1 are user defined scalars. The substitution $\mathbf{f_m} = \mathbf{B}\dot{\mathbf{q}}$ is then made where $\mathbf{B}$ is the diagonal damping matrix with each k_i making up the diagonal elements. Moving all terms to the left results in the closed loop modal equation of motion (6.15). This assumption is checked by Section 6.4.3.

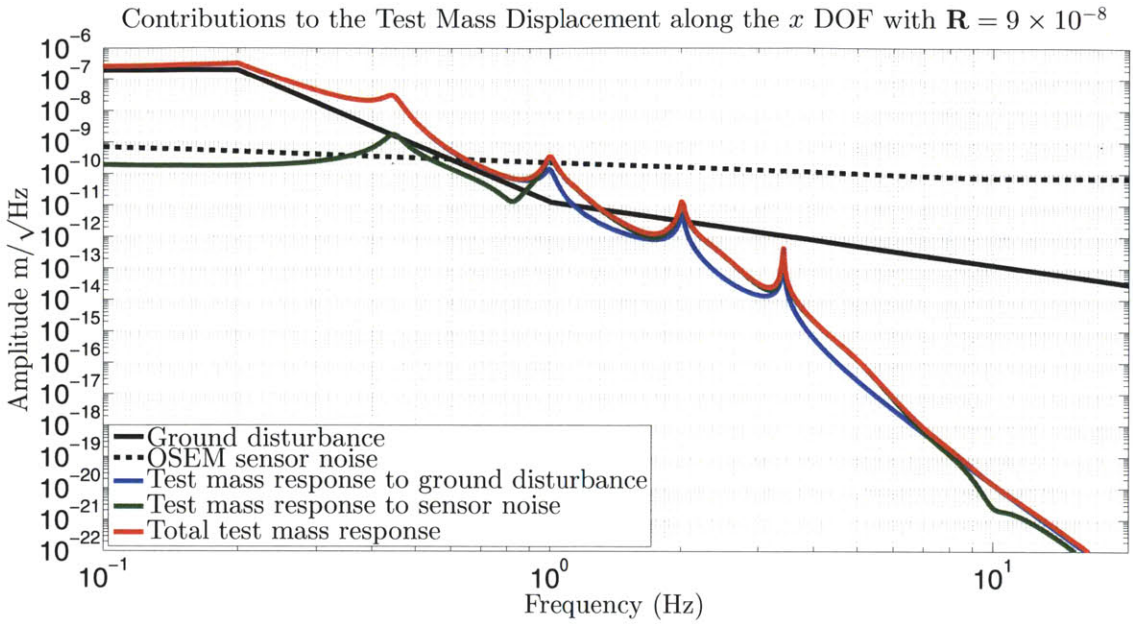

Figure 6-4: An amplitude spectrum showing a simulation of the test mass displacement along the x DOF under the influence of the optimized modal damping loop with $\mathbf{R} = 9 \times 10^{-8}$. The black dashed line is the sensor noise and the green line is its contribution to the test mass displacement. The solid black line is the ground disturbance and the blue line is its contribution to the test mass displacement. The red line is the uncorrelated stochastic sum of both contributions.

$$\mathbf{M_m}\ddot{\mathbf{q}} + \mathbf{\Phi}^T\mathbf{\Gamma}(\mathbf{\Phi}^{-1})^T\mathbf{B}\dot{\mathbf{q}} + \mathbf{K_m}\mathbf{q} = 0 \tag{6.15}$$

An equivalent representation of Eq. (6.15) is the state space form in Eqs. (6.16) and (6.17).

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \ddot{\mathbf{q}} \end{bmatrix} = \mathbf{A_m} \begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix} \tag{6.16}$$

$$\mathbf{A_m} = \left[\begin{array}{c|c} \mathbf{0}_{n\times n} & \mathbf{I}_{n\times n} \\ \hline -\mathbf{M_m}^{-1}\mathbf{K_m} & \mathbf{M_m}^{-1}\mathbf{\Phi}^T\mathbf{\Gamma}(\mathbf{\Phi}^{-1})^T\mathbf{B} \end{array}\right] \tag{6.17}$$

The index n is the number of modes, $\mathbf{0}$ is a zero matrix, and $\mathbf{I}$ is the identity matrix.

Note that for a fully actuated system, the term $\mathbf{\Phi}^T\mathbf{\Gamma}(\mathbf{\Phi}^{-1})^T$ reduces to the identity matrix and the closed loop modal system becomes fully decoupled. In that case, the user can tune each k_i to apply arbitrary damping to each mode. In this case however, the term does not reduce to identity. We thus proceed to investigate the resulting damping-induced coupling between modes.

6.4.1 Two Mode Case

The damping limitations imposed by the underactuation are first investigated with the simple case of a two mode system. For such a system, the state space $\mathbf{A_m}$ matrix is

$$\mathbf{A_m} = \left[\begin{array}{c|cc} \mathbf{0}_{2\times 2} & \multicolumn{2}{c}{\mathbf{I}_{2\times 2}} \\ \hline -\mathbf{M_m}^{-1}\mathbf{K_m} & \frac{\Phi_{11}(\Phi^{-1})_{11}k_1}{M_{m11}} & \frac{\Phi_{11}(\Phi^{-1})_{21}k_2}{M_{m22}} \\ & \frac{\Phi_{12}(\Phi^{-1})_{11}k_1}{M_{m11}} & \frac{\Phi_{12}(\Phi^{-1})_{21}k_2}{M_{m22}} \end{array}\right] \tag{6.18}$$

Double numerical subscripts indicate the row and column location of matrix elements. Solving for the characteristic equation associated with matrix $\mathbf{A_m}$ yields

$$\left(s^2 + b_1 s + \omega_1^2\right)\left(s^2 + b_2 s + \omega_2^2\right) - b_1 b_2 s^2 = 0 \tag{6.19}$$

$$\omega_i = \frac{(K_m)_{ii}}{(M_m)_{ii}},\ b_i = \frac{\Phi_{1i}(\Phi^{-1})_{i1}k_i}{(M_m)_{ii}},\ i = 1,\ 2 \tag{6.20}$$

s is the Laplace variable. The roots of Eq. (6.19) are the closed loop poles with damping ratios ζ_{CL} and frequencies ω_{CL}. Since this equation is fourth order it is not easy to solve analytically. Nonetheless, plenty of useful information can be obtained without finding a general solution.

The structure of the equation is that of the product of two second order damped resonant systems subtracted by a term in s^2. The values b_1 and b_2, controlled by the user through k_1 and k_2, act like the damping terms in each damped second order component. They also directly determine the final s^2 term. It is this term that

introduces the cross coupling between the modes. Indeed, if $b_1 = b_2 = 0$ implying no feedback, the equation recovers the open loop poles.

If this extra s^2 term simply did not exist we would have arbitrary control over the amount of damping applied to each mode. Interestingly, arbitrary damping can be achieved on a given mode by setting either b_1 or b_2 to zero. It is only when both modes are damped simultaneously that coupling is introduced. In the single mode damping case, the damping simply acts as a disturbance to the undamped mode. Since the 'modal sensor' sees only its own mode, no information about the disturbed mode influences the damped mode. However, when the second damping loop is closed both modes disturb each other, allowing information to flow back and forth. This two mode analysis may apply to any multi-mode system if only two of the modes are damped.

Another important property is that the closed loop poles of Eq. (6.19) depend directly on the mode frequencies, ω_i, but not the mode shapes, $\mathbf{\Phi}$. Consequently, the limitations on the closed loop damping ratios for a given mode are independent of how much stage 1 participates in that mode. The limitations arise purely due to the mode frequencies. The lack of mode shape dependence arises because $\mathbf{\Phi}$ enters Eq. (6.19) indirectly through the values b. Modes with small coupling to stage 1 tend to decrease the value of b. Since b is also proportional to the arbitrary damping gains k, these small couplings can be compensated for simply by choosing larger values of k. The only theoretical limitation is that the mode have some non-zero coupling to stage 1, i.e. $\Phi_{1i} > 0$. In practice there will be other limitations that prevent k from scaling arbitrarily large, such as actuator range, noise, and higher order dynamics. Realistically, the modes must have some minimal coupling to stage 1 so they are observable enough for the sensors to detect them and controllable enough for the actuators to influence them. Here we will ignore these practical limitations to focus on the theoretical ones.

The figures of merit for these limitations are the closed loop damping ratios $\zeta_{CL,1}$ and $\zeta_{CL,2}$ for modes 1 and 2 respectively. To observe how they respond to two nonzero feedback gains, Eq. (6.19) can be modified to a form more intuitively related to these

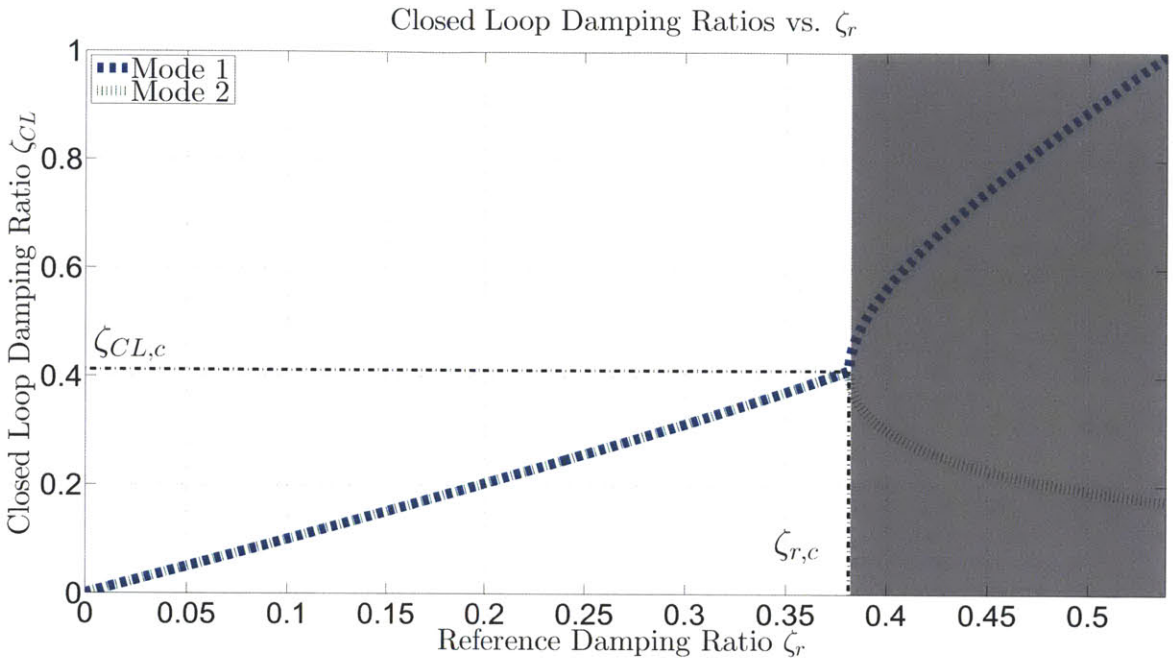

Figure 6-5: The dashed and dotted lines are the closed loop damping ratios as a function of ζ_r. The open loop mode frequencies are 0.440 Hz and 0.982 Hz. The shaded region represents the area beyond the critical point.

damping ratios. Let,

$$\zeta_{r,1} = \frac{b_1}{2\omega_1}, \ \zeta_{r,2} = \frac{b_2}{2\omega_2} \tag{6.21}$$

Where each ζ_r is the reference damping ratio set by the feedback gains. These values are exactly the closed loop damping ratios of the single mode case. When both modes are damped, the closed loop ratios deviate from these reference values.

We explore this deviation using a state space model of the quadruple pendulum in MATLAB® (see Appendix A for details). Modal damping is applied to just the first two modes at 0.440 Hz and 0.982 Hz. We can then vary the values of $\zeta_{r,1}$ and $\zeta_{r,2}$ and observe how the properties of the closed loop poles respond. Good regions of this two dimensional parameter space occur when the dependence of a closed loop damping ratio is dominated by its respective reference damping value.

To reduce the parameter space to a dimension of 1 for the purposes of visualization, let $\zeta_{r,1} = \zeta_{r,2} = \zeta_r$. Figure 6-5 plots the closed loop damping ratios against increasing

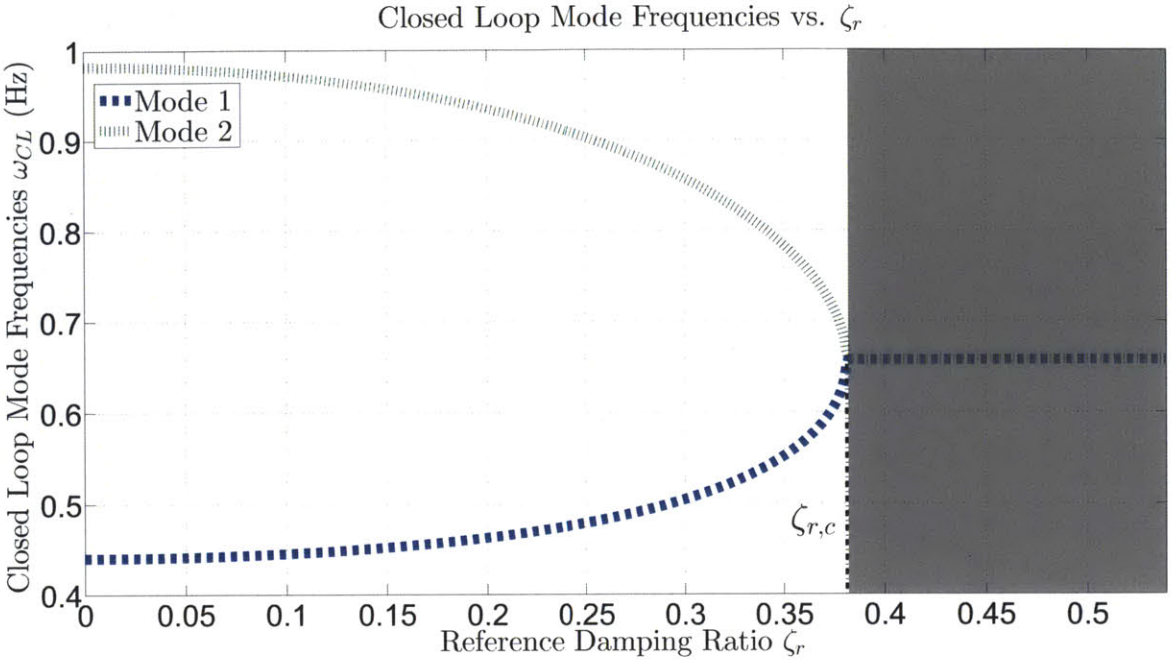

Figure 6-6: The dashed and dotted lines are the closed loop mode frequencies as a function of ζ_r. The open loop mode frequencies are 0.440 Hz and 0.982 Hz. The shaded region represents the area beyond the critical point.

values of ζ_r. The dashed and dotted lines represent these closed loop ratios. Figure 6-6 plots the closed loop mode frequencies against ζ_r. The dashed and dotted lines are these frequencies. Together these two plots reconstruct the information in a root locus diagram.

In both these plots there is a behavioral transition at a critical value $\zeta_r = \zeta_{r,c}$. Up to this point both closed loop damping ratios maintain the same, approximately linear, value and the mode frequencies converge together. Beyond this point, indicated by the shaded regions, the damping ratios diverge and the mode frequencies maintain the same constant value.

At the critical point itself, both modes have exactly the same pole locations and Eq. (6.19) reduces to the squared second order Eq. (6.22). The damping ratios and pole locations are easily solved analytically by multiplying out both equations and comparing like terms. These solutions are listed in Eqs. (6.23) to (6.25). For this example, $\zeta_{r,c} = 0.3812$, $\zeta_{CL,c} = 0.4124$, and $\omega_c = 4.1305\,\text{rad/s} = 0.6574\,\text{Hz}$. The subscript c indicates critical point values.

$$\left(s^2 + 2\zeta_c\omega_c s + \omega_c^2\right)^2 = 0 \tag{6.22}$$

$$\zeta_{r,c} = \frac{\omega_2 - \omega_1}{\omega_2 + \omega_1} \tag{6.23}$$

$$\zeta_{CL,c} = \frac{\omega_2 - \omega_1}{2\sqrt{\omega_2\omega_1}} = \frac{1}{2}\left[\sqrt{\frac{\omega_2}{\omega_1}} - \sqrt{\frac{\omega_1}{\omega_2}}\right] \tag{6.24}$$

$$\omega_c = \sqrt{\omega_1\omega_2} \tag{6.25}$$

Eq. (6.24) is particularly important because it puts a limit on the available closed loop damping while moving along this trajectory in the parameter space. This limit depends only on the open loop mode frequencies and grows as a function of their ratios. In fact, overdamping where $\zeta_{CL,c} \geq 1$ is possible when $\omega_2 \geq \left(3 + 2\sqrt{2}\right)\omega_1$.

The fact that the closed loop damping ratios are roughly proportional to ζ_r up to the critical point implies that there is at least some region in the parameter space where the closed loop damping ratios are similarly proportional to the modal damping gains, k_i. This proportionality is defined by the ratio of the arithmetic mean and geometric mean of the open loop mode frequencies, $\frac{\omega_2+\omega_1}{2\sqrt{\omega_2\omega_1}}$, found by dividing Eq. (6.24) by (6.23). For Figure 6-5 the result is a slope of 1.082 to the critical point.

Other directions in the parameter space do not have similar critical points. However, a generalization can be made that if $\zeta_{r,1}$ and $\zeta_{r,2}$ are less than $\zeta_{r,c}$, each mode's behavior will be dominated by its respective ζ_r. This generalization is tested by Figures 6-7 and 6-8 which expand Figure 6-5 into the full two dimensional parameter space for each closed loop damping ratio. The reference damping values, $\zeta_{r,1}$ and $\zeta_{r,2}$ are restricted to be less than the critical value $\zeta_{r,c}$, and the space within this square region is sampled. The ideal response would be an exactly planar shape where $\zeta_{CL,1} = \zeta_{r,1}$ and $\zeta_{CL,2} = \zeta_{r,2}$. Deviations from the plane indicate coupling between the modes where $\zeta_{CL,1}$ has some dependence on $\zeta_{r,2}$ and $\zeta_{CL,2}$ has some dependence on $\zeta_{r,1}$.

For this example, both plots are nearly planar, where the greatest dependence is by far on each mode's respective reference damping ratio. The greatest deviation

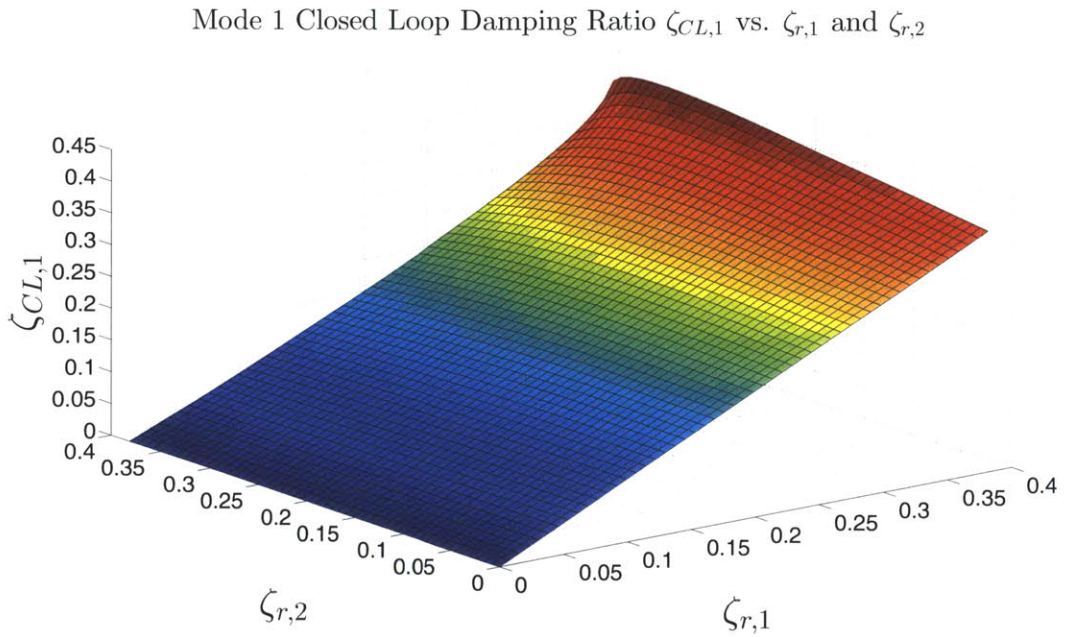

Figure 6-7: The mode 1 closed loop damping ratio $\zeta_{CL,1}$ versus the reference damping ratios for both modes 1 and 2. Both reference damping ratios are less than or equal to the critical value $\zeta_{r,c} = 0.3812$. In this example $\zeta_{CL,c} = 0.4124$.

from the ideal planar response (about 20%) occurs as the reference damping ratios approach the critical value. If smaller deviations are desired to minimize modal coupling further, each ζ_r can be bounded by some fraction of the critical value.

6.4.2 Four Mode Case

In the context of Advanced LIGO, we explore all four x modes of the quadruple pendulum at the frequencies of 0.440 Hz, 0.982 Hz, 1.987 Hz, and 3.394 Hz. However, one can in principal study a system with any number of modes.

Figures 6-9 and 6-10 repeat the analysis in Figures 6-5 and 6-6 for all four modes. The four reference damping ratios are increased equally by scaling the damping gains, k_i, and the closed loop damping ratios and frequencies are plotted.

This four mode case has both similarities and differences to the two mode case. The most obvious difference is that there is no longer a critical point. Instead, there is a gradual divergence of the damping ratios. The mode frequencies, while they tend towards each other, do not reach a common value. The damping ratios are similar to

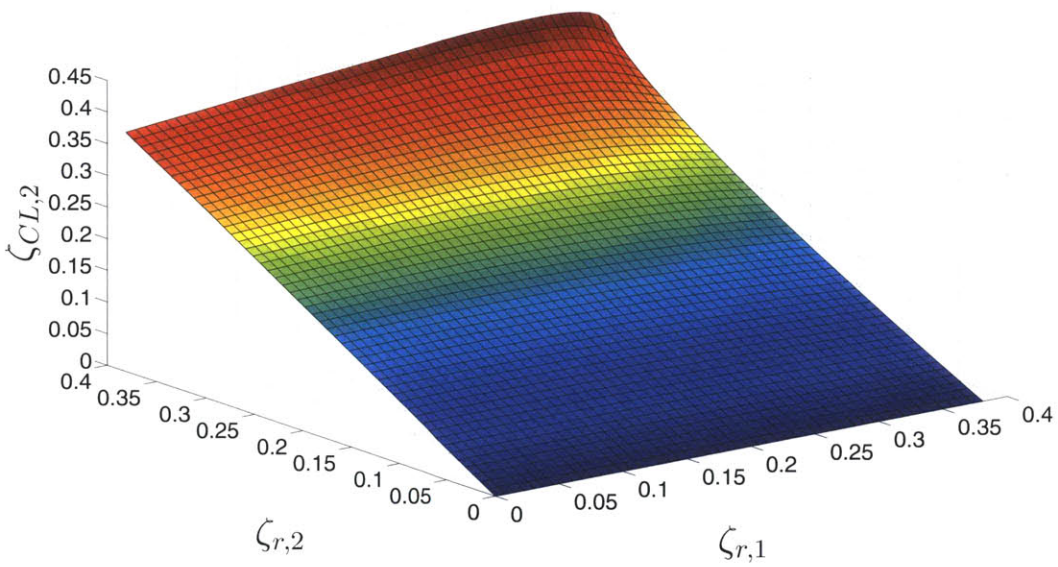

Figure 6-8: The mode 2 closed loop damping ratio $\zeta_{CL,2}$ versus the reference damping ratios for both modes 1 and 2. Both reference damping ratios are less than or equal to the critical value $\zeta_{r,c} = 0.3812$. In this example $\zeta_{CL,c} = 0.4124$.

the two mode case in that only one mode reaches arbitrarily large damping while the others tend to zero after an initial increase.

Since there is no clear critical point in the four mode case (or more generally beyond the two mode case), one is chosen to divide the parameter space into 'good' and 'bad' regions. The first point where one of the damping ratios achieves zero slope is chosen as this value. The fourth mode here is the first to level off with zero slope at $\zeta_r = 0.203$. This value is indicated in Figure 6-9 by the vertical dotted line labeled $\zeta_{r,c}$. The area beyond the chosen critical value is shaded to indicate the 'bad' region where too much coupling exists between damped modes.

Following the analysis of the two mode case, Figures 6-11 to 6-14 sample the damping response of each mode over an entire region of the parameter space within the hypercube where each ζ_r is less than $\zeta_{r,c}$. The right horizontal axis of these plots is one of the reference damping ratios and the vertical axis is the respective closed loop damping ratio. To compress all the information into a format that can be plotted in 3 dimensions, the left horizontal $\zeta_{r,2,3,4}$ axis combines the remaining three reference damping ratios. For example, Figure 6-11 plots the damping of mode 1. For every $\zeta_{r,1}$,

the left horizontal axis considers all values within the cube defined by $\zeta_{r,2} \leq \zeta_{r,2,3,4}$, $\zeta_{r,3} \leq \zeta_{r,2,3,4}$, $\zeta_{r,4} \leq \zeta_{r,2,3,4}$. The closed loop damping value that deviates from $\zeta_{r,1}$ the most is the value plotted on the vertical axis. The deviation is quantified as $|\frac{\zeta_{CL,1}}{\zeta_{r,1}} - 1|$. This logic is given in equation form for Figure 6-11 by Eq. (6.26) subject to the constraint (6.27). Figures 6-12 to 6-14 follow similar logic.

$$\zeta_{CL,1}(\zeta_{r,1}, \zeta_{r,2,3,4}) = \underset{\zeta_{CL,1}(\zeta_{r,2},\zeta_{r,3},\zeta_{r,4})}{\operatorname{argmax}} \left| \frac{\zeta_{CL,1}}{\zeta_{r,1}} - 1 \right| \tag{6.26}$$

$$\text{s.t. } \zeta_{r,2} \leq \zeta_{r,2,3,4},\ \zeta_{r,3} \leq \zeta_{r,2,3,4},\ \zeta_{r,4} \leq \zeta_{r,2,3,4} \tag{6.27}$$

As in the two mode case, the ideal response for all modes would be on the hyperplane $\zeta_{CL,i} = \zeta_{r,i}$, where i is the index of the mode. The overall response here is very similar to the two mode case. For small damping ratios the deviations from this hyperplane are very small. The deviations increase with the size of the damping ratios. For the maximum value of $\zeta_{r,c} = 0.203$ here, the greatest deviation is on mode 3 at about 35%. Mode 1 has the smallest maximum deviation at 12%.

The damping requirements state that stage 4 must ringdown from an impulse to e^{-1} (36.8%) of the peak in 10 s. This decay rate corresponds to a damping ratio of

$$\zeta_{CL,i} = \frac{1}{10(2\pi f_i)} \tag{6.28}$$

where f_i is the resonant frequency in Hz for mode i. If the ζ_r values are set to these desired closed loop damping ratios for the quadruple pendulum, the resulting values and corresponding closed loop damping ratios are

$$\zeta_{r,1} = 0.0362 \qquad \zeta_{CL,1} = 0.0364 \tag{6.29}$$

$$\zeta_{r,2} = 0.0162 \qquad \zeta_{CL,2} = 0.0163 \tag{6.30}$$

$$\zeta_{r,3} = 8.00 \times 10^{-3} \qquad \zeta_{CL,3} = 8.05 \times 10^{-3} \tag{6.31}$$

$$\zeta_{r,4} = 4.69 \times 10^{-3} \qquad \zeta_{CL,4} = 4.71 \times 10^{-3} \tag{6.32}$$

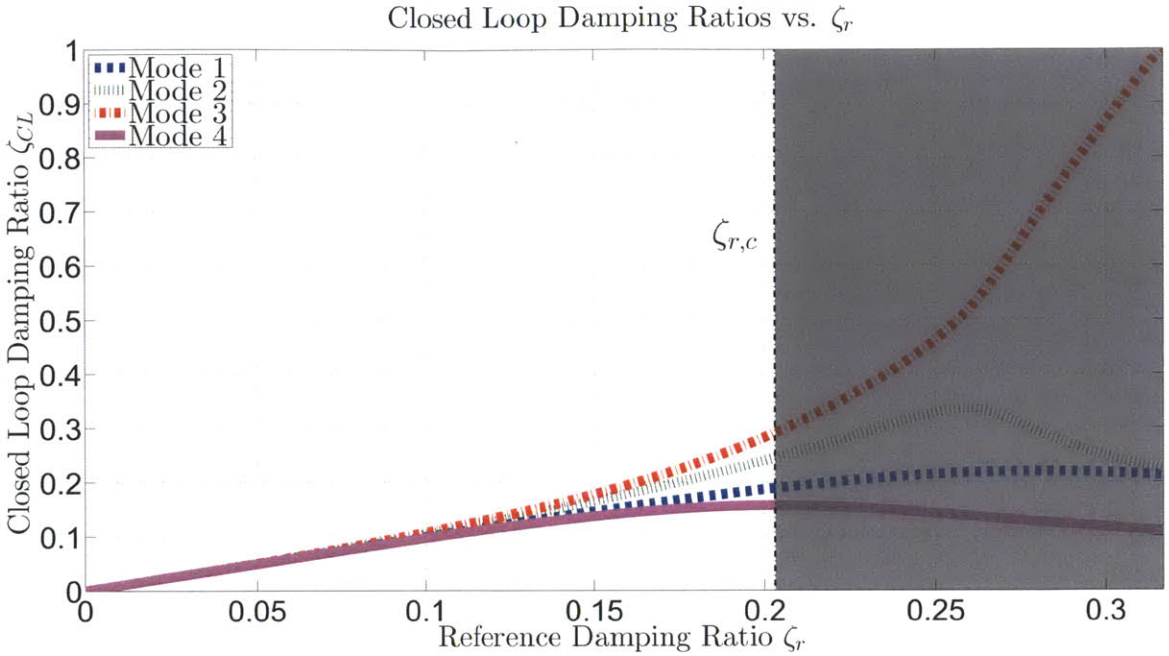

Figure 6-9: The dashed, dotted, and thick solid lines are the closed loop damping ratios as a function of ζ_r. The open loop mode frequencies are 0.440 Hz, 0.982 Hz, 1.9873 Hz, and 3.3942 Hz.

All damping ratios are well below the critical value, allowing the relation $\zeta_r \approx \zeta_{CL}$ to hold true. Here the deviation between ζ_r and ζ_{CL} is only about 0.5% for each mode.

6.4.3 Implementing Realistic Feedback

Sections 6.4.1 and 6.4.2 employed control filters $G_i = 1$ in the damping analysis. Practical applications require filters with some amount of lowpassing to filter noise, prevent instability from higher order dynamics, and avoid actuator saturations. The primary concern for the quadruple pendulum case is the noise filtering issue. Section 2.1 stated that at 10 Hz stage 4's motion due to sensor noise must be no greater than $10^{-20}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. Since our goal is to apply damping to resonances up to about 3.5 Hz, and since the noise of our sensors is relatively high, rather aggressive lowpass filters are required.

As a consequence of these requirements and limitations, the feedback filters employed in reality deviate from those ideal filters in the previous sections. Nonetheless,

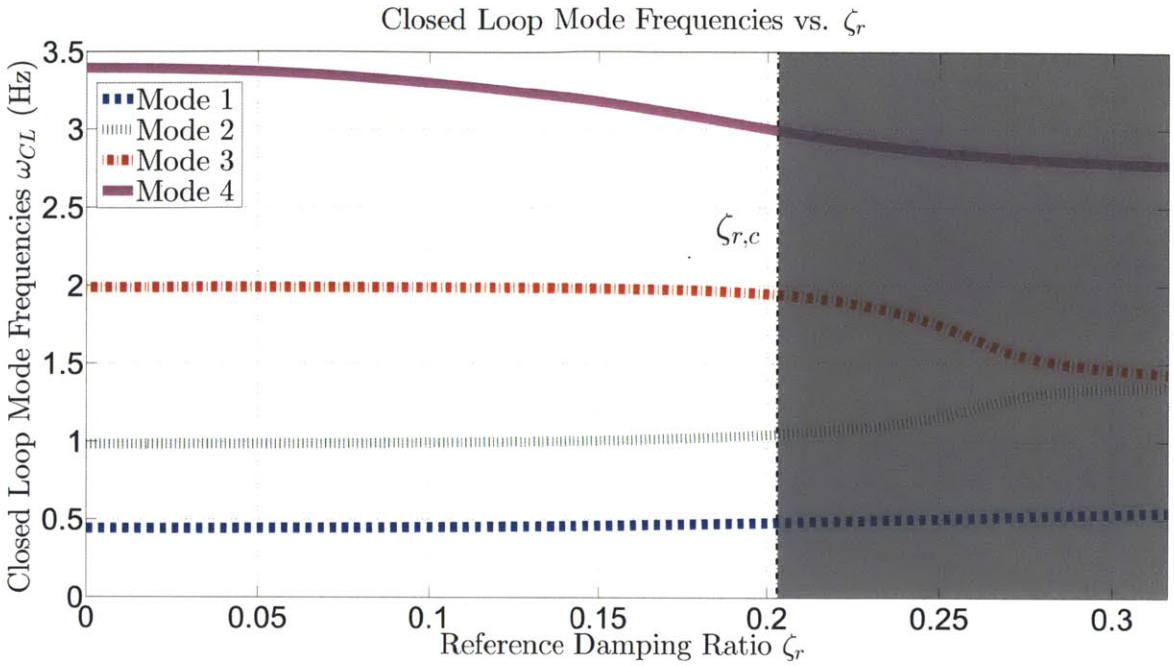

Figure 6-10: The dashed, dotted, and thick solid lines are the closed loop mode frequencies as a function of ζ_r. The open loop mode frequencies are 0.440 Hz, 0.982 Hz, 1.9873 Hz, and 3.3942 Hz.

practical filters can be employed that produce similar responses provided two conditions are met. Those conditions are: the gain of the filter, G, at the resonant frequency equals unity; and there are sufficient phase and gain margins. The design of the filters in Section 6.2 meet these requirements (when properly scaled).

Figure 6-15 illustrates an impulse response to the quadruple pendulum with these feedback filter designs. The solid line shows the stage 4 response when the ideal unity gain feedback filters are employed. The dashed line shows the response when the more practical design from Section 6.2 is employed. The unity filter response meets the damping requirement well since it falls below the dotted $\frac{1}{e}$ amplitude lines just before 10 s. The response with the practical filter is very similar. It does not perform quite as well, but its last crossing of the $\frac{1}{e}$ lines still occurs by 10 s.

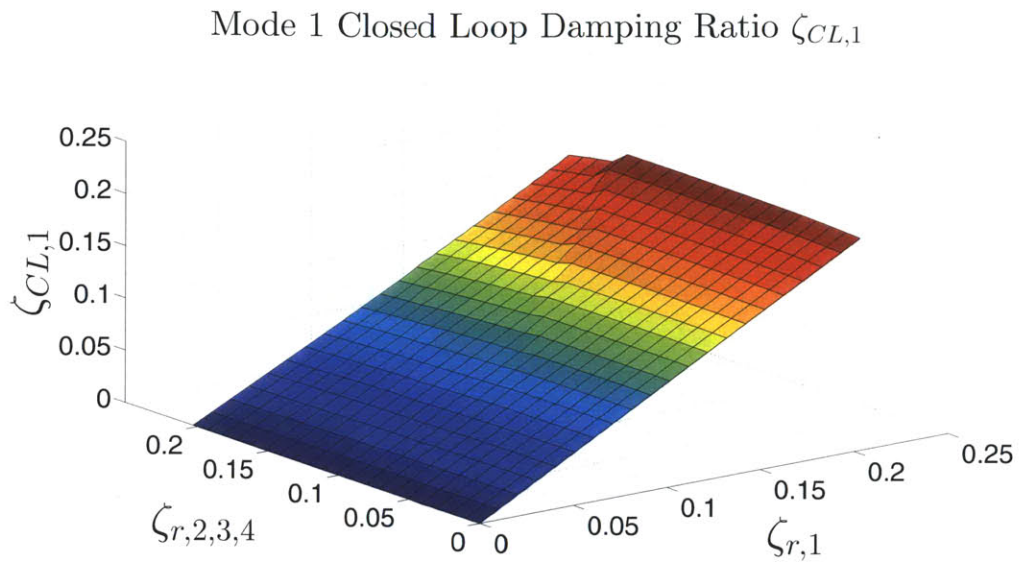

Figure 6-11: The mode 1 closed loop damping ratio $\zeta_{CL,1}$ versus the reference damping ratios for modes 1, 2, 3, and 4. The $\zeta_{r,2,3,4}$ axis considers all $\zeta_{CL,1}$ values at $\zeta_{r,1}$ within the cube defined by $\zeta_{r,2} \leq \zeta_{r,2,3,4}$, $\zeta_{r,3} \leq \zeta_{r,2,3,4}$, $\zeta_{r,4} \leq \zeta_{r,2,3,4}$. The $\zeta_{CL,1}$ value plotted represents the largest deviation from $\zeta_{r,1}$, where the deviation is measured as $|\frac{\zeta_{CL,1}}{\zeta_{r,1}} - 1|$. The largest deviation over the entire plot is 12%.

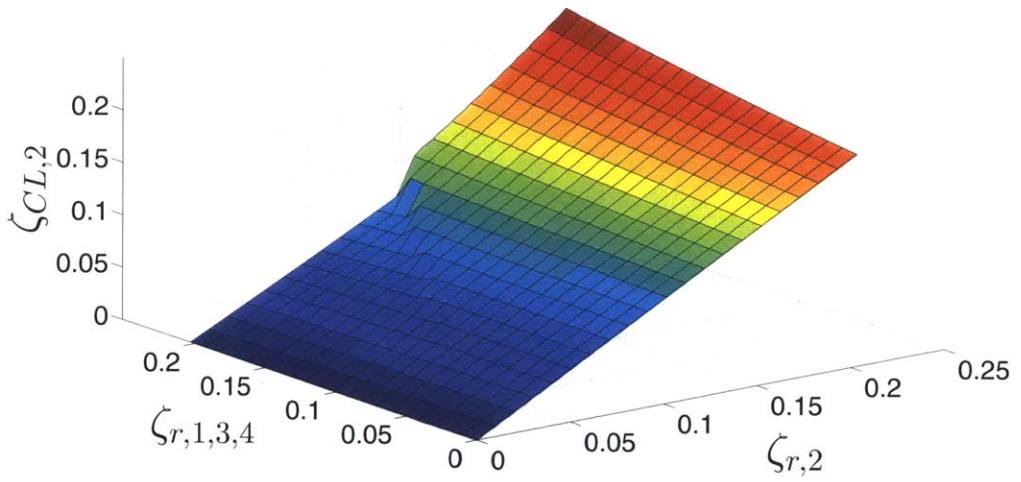

Figure 6-12: The mode 2 closed loop damping ratio $\zeta_{CL,2}$ versus the reference damping ratios for modes 1, 2, 3, and 4. The $\zeta_{r,1,3,4}$ axis considers all $\zeta_{CL,2}$ values at $\zeta_{r,2}$ within the cube defined by $\zeta_{r,1} \leq \zeta_{r,1,3,4}$, $\zeta_{r,3} \leq \zeta_{r,1,3,4}$, $\zeta_{r,4} \leq \zeta_{r,1,3,4}$. The $\zeta_{CL,2}$ value plotted represents the largest deviation from $\zeta_{r,2}$, where the deviation is measured as $|\frac{\zeta_{CL,2}}{\zeta_{r,2}} - 1|$. The largest deviation over the entire plot is 15%.

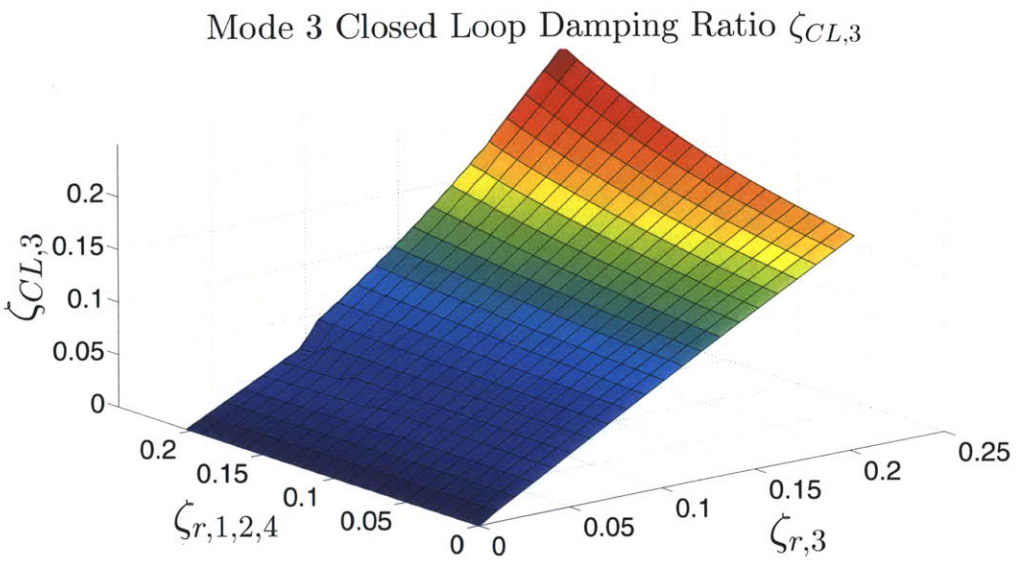

Figure 6-13: The mode 3 closed loop damping ratio $\zeta_{CL,3}$ versus the reference damping ratios for modes 1, 2, 3, and 4. The $\zeta_{r,1,2,4}$ axis considers all $\zeta_{CL,3}$ values at $\zeta_{r,3}$ within the cube defined by $\zeta_{r,1} \leq \zeta_{r,1,2,4}$, $\zeta_{r,2} \leq \zeta_{r,1,2,4}$, $\zeta_{r,4} \leq \zeta_{r,1,2,4}$. The $\zeta_{CL,3}$ value plotted represents the largest deviation from $\zeta_{r,3}$, where the deviation is measured as $|\frac{\zeta_{CL,3}}{\zeta_{r,3}} - 1|$. The largest deviation over the entire plot is 35%.

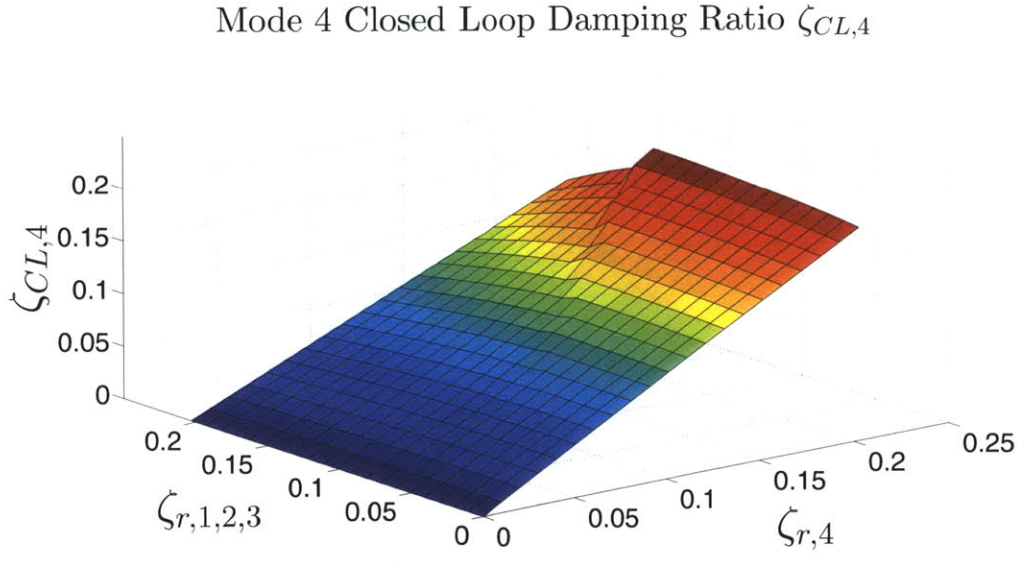

Figure 6-14: The mode 4 closed loop damping ratio $\zeta_{CL,4}$ versus the reference damping ratios for modes 1, 2, 3, and 4. The $\zeta_{r,1,2,3}$ axis considers all $\zeta_{CL,4}$ values at $\zeta_{r,4}$ within the cube defined by $\zeta_{r,1} \leq \zeta_{r,1,2,3}$, $\zeta_{r,2} \leq \zeta_{r,1,2,3}$, $\zeta_{r,3} \leq \zeta_{r,1,2,3}$. The $\zeta_{CL,4}$ value plotted represents the largest deviation from $\zeta_{r,4}$, where the deviation is measured as $|\frac{\zeta_{CL,4}}{\zeta_{r,4}} - 1|$. The largest deviation over the entire plot is 23%.

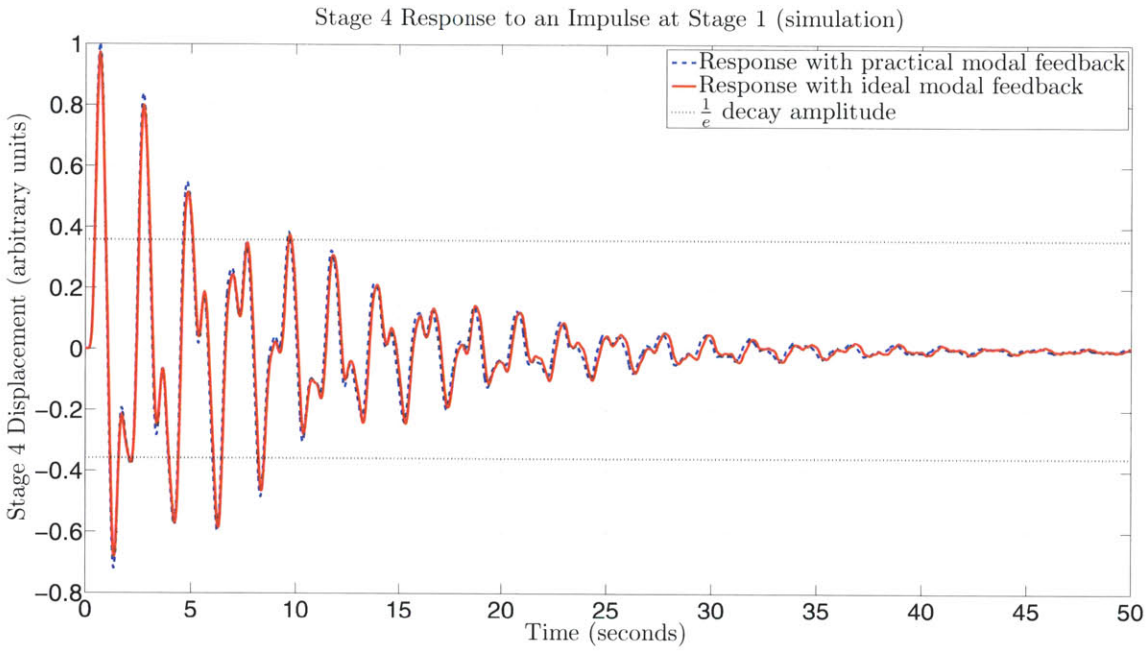

Figure 6-15: The response of stage 4 to an impulse at stage 1 along the x axis with modal damping engaged. The damping values are set to the values listed in Eqs. (6.29) to (6.32) so that each mode damps to $\frac{1}{e}$ in 10 s. The solid red line is the response with the ideal unity feedback filters. The dashed blue line is the response with the filters designed in Section 6.2. The dotted lines represent the required $\frac{1}{e}$ decay amplitude for the practical feedback.

6.5 Conclusion

Modal damping provides a convenient way to decouple the dynamics of a complex many degree of freedom system into many second order single degree of freedom systems. The damping can then be tuned for each mode independently to apply the minimum required damping, minimizing the sensor noise amplification within the gravitational wave detection band above 10 Hz. The design of the estimator used to reconstruct the incomplete sensor information is found by solving a customized cost function that optimizes the overall trade-off between test mass ringdown time and displacement within the GW frequency band.

Another important advantage of modal damping is the simplicity of tuning the system performance. The decoupled nature of each modal damping loop means that more or less damping for each mode can be chosen simply by adjusting the respective loop gain. For example, if it is desired at a later time to reduce the sensor noise

amplification, the overall loop gain on the highest frequency mode (a single parameter, adjustable in real-time) can simply be reduced, rather than redesigning the damping control for the entire system.

In practice many systems are underactuated, a property that degrades the otherwise ideal decoupling of a system's modes. The analysis in this chapter determines a number of important properties regarding the damping available to a multi-degree of freedom system when actuation is limited to a single degree of freedom.

One of these properties states that arbitrary amounts of damping can be applied to only one mode at a time. Nonetheless, all modes with nonzero participation at the actuated degree of freedom can be damped approximately independently if the desired damping ratios are kept below some maximum value. The smaller the maximum value, the greater the independence. This maximum value is purely limited by the mode frequencies. Modes that are proportionally farther apart in frequency have higher maximum damping limits. If modes are separated enough, the maximum limit may go beyond critical damping.

The fact that mode frequencies govern the damping limitation indicates another, less intuitive property, that the mode shapes play no role. In practice however, the mode shapes are important because the participation of a mode along the actuated degree of freedom determines its observability and controllability. A mode must be sufficiently controllable for damping to be practical. Otherwise, excessively large control forces risk saturating the actuators, amplifying noise, and interacting with higher order dynamics. At the limit where a mode's participation factor goes to zero, it becomes completely inaccessible to damping.

A final important point is that the poles of undamped modes are isolated from those that are damped. Modes must have simultaneous active damping for them to couple together. At the limit where all modes are uncontrolled except one, the system reduces to a fully decoupled second order system. In fact, any mode (provided it has non-zero coupling to the actuated degree of freedom) can be arbitrarily damped simply by turning off the control applied to all other modes.

This modal damping technique is applied to an Advanced LIGO quadruple pen-

dulum to damp the four modes that couple strongly to the axis of the interferometer. It is shown that the damping ratios required to meet the necessary decay rates are well within the range of what is available when all modes are damped simultaneously. Practical feedback with significant lowpass filtering is also tested, breaking the ideal assumption of scalar feedback in the preceding analysis. These filters are shown to function similarly to the ideal case provided they have similar gains at the modal frequencies and sufficient stability margins.

Chapter 7

Adaptive Modal Damping

This chapter extends the modal damping design from Chapter 6 by building an adaption algorithm around it. This adaptation will provide real-time optimization of the quadruple pendulum damping as the disturbances from the environment continuously evolve. Section 7.1 lists the challenges that motivate this adaption. Section 7.2 maps the top-level basic components of a complete adaptive modal damping control system. The cost function component of this system is described by Section 7.3, while the adaption kernel is given by Section 7.4. An experimental setup to test adaptive modal damping on a prototype pendulum is described in Section 7.5. Section 7.6 presents the results from this experiment.

7.1 Challenges

The active damping design must compete with three challenges. These challenges are sensor noise, non-stationary disturbances, and nonlinear interferometer behavior. Meeting the sensor noise challenge is essential to satisfying the requirements in the previous chapter. The other two challenges limit the sensitivity of the interferometer at various times.

The most significant form of electronic noise in the damping feedback loop is sensor noise. This noise term has been measured to $10^{-10}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ for each OSEM beyond 10 Hz [37]. The consequence is that the damping loop gain for the stage 1

x DOF must be less than 5×10^{-6} at 10 Hz. However, the highest frequency mode along that DOF occurs at 3.4 Hz, about half a decade away. Therefore, the loop gain transfer function must drop at least five orders of magnitude in this space if any damping is to occur at that mode. This task poses a challenge for stable feedback loop designs.

The next challenge is the variable nature of the seismic disturbances impinging on the interferometer. These disturbances constantly evolve in time at many different time scales and frequencies over large ranges of amplitudes. Examples of the four dominant frequency bands of these disturbances are plotted in Figure 7-1. This figure shows the actual measured ground vibration at the Livingston, LA observatory over a 24 hour period on November 24, 2009. The measurements were taken at three locations: the ends of the two arms (X end and Y end), and the vertex station where the arms meet. Each frequency band tends to represent a different disturbance class. The lowest frequency band from 0.03 Hz to 0.1 Hz in the top left plot represents actual seismic activity such as earthquakes. The next band from 0.1 Hz to 0.35 Hz in the bottom left plot represents weather disturbances. The large peak is due to a passing storm, which dominates the total disturbance on this day. The storm in fact prevented the interferometer from operating for at least two full days. The next two bands on the right from 1 Hz to 10 Hz represent anthropogenic disturbances, i.e. created by humans. These disturbances have a diurnal cycle, picking up during the day as people begin the work day and dropping off in the evening as people return home. The weekends are typically quieter than the week. The largest spikes in the two anthropogenic plots represent trains passing near the end of one of the arms. These spikes often caused the Initial LIGO interferometer control to go unstable for a few minutes until the train passed, except in this case the interferometer was already out of commission due to the storm.

The final challenge to active damping is the nonlinear response of the interferometer. At small amplitudes of motion, the interferometer response is very linear. However, nonlinear terms become visible in its response when the system oscillates with large amplitudes. Most of these nonlinearities are either unsensed or poorly

understood. Consequently, it is not possible to subtract them from the system's output. These nonlinearities cause the seismic disturbances to upconvert to frequencies beyond 10 Hz where LIGO searches for GWs. Some of the suspected sources of these nonlinearities come from laser light scattering off the vacuum system walls, laser beams falling off optics, higher order dynamics in the isolation systems, creak (sliding dislocations) in the materials of the isolation systems, nonlinear actuator responses, and the near quadratic response of the interferometer output with test mass displacement. Large enough oscillations will ultimately prevent the interferometer from operating.

These three challenges couple together in a way that demands adaptive tuning of the damping feedback to respond to changing environmental conditions. During times of large ground vibrations, large feedback gains are desirable to suppress the disturbances and reduce seismic upconversion at the expense of sensor noise amplification. During times of the largest amplitude seismic events, these high gains help keep the interferometer operational where it would not be otherwise. During times of small ground vibrations it is better to reduce the feedback gains to the basic $\frac{1}{e}$ in 10 s damping requirement to minimize the sensor noise amplification.

7.2 Overview

The block diagram in Figure 7-2 provides a top level basic view of the signal flow for the adaptive modal damping control system. This diagram represents a general case in that the interferometer control is not necessarily applied to all quadruple pendulums, though all quadruple pendulums may have modal damping. Two of the three challenges are shown in this diagram, the variable seismic disturbances and the stage 1 sensor noise. The third challenge of nonlinear upconversion of the seismic disturbances has no unique source, so it is not shown.

Moving right to left in the diagram, the modal damping developed in Chapter 6 is applied to stage 1 for modes that couple strongly to the cavity axis. For the quadruple pendulum there are at most four of these modes. There may be fewer modes if the

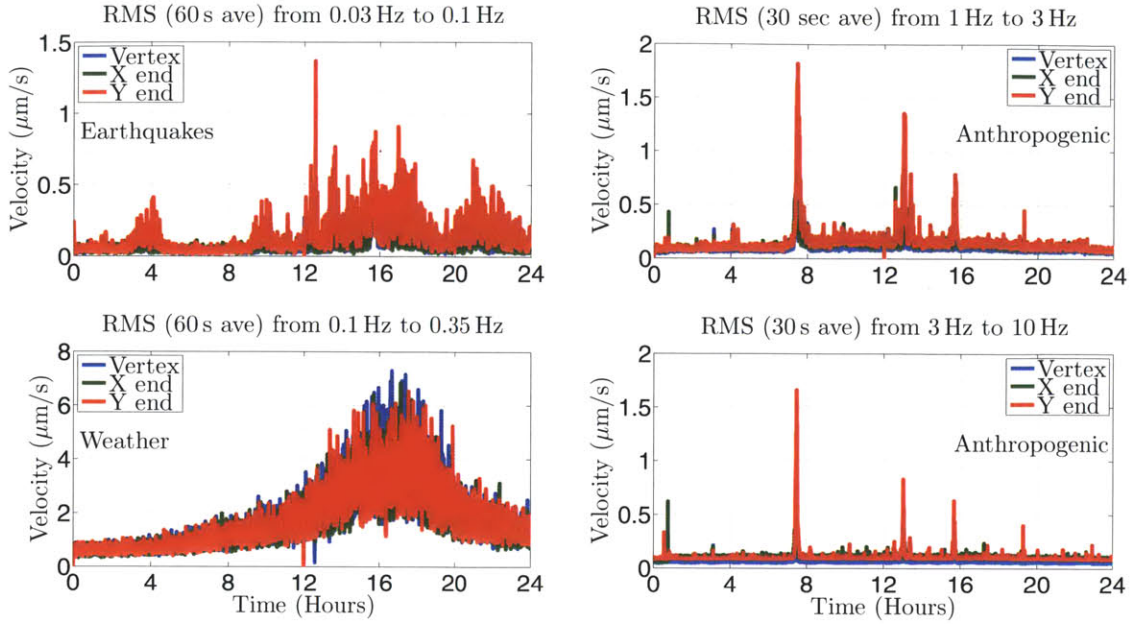

Figure 7-1: Seismic disturbances at the Livingston, LA observatory, Nov. 21, 2009 measured with commercial Streckeisen STS-2 seismometers. Courtesy LIGO Lab.

bandwidth of the interferometer control applied to the pendulum is higher than the frequencies of these modes. In this case, the interferometer control forces some of the stages to follow the interferometer more than the pendulum itself, effectively removing DOFs from the pendulum dynamics seen by the top stage.

The adaption box then sets the modal damping feedback gains. The feedback filters themselves remain constant, so the quantity of damping is tuned simply by scaling the output of the filters with the gains k_i. Two adaptive gains are shown in this diagram. The adaptation box uses a Gauss-Newton algorithm to decide which way and how much to step these gains. Section 7.4 details this algorithm.

The cost box provides the adaptation with the information it needs to make these decisions. This information is organized in the form of two competing components of a cost function that the adaptation box attempts to minimize for each mode. Section 7.3 below provides the details of this box.

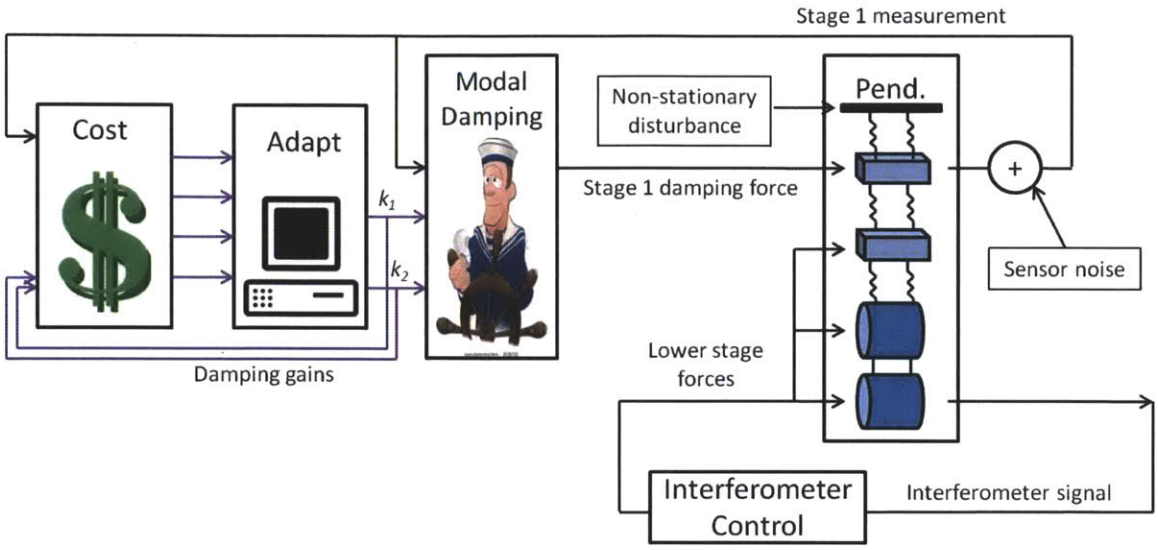

Figure 7-2: Top level block diagram of adaptive modal damping applied to a quadruple pendulum. The interferometer control is included for reference.

7.3 Cost Box

Figure 7-3 visualizes the flow of the cost information by peering one level down inside the cost box. This flow is repeated for each modally damped mode. Following this diagram, the two competing parts of the cost function follow separate paths. The upper path determines the cost for the amplitude of a modal signal. The bottom path determines the cost of applying feedback with a noisy sensor.

7.3.1 Dual Cost Paths

The upper path is divided into three steps. At the first step, a narrow bandpass filter applied to the stage 1 sensor signal selects one of the modal responses. The amplitude of the output of this filter is proportional to the amplitude of the selected modal response. The bandpass filters are implemented with MATLAB®. Figure 7-4 displays two examples of this filter. A simulated undamped example amplitude spectrum of stage 1 with two modes is superimposed on the plot for reference. Each filter has an elliptical bandpass component designed using the 'ellip' function. For each mode the call to the function is 'ellip(6,1,60,2*pi*f_i*[0.925 1.075],'bandpass','s')'. That is a sixth order filter with 1 dB of passband ripple and a stopband 60 dB down from the

passband. The corners of the bandpass occur 7.5% below and above the frequency of the mode, f_i in units of Hz, resulting in a pass bandwidth about 15% percent of the mode frequency. This width is enough to capture the peak frequency within all reasonable damping ratios. The largest reasonable damping ratio is established as 0.2 by Section 6.4. Beyond 0.2 the modal damping can no longer be considered independent for each mode due to the stage 1 limited actuation. Note that this function call outputs a filter with 12 poles rather than 6 because both the lowpass and highpass components are each sixth order. Each bandpass has an additional lowpass and highpass component implemented with two zeros at 0 Hz and four poles at the mode frequency f_i. These additional components force the filter to roll off at 40 dB per decade away from the modal frequency. The overall bandpass filter is set so that the gain at f_i is unity.

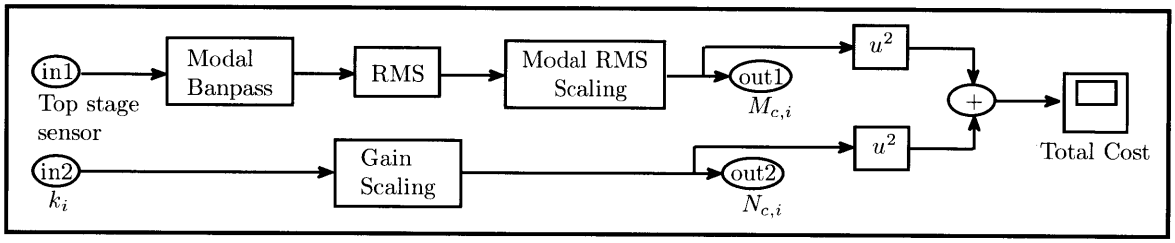

Figure 7-3: Signal flow block diagram inside the cost box.

The amplitude of the bandpass output is measured by an RMS filter in the second step, shown in detail by Figure 7-5. The bandpass filtered signal enters the RMS filter and the value of each sample immediately gets squared. This squared signal then goes through a lowpass filter. The output of this lowpass filter approximates the average of the squared signal. As a result, the lowpass output gives a real-time approximation of the variance of the modal signal. A following square root operation converts this variance into the RMS of the modal amplitude as seen by the top stage sensor.

The RMS lowpass filter is simple in structure but involves a significant design trade-off. The filter has just two poles and no zeros. The poles are set to an angular frequency of ω_c. The gain of the filter is scaled to be unity at zero frequency. The design trade-off involves the choice of ω_c. Smaller ω_c corresponds to more precise averaging. However, the longer time constant also results in a slower response. If ω_c

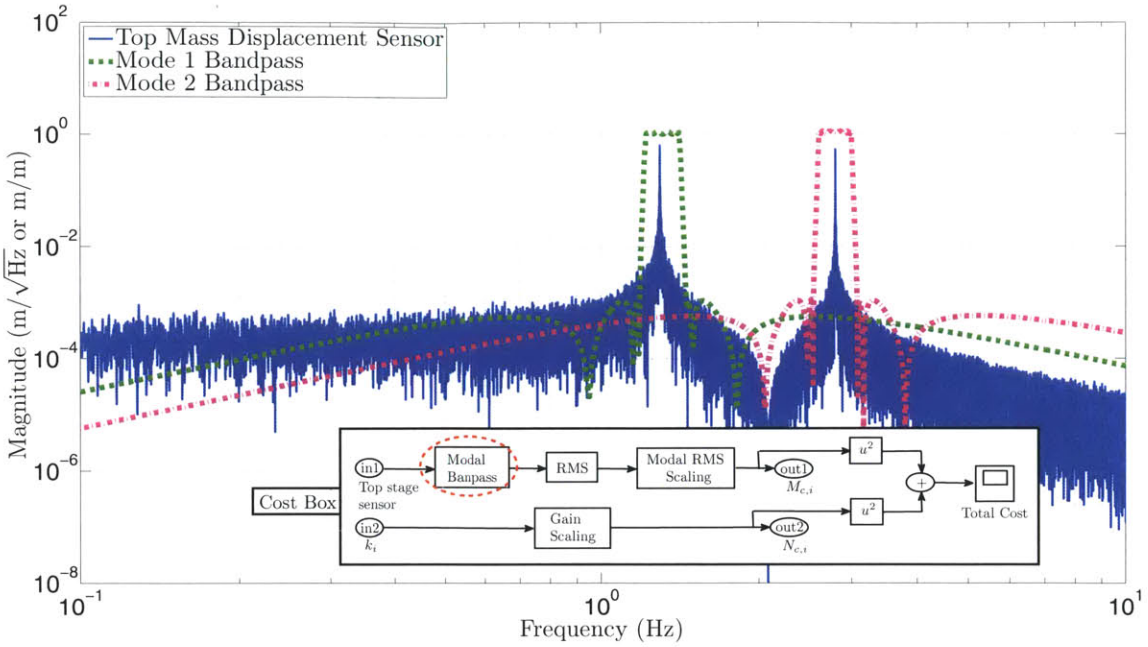

Figure 7-4: Illustration of the modal bandpass filters. The blue line in the background is a simulation of an example stage 1 displacement spectrum for a two mode system. The magnitude of this line is arbitrary. The dotted and dashed lines are the unitless bandpass filter transfer functions.

is small enough, large sudden disturbances will be missed. Larger ω_c results in faster responses that will capture sudden disturbances, however the filter will output poor quality averages. Quantifying this trade-off, the time constant of the RMS response is inversely proportional to ω_c whereas the precision of the RMS output is proportional to the square root of ω_c. Figure 7-5 displays this trade-off. The blue line is an example of the output of the modal bandpass for a 1.3 Hz mode. The black dashed line is the RMS filter output when $\omega_c = 0.1\,\text{rad/s}$, corresponding to a time constant of 62.8 s. The red solid line is the RMS filter output when $\omega_c = 0.025\,\text{rad/s}$, corresponding to a time constant of 251.3 s. Since the red solid has an ω_c that is 4 times smaller it rises 4 times slower, but has twice the precision.

The final box in the upper path then scales this measured modal amplitude according to the chosen shape of the cost function. A discussion of this scaling and the cost function shape is found in Section 7.3.2.

The lower path determines the cost of applying active damping with a noisy sensor. This path has just one single step because the sensor noise from the OSEMs is known

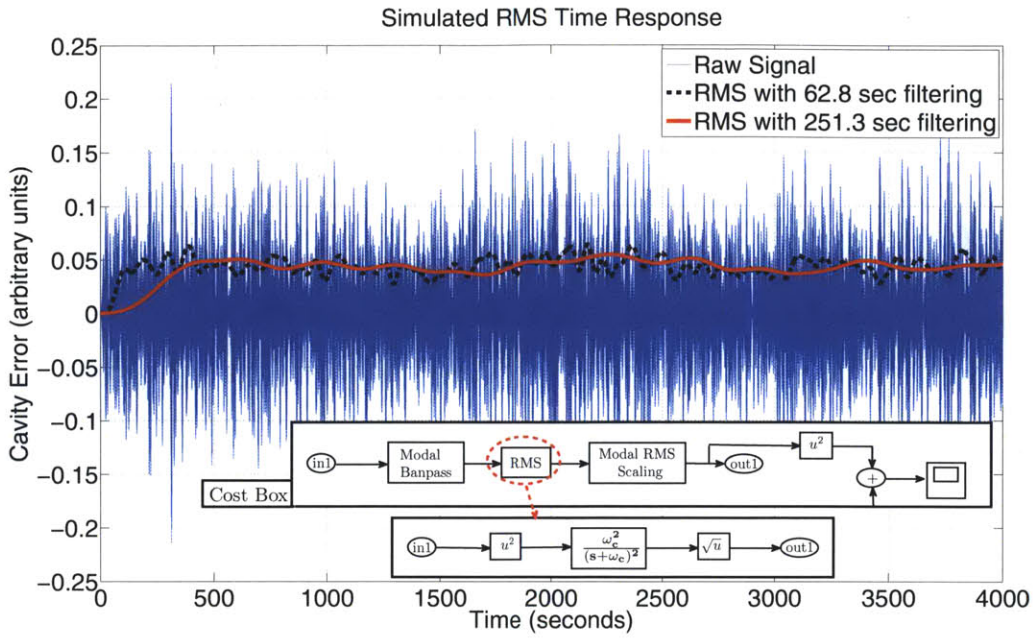

Figure 7-5: Simulated example of the response of an RMS filter with different time constants. The blue trace in the background is the input to the filter, the dashed black and red solid lines are the outputs with different time constants.

and constant and the pendulum is well modeled. As a result, the influence of this sensor noise on the cavity signal can be predicted. This prediction can be updated immediately at every sample time where a similar quality measurement of the noise amplification would require many samples. The transfer functions in Eqs. (7.1) and (7.2) formulate this prediction for the damping of mode i. x_4 is the motion of stage 4 along the cavity axis, ν is the sensor noise, C_i is the combined estimator and modal feedback filter transfer function for mode i, Q_{41} is the quadruple pendulum transfer function from stage 1 to stage 4 along the cavity axis, Q_{11} is the quadruple pendulum transfer function from stage 1 to stage 1 along the cavity axis, $(\phi^{-1})_{i1}$ is the mode i element of the modal to Cartesian coordinate transformation, $E_{s,i}$ is the estimator transfer function between the sensor signal and mode i, $E_{a,i}$ is the estimator transfer function between stage 1 actuation and mode i, G_i is the feedback filter for mode i, and k_i is the adaptive gain factor on G_i.

$$x_4 = \frac{C_i Q_{41}}{1 + C_i Q_{11}} \nu \tag{7.1}$$

$$C_i = \frac{k_i (\phi^{-1})_{i1} G_i E_{s,i}}{1 + k_i (\phi^{-1})_{i1} G_i E_{a,i}} \tag{7.2}$$

Eq. (7.3) states that when the loop gain $C_i Q_{11}$ for mode i is small, then the displacement of the test mass is directly proportional to the adaptive scale value k_i. In this regime, the feedback simply amplifies sensor noise. 10 Hz where the noise requirement begins is always within this regime because each unity gain crossing tends to stay near the modal frequencies. The 'Gain Scaling' step in the lower path takes advantage of this relation and simply calculates the sensor noise cost directly from the current value of k_i. Section 7.3.2 discusses this scaling.

$$x_4 \propto k_i, \quad \mathrm{C_i Q_{11}} < 1 \tag{7.3}$$

7.3.2 Cost function

The shapes of the cost function components are implemented by the 'Gain Scaling' step in the lower path of the cost box and the 'Modal RMS Scaling' step in the upper path. Figure 7-6 illustrates an example of these cost components as functions of the damping gain k for a given level of seismic disturbance D. The dashed green line is the sensor noise component implemented by the 'Gain Scaling' step and the solid blue line is the modal amplitude component implemented by 'Modal RMS Scaling' step. The dotted red line is the overall cost function, which is the sum of the squares of each component.

The sensor noise cost component for each mode i follows Eq. (7.4). This cost component is chosen to be linear in the damping gain because the influence of sensor noise is also linear in the damping gain, as was shown by Eqs. (7.1) to (7.3). A zero crossing for this linear cost is chosen at the gain value of $k_{0,i}$, which is set to meet the noise amplification requirements in Section 2.1. This value was chosen to have zero

cost because there is no reason for the damping gain to be smaller than $k_{0,i}$ if sensor noise is not an issue. In practice, the adaptive algorithm is constrained to $k_i \geq k_{0,i}$.

The modal amplitude component for each mode i follows Eqs. (7.5) to (7.7). $M_{RMS,i}$ is the true RMS modal amplitude at any given time. Eq. (7.5) models this value as being proportional to the unknown amplitude of the seismic disturbance D_i and inversely proportional to the square root of the damping gain k_i. The cost function chosen to weight this modal amplitude is given by Eq. (7.6). The value $M_{0,i}$ is a value chosen by the user to scale the importance of the modal amplitude relative to the sensor noise. $erf()$ is the error function which equals 0 for an argument of 0 and asymptotically approaches 1 as the argument approaches infinity. The cost function is structured such that for small modal amplitudes the denominator approaches 1. In this regime the modal amplitude cost is linearly proportional to the modal amplitude. Since both cost components scale linearly, there is a symmetric trade-off in the optimization between modal amplitude and sensor noise amplification. Large modal amplitudes enter a regime where the denominator begins approaching zero and the modal amplitude cost scales faster than a linear function. This nonlinear increase puts more weight on suppressing large modal amplitudes to prevent the nonlinearities in the interferometer from dominating the response. At the largest modal amplitudes, the modal cost approaches infinity and the adaptive algorithm ignores sensor noise entirely to do everything it can to keep the interferometer operational. In practice it is not possible to implement Eq. (7.6) because D_i is unknown. Instead, the adaption implements Eq. (7.7), which employs the measured approximation $\hat{M}_{RMS,i}$ given by the output of the RMS filter.

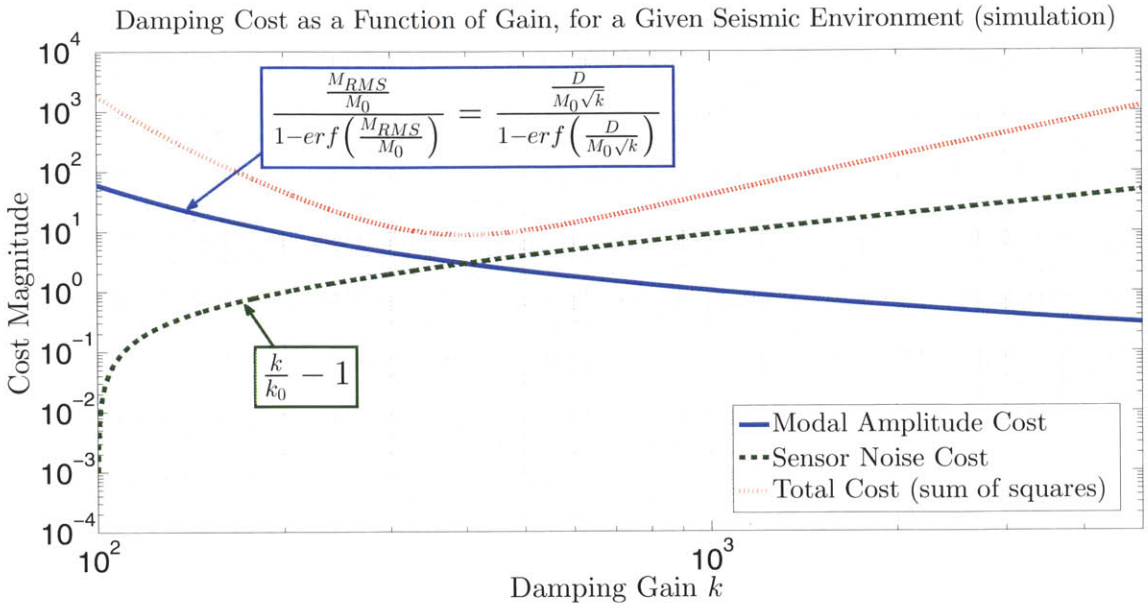

Figure 7-6: Example cost function employed by the adaptive modal damping method for each mode. In this example, $D = 94$, $M_0 = 6$, and $k_0 = 100$.

$$N_{c,i}(k_i) = \frac{k_i}{k_{0,i}} - 1 \tag{7.4}$$

$$M_{RMS,i}(k_i) = \frac{D_i}{\sqrt{k_i}} \tag{7.5}$$

$$M_{c,i}(M_{RMS,i}) = \frac{\frac{M_{RMS,i}}{M_{0,i}}}{1 - erf\left(\frac{M_{RMS,i}}{M_{0,i}}\right)} \tag{7.6}$$

$$M_{c,i}(\hat{M}_{RMS,i}) = \frac{\frac{\hat{M}_{RMS,i}}{M_{0,i}}}{1 - erf\left(\frac{\hat{M}_{RMS,i}}{M_{0,i}}\right)} \tag{7.7}$$

7.4 Adaptation Box

7.4.1 Gauss-Newton Algorithm

The adaptation box receives the two cost components for each mode i and calculates how to step the corresponding damping gain k_i. The Gauss-Newton algorithm in Eq. 7.8 performs this calculation.

$$k_{i,n+1} = k_{i,n} - \alpha_{i,n}\mathbf{J}_{i,n}^{\dagger}\mathbf{c}_{i,n} \tag{7.8}$$

$$\mathbf{c}_{i,n} = \begin{bmatrix} M_{c,i,n} \\ N_{c,i,n} \end{bmatrix} \tag{7.9}$$

$$\begin{aligned} \mathbf{J}_{i,n} &= \frac{d}{dk_i}\mathbf{c}_{i,n} \\ &= \begin{bmatrix} -\dfrac{M_{c,i,n}}{2k_{i,n}}\left(1 + \dfrac{2}{\sqrt{\pi}}M_{c,i,n}e^{-\left(\frac{\hat{M}_{RMS,i,n}}{M_{0,i}}\right)^2}\right) \\ \dfrac{1}{k_{0,i}} \end{bmatrix} \end{aligned} \tag{7.10}$$

The subscript n refers to the time of the current adaption step. The adaptation stepping rate is different from the system's sampling rate. This adaption rate is variable, but must always be slower than the sampling rate since a collection of samples are required to build an average at each step. The adaption rate should also be slower than the slowest closed loop pole to maintain modal damping stability. Further analysis, with Lyapunov functions for example, is required for faster adaptation rates. Section 7.4.2 details the switching of the adaption rate.

The cost function components are listed in the two element column vector $\mathbf{c}_{i,n}$. The 2×1 matrix $\mathbf{J}_{i,n}$ is a Jacobian matrix that lists the gradients of $\mathbf{c}_{i,n}$ with respect to $k_{i,n}$. This matrix is model based, depending on Eqs. (7.5) and (7.6).

The † symbol signifies the Moore-Penrose pseudoinverse operation. The scalar $\alpha_{i,n}$ is the Gauss-Newton step size. $\alpha_{i,n}$ is also variable and discussed in the next section.

7.4.2 Switching Step Rates and Sizes

As Figure 7-1 suggests, the adaptation must be able to handle seismic disturbances that vary at a large range of time scales. To better handle this range, the adaptation switches its step size and step rate so that it can respond quickly when necessary while still converging accurately to the optimal damping gain when it can.

Figure 7-7 illustrates how this switching works. It is driven by a boundary layer always referenced a factor of $\beta > 1$ above the sensor noise cost at step n. The sensor noise cost makes a good reference because its shape is always the same, where the modal amplitude cost changes according to unknown variable seismic disturbances. This moving boundary splits the adaption into two regimes. If the modal amplitude cost is below this boundary, the adaption remains in the regime of small α and slow step rates. In this regime, the adaption converges accurately to the most optimal damping gain at the expense of speed. If the modal amplitude cost jumps above this boundary at step $n+1$ due to a large sudden increase in the disturbance, then the adaption enters the regime of large α and short step rates. Such quick responses generally occur at the expense of poor convergence.

Since the boundary layer is always referenced to the current sensor noise cost, it will eventually move back above the modal amplitude cost if the disturbance remains at an increased amplitude. Then, the adaptation switches back into the slow, small step regime allowing good convergence to the optimal solution. In essence, the fast regime allows the adaption to quickly bring the modal damping into the 'neighborhood' of the optimal solution, while the slow regime brings it in the rest of the way.

The choice of β involves a design trade-off. Any values greater than 1 are valid. Values closer to 1 result in a more aggressive adaptation. Too close to 1 also risks the chance of triggering 'false alarm' crossing of the boundary. These false crossings occur because the modal amplitude cost is an estimate of the properties of a stochastic variable and is thus stochastic itself. Thus, β must be chosen large enough such that the chance of a false crossing is sufficiently improbable, but small enough that true

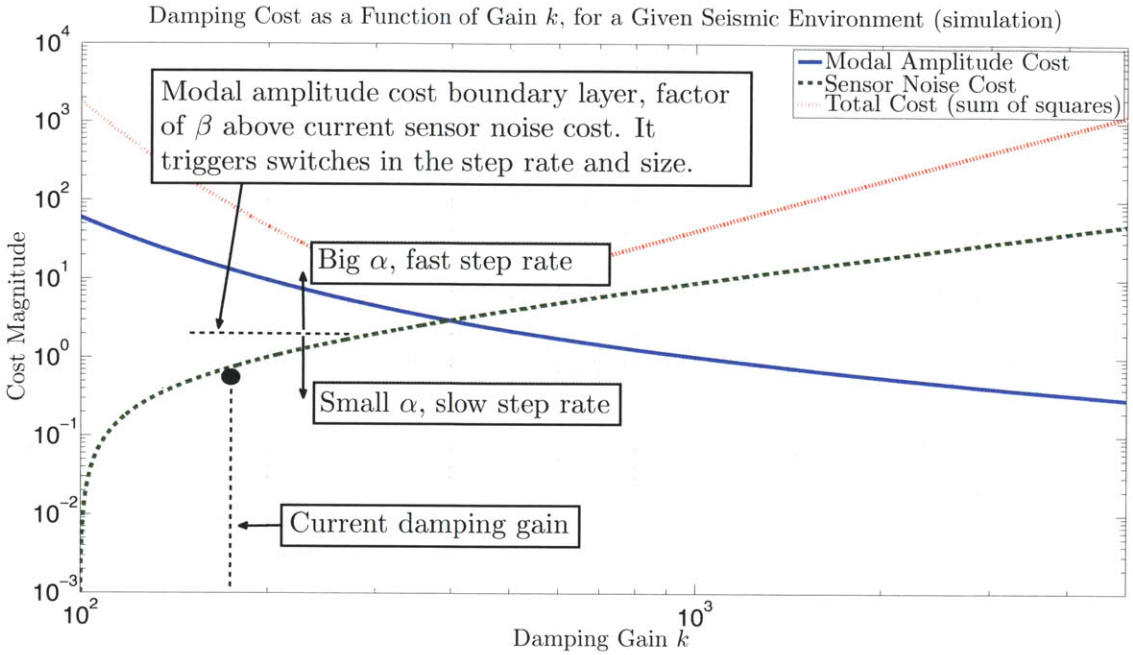

Figure 7-7: Boundary layer separating the cost function space into slow and fast adaptive regions.

disturbance changes are caught.

Similarly, the choices of the two step sizes and step rates involve a design choice. Though switching between two regimes allows more freedom, bigger steps and faster rates still yield more aggressive adaption. The rates and sizes should be big enough for the system to respond with the proper speed in each regime while still preserving the appropriate accuracy in the modal amplitude cost.

7.5 Experimental Setup

Figure 7-8 sketches the experimental setup built at the LIGO Advanced Systems Test Interferometer (LASTI) at MIT. This setup models on a smaller scale the 4 km Fabry-Perot cavities of the Advanced LIGO observatories. These cavities consist of two quadruple pendulums facing each other on opposite ends of the 4 km space. The MIT setup has tighter space constraints and so consists of a full scale quadruple pendulum prototype facing a smaller triple pendulum with 16 m between them. The triple pendulum has three 3 kg stages and is about 80 cm tall.

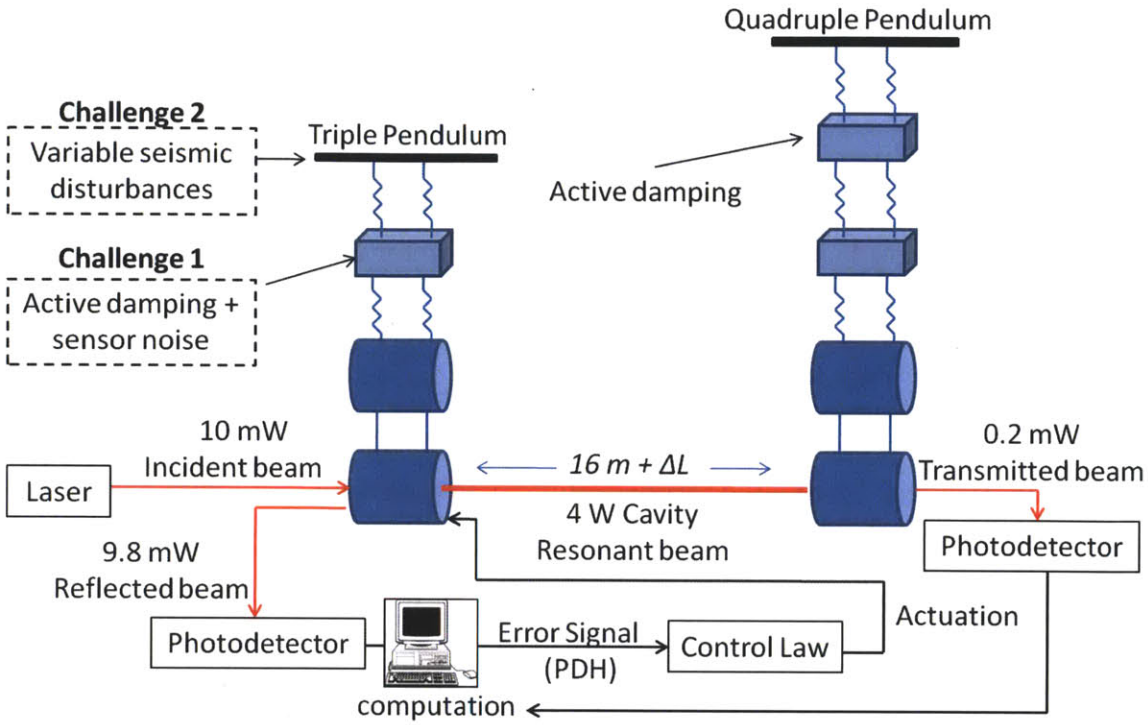

Figure 7-8: Experimental Fabry-Perot cavity setup at MIT used to test the adaptive modal damping technique. The input laser power is 10 mW, the power transmission coefficient of the triple pendulum test mass is 0.01 and the transmission coefficient of the quadruple pendulum test mass is 5×10^{-5}.

The cavity length measurement is generated from the laser light reflecting off the back of the triple pendulum test mass by using a standard Pound-Drever-Hall (PDH) technique [50]. The cavity length between the test masses is maintained by applying feedback forces to the bottom stage of the triple pendulum using the measured cavity length as the error signal. Since the cavity measurement only has sufficient signal when the cavity is near optical resonance, there is a 'lock acquisition' problem that must be solved. This is the problem that to get a good interferometer signal the cavity must be controlled near optical resonance, but to apply the control a good interferometer signal must exist. The cavity is said to be 'locked' when it reaches the state where it is controlled near resonance. The solution to this problem is to use the laser beam transmitted through the partially reflective quadruple pendulum test mass as a trigger because the power in this beam is proportional to the power in the cavity. When this transmitted power is measured beyond a certain threshold, the

cavity measurement is assumed good and the control is engaged. Typically with this triggering the cavity will swing through a few optical fringes, before finally acquiring lock. The control here has a bandwidth of 100 Hz, which suppresses the nominal cavity displacement under 1.5×10^{-12} m RMS for typical seismic disturbance magnitudes. This closed loop displacement is small enough that nonlinearities are not visible in the output.

This setup is significantly less sensitive than the optical cavities comprising a true GW observatory. Up to about 20 Hz the system is dominated by seismic disturbances coming through the relatively poorly isolated triple pendulum. At 10 Hz without the cavity control the test mass of the triple pendulum moves at $10^{-11}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$, about 8 orders of magnitude more than the true advanced LIGO test masses. Above 20 Hz the cavity is dominated by laser frequency noise. This noise constitutes a type of sensor noise at the level of $10^{-13}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$, over 6 orders of magnitude greater than the desired Advanced LIGO sensitivity. Suffice it to say, this cavity will not be measuring GWs. Nonetheless, it is a good setup to test the various challenges associated with the damping loops. The advantage of the MIT setup is that its cavity controller has a large dynamic range, so that these disturbances and noises can be suppressed to maintain a robust cavity lock. We will be interested in how the cavity signal responds to various levels of seismic disturbances and active damping. Further details of this cavity are found in reference [12].

The triple pendulum in this setup is used as the testbed for the adaptive modal damping. The reason for this choice is that the OSEM sensor noise (even artificially enhanced noise) is simply not visible in the relatively insensitive cavity measurement once filtered through the well isolating quadruple pendulum. To test the adaptive modal damping, artificial variable seismic disturbances are applied by driving a hydraulically actuated table underneath the triple pendulum. The quadruple pendulum is given basic, non-adaptive active velocity damping. This damping benefits the cavity control by reducing the motion of the quadruple pendulum test mass.

The triple pendulum has three modes that couple strongly to the cavity axis. However, the cavity control reduces this number to two. Because the cavity control's

bandwidth of 100 Hz is well above the highest frequency mode of the triple pendulum at 2.8 Hz, the response of the bottom stage along the cavity axis follows the quadruple pendulum much more than the rest of the triple pendulum. Consequently, the triple pendulum perceives its bottom stage as an effectively rigid constraint along this axis. For this reason, the modal damping in this experiment is applied to two modes, one at 1.31 Hz and the other at 2.80 Hz. Logistically, this change in the triple pendulum's response creates a problem because the modal damping must be turned on after the cavity is already locked. In practice a robust, non-modal damping loop is applied to the top stage to help obtain cavity lock. This loop is designed to be stable both before and after the cavity is locked and does not need to meet any noise amplification requirements. Then after the cavity is locked, the top stage damping switches to the modal configuration.

Even though modal damping is applied to the triple pendulum, true OSEM sensor noise is not visible in the cavity signal. Consequently, for the purposes of this experiment, the sensor noise above 10 Hz is increased to about $2 \times 10^{-5}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. The enhanced sensor noise is implemented by summing with the sensor signal digitally generated white noise filtered through a fourth order highpass elliptical filter with a 10 Hz corner frequency, 1 dB of passband ripple, and a 60 dB stopband attenuation. The corner frequency was chosen higher than the mode frequencies because this enhanced sensor noise is so large that the damping performance would be significantly impaired. The corner is also low enough that the enhanced noise will not be lost to the cavity signal by the mechanical filtering of the triple pendulum (as it would inevitably be with the quadruple pendulum).

The damping feedback filters are an earlier variation of the design presented in Section 6.2. Equations (7.11) and (7.12) list the filters for modes 1 and 2 respectively. The main difference in the design here is that they have more aggressive lowpass filtering for more sensor noise isolation. The design in Section 6.2 considers that the greater phase margin afforded by filters with less aggressive filtering can reach greater closed loop damping ratios. The loss in noise filtering is compensated for by the adaptation. As a result, the maximum damping gains in this experiment do not

reach the upper limits that provide maximum damping.

$$G_1 = \frac{(s + 80.54)(s + 0.8219)}{(s^2 + 8.136s + 66.2)(s^2 + 16.7s + 595.8)} \tag{7.11}$$

$$G_2 = \frac{(s + 173.7)(s + 1.772)}{(s^2 + 17.54s + 307.8)(s^2 + 36s + 2770)} \tag{7.12}$$

The analysis in Section 6.4 establishes the link between damping gains and damping ratios. The damping gains were given upper bounds set to $k_1 = 2000$ corresponding to an approximate damping ratio of 0.32, and $k_2 = 5000$ corresponding to an approximate damping ratio of 0.10. These upper bounds were chosen to be small enough to ensure stability, robustness, and sufficient decoupling between the modes. The analysis in Section 6.4, developed after this experiment, indicates that both upper bounds could be pushed to damping ratios of $\eta_{max} = 0.39$ if the feedback design in Section 6.2 was employed. Beyond this limit the modal damping feedback loops can no longer be considered independent because the underactuated nature of the damping control causes coupling between the modes.

The sensor noise costs were calculated with $k_{0,1} = 500$ and $k_{0,2} = 100$. These are the experimentally determined upper limits where sensor noise is not visible in the interferometer output. These values also establish the lower bounds in the range of possible adaptive damping gains. These lower bounds correspond to a damping ratio of 0.080 for the 1.31 Hz mode and a damping ratio of 0.002 for the 2.80 Hz mode. The lower frequency mode has a higher minimum damping ratio because it is farther from the frequencies where sensor noise first becomes visible.

The modal amplitude costs were calculated with scale factors $M_{0,1} = 15$ and $M_{0,2} = 4$. These values were chosen experimentally to provide an approximately equal trade-off between sensor noise amplification and upconversion for all possible damping gains.

The switching boundary layer was chosen to have a height of $\beta = 2$ above the sensor noise cost for both modes. This value was determined by experimentally observing the smallest β that would not result in false boundary crossings for at least an hour of constant disturbance. Also, both modes were given slow step rates of 60 s

and fast step rates of 30 s. In the slow step regime the adaptation employs RMS filters with poles at $\omega_c = 0.1\,\text{rad/s}$ in the modal amplitude cost path. In the fast step regime the RMS filters have poles at $\omega_c = 0.2\,\text{rad/s}$. Both slow and fast modal amplitude costs are calculated simultaneously so that when the step rate changes, the system merely chooses the cost value appropriate to the new rate. The small step sizes were chosen as $\alpha = 0.125$ for the first mode and $\alpha = 0.25$ for the second mode. The large step sizes were chosen to be $\alpha = 2$ for the first mode and $\alpha = 4$ for the second mode.

The switching from a slow step rate to a fast rate is set to trigger when either one of the slow modal amplitude costs jumps above its boundary layer. As a result, both modes always take simultaneous steps. Triggering on the slower cost measurement also reduces the likelihood of 'false alarms'. The switching of step sizes however occurs independently for each mode and is triggered by whichever modal cost measurement is active, slow or fast. Thus, it is possible for three cases to exist for each mode. A slow step rate with a small step size, a fast step rate with a large step size, or a fast step rate with a small step size.

7.6 Results

This section presents the results from two types of tests performed on the experimental setup at MIT described in the previous section. Both tests apply disturbances by driving the hydraulically actuated table mounting the triple pendulum. The velocity of the support table tracks the input disturbance signal. The first test emphasizes the adaptation's time response, the second its steady state frequency response.

The first test evaluates the system's time response to large amplitude, short time constant disturbances. Specifically, the disturbance induced by trains passing over the tracks near the LIGO Livingston Observatory was reproduced by driving the table underneath the triple pendulum with the same approximately 15 minute waveform observed by local seismometers on November 21, 2009. The results from this test are shown in Figure 7-9. The figure includes the results from three trials of the test. The first trial was run with non-adaptive modal damping by holding the damping gains

fixed at the values of k_0. The second and third trials were run with the adaption engaged, using the k_0 values as the initial conditions.

The lower plot of Figure 7-9a displays the disturbance waveform as measured by seismometers mounted to the triple pendulum's table. The top plot displays the measured cavity displacement between the two test masses. 7-9b displays the power of the transmitted laser beam exiting the back of the quadruple pendulum test mass. This value is proportional to the laser power resonating between the two test masses. The normalized optimal value of 1 indicates maximum power, where 0 indicates no power. 7-9c plots the evolution of the adaptive damping gains for both modes.

The first trial without adaption occurs before the vertical dotted line at 10 minutes on the time axis. This trial was cut short because the disturbance was strong enough to break the cavity lock, meaning a failure of the control system. The lock loss is indicated in 7-9b by the drop in transmitted power to zero. The subsequent two trials with adaptation occur from 10 to 25 minutes and 48 to 63 minutes. Though a visible reduction in the transmitted power is visible, the adaptation during these trials was successful at maintaining cavity lock. This success is a result of the rapid response of the adaption shown in 7-9c to the initiation of the disturbances. Each trial reached the maximum allowed damping gain for mode 1 within about 3 steps, or about 1.5 minutes. The second mode's damping gain followed a few minutes later. The apparent delay in the second mode is a result of the character of the disturbance's frequency content. The disturbance contains lower frequencies when the train is relatively distant, but increases in frequency as the train gets nearer. The adaption responds accordingly by only increasing the damping gains required at the appropriate time, minimizing the total sensor noise amplification. First it increases the gain on the lower frequency mode and later increases the gain on the higher frequency mode as the frequency content evolves. The gains drop in magnitude relatively slowly after the disturbance has passed because the modal amplitude cost falls below the slow/fast boundary layer. The disturbance from 65 to 75 minutes is not part of the test and represents a true disturbance that just happened to occur near the test's conclusion. The adaption responds to this disturbance as well.

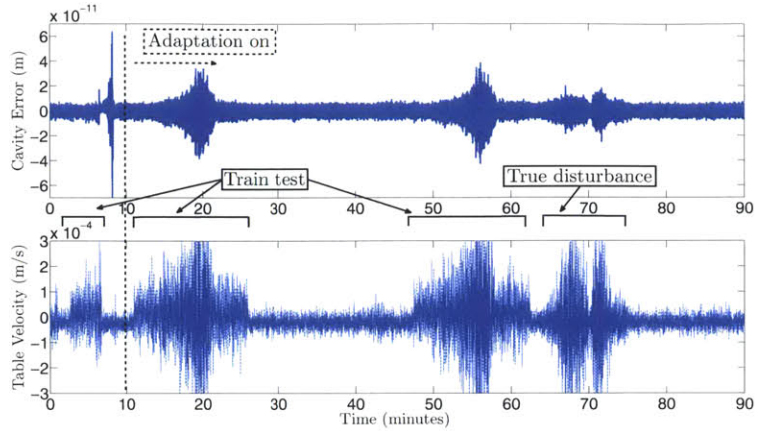

(a) Cavity Error Response to Train Test

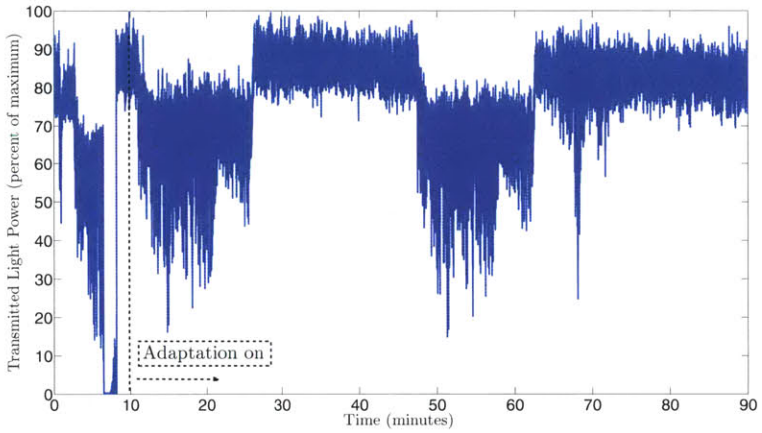

(b) Transmitted Light Response to Train Test

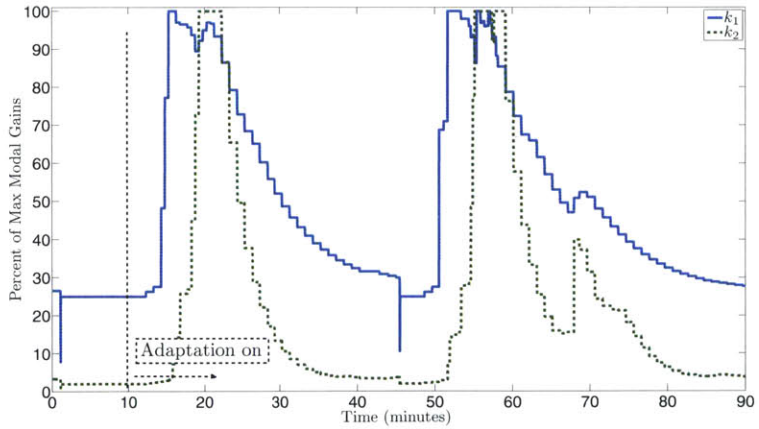

(c) Damping Gain Response to Train Test

Figure 7-9: This figure shows the response of the MIT setup to disturbances applied by shaking the table underneath the triple pendulum. The top half of plot (a) shows the measured change in displacement between the two test masses due to the velocity of these disturbances shown in the bottom plot of (a). Plot (b) shows the measured power of the laser transmitted through the quadruple pendulum, which is proportional to the power resonant in the cavity between the test masses. Plot (c) shows the response of the adaptive damping gains.

The second test evaluates the interferometer's frequency response with adaptive modal damping to varying amplitudes of a seismic disturbance. This disturbance is applied to the triple pendulum in a frequency band between 1 Hz and 3 Hz. This band was chosen because it overlaps both with common anthropogenic disturbance frequencies and the pendulum mode frequencies. The signal is generated by filtering white noise through a sixth order bandpass elliptical filter with a 60 dB attenuated stopband and 1 dB ripple passband.

Figure 7-10 illustrates these results. The black solid line serves as the reference when there is no enhanced disturbance and the adaption is at the initial conditions of $k_1 = k_{0,1}$ and $k_2 = k_{0,2}$. The disturbance amplitude is subsequently increased in 4 steps. Each step adds a factor of 1.7 in amplitude from the background reference level. Between each step the adaption is allowed to settle to its steady state optimal damping gains. The amplification of sensor noise is visible in the band from 6 Hz to 15 Hz. The upconversion of the seismic disturbance is visible between 3 Hz and 6 Hz. The cost function was tuned by adjusting the $M_{0,1}$ and $M_{0,2}$ values so that the upconversion increases at approximately the same rate as the sensor noise amplification with this type of disturbance.

7.7 Conclusion

The adaptive modal damping algorithm presented in this chapter offers a real-time optimal solution to the damping of suspension systems in gravitational wave interferometers. This algorithm finds a set of optimal gains on modal damping feedback filters, where the optimal point evolves in time in response to the changing environmental disturbances. The adaptation converges to the optimal point at different rates depending on the rate of change of these disturbances. It responds with fast adaption to sudden increases in disturbance amplitudes to quickly reach the neighborhood of the optimal solution. This fast response also allows the interferometer to remain operational where it otherwise would not. Slow adaptation then allows the system to more precisely close in on the optimal set of gains.

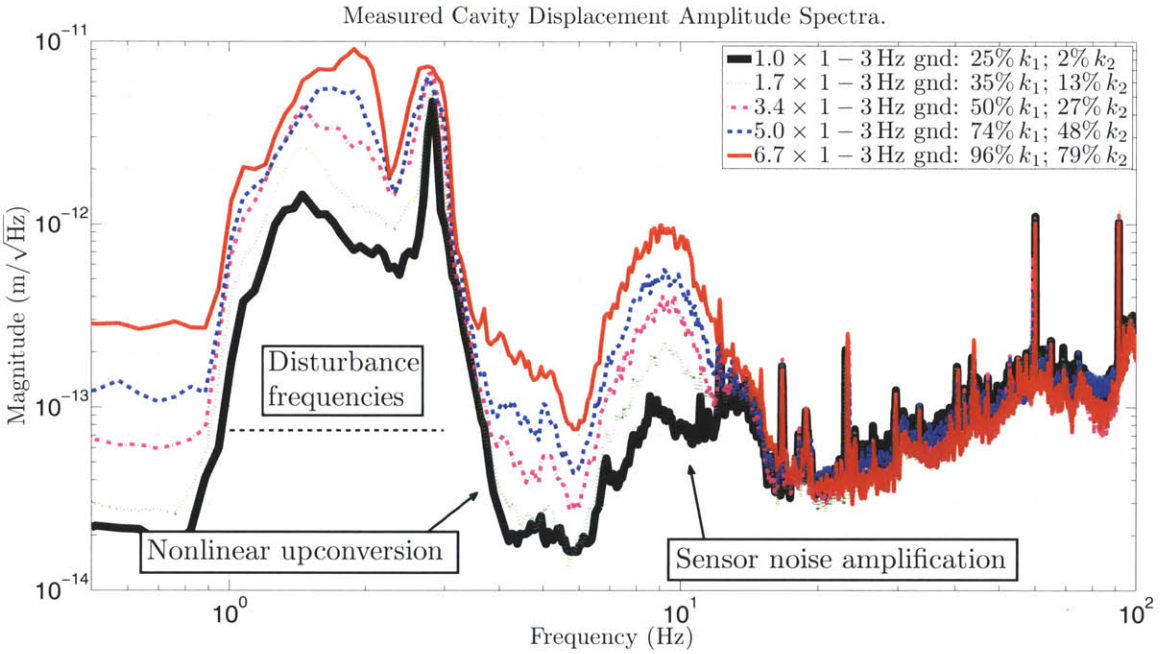

Figure 7-10: The steady state measured amplitude spectra of the output of the MIT interferometer with varying amplitudes of seismic disturbance. These disturbances are applied by shaking the table supporting the triple pendulum and are designed to be flat in velocity between 1 Hz and 3 Hz and close to zero elsewhere. The disturbance signal comes from filtering white noise through a bandpass elliptical filter. The modal adaptive damping is allowed to respond and settle to each amplitude of seismic disturbance.

The adaptive method described here is versatile in that it is easily modified to other classes of interferometer control. Examples of other control loops that could benefit from similar adaption are the test mass steering control, the control of auxiliary cavities, and the control of the vibration isolation tables supporting the pendulums. The adaption can also be modified to optimize the interferometer sensitivity to specific classes of astronomical events, as shown in Chapter 8.

There are a few significant limitations to this adaptive method. First, stability is only guaranteed if the adaption rates are well below the slowest eigenvalues of the closed loop system. From the point of view of the control system, the changing parameters can be thought of as quasi-static. If the parameters change more quickly, additional analysis is required, for example with a Lyapunov function. Second, faster adaption rates lead to poorer convergence due to the necessity of estimating the variance of a stochastic variable. Thus, there is a fundamental limit at how quickly any adaption can respond. This estimate requires the measurement of a certain number oscillations. As a result, lower frequency disturbances will inevitably result in slower adaptation. The final significant limitation is obtaining the Jacobian matrix of the Gauss-Newton algorithm. The adaption here relies on a model of the gradient of the cost functions with respect to the damping gains. Calculating this model was possible because the cost functions have a simple relation to the damping gains. More complicated relations might make a model difficult to obtain or require the Jacobian matrix to be estimated either off-line or on-line. On-line estimation may greatly degrade the performance and stability of the adaptation.

Future work on this adaptation technique should pursue multiple directions. One direction is an optimal way to transition the modal damping feedback gains. The adaption in the experiment here updates the gains with abrupt steps in the values. Such abrupt transitions can cause momentary noise and disturbance upconversion, exactly what we are trying to avoid. Work needs to be done to determine how much this transition needs to be smoothed out since it will impact the adaption rate.

Chapter 8

Adaptive Modal Damping for Targeted Searches

Each class of the GW sources produces waves with a certain signature in the frequency domain. Consequently, adaptive modal damping does not have the same impact on all sources. This chapter employs a simulation of an Advanced LIGO 4 km Fabry-Perot cavity to establish adaptive modal damping cost functions that are optimal to the detection of waves from these sources. Section 8.1 summarizes the details of this simulation. Section 8.2 develops the metrics used to quantify these sensitivities. Sections 8.3 and 8.4 use these metrics to present the simulation results for the source classes.

8.1 Simulation Details

The adaptive modal damping performance is explored through MATLAB® and Simulink® simulations of an Advanced LIGO 4 km arm cavity. This simulation, illustrated by the diagram in Figure 8-1 follows a similar structure to Figure 7-2 where one of the pendulums in the cavity receives modal damping and variable seismic disturbances along the x axis. However, the cavity control forces that maintain the fixed 4 km between the test masses are applied to the pendulum at the other end of the cavity using the feedback filters C_1 to C_4. Separating the damping and cavity control

allows us to study modal damping on a greater number of modes, in this case four. The state space models of these identical pendulums and the cavity control feedback filters employed by this simulation are listed in Appendices A and C respectively. The frequencies of pendulum model's modes are 0.443 Hz, 0.996 Hz, 2.01 Hz, and 3.42 Hz.

The damping filter design of Section 6.2 is employed. The feedback gains k are replaced with the ζ_r values from Section 6.4. The relationship between the two is given by Eq. (8.1), which combines Eqs. (6.20) and (6.21). This transformation provides a more intuitive understanding of the feedback gains since $\zeta_{r,i}$ approximates the closed loop damping ratio of mode i.

$$\zeta_{r,i} = \frac{\Phi_{1i}(\Phi^{-1})_{i1}k_i}{2\omega_i(M_m)_{ii}} \tag{8.1}$$

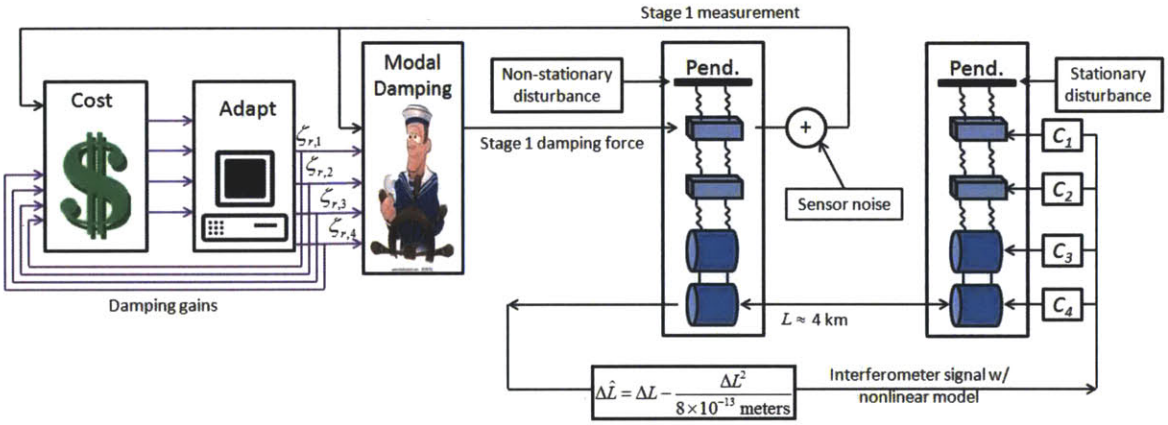

Figure 8-1: Top level block diagram of the simulation used to investigate the influence of adaptive modal damping on the sensitivity of Advanced LIGO to particular GW sources.

To simulate the effect of upconversion, a nonlinearity must be introduced into the simulated interferometer. It is known from experience with the Initial LIGO interferometers that there are many possible sources of nonlinearity, many of which are poorly understood or poorly measured as described in Section 7.1. Further, the Advanced LIGO interferometers are not yet operational (under construction at time of writing), so their behavior has yet to be characterized. Consequently, a nonlinearity must be assumed for the purposes of this work. A simple nonlinear function (8.2) is

chosen for the interferometer output in terms of the interferometer arm length.

$$\hat{\Delta L} = \Delta L - \frac{\Delta L^2}{8 \times 10^{-13}\,\text{meters}} \tag{8.2}$$

ΔL is the true deviation in the interferometer arm length and $\hat{\Delta L}$ is the interferometrically measured deviation. For small ΔL this equation is linear. For larger ΔL, the quadratic term becomes significant. This form was chosen because it is straight forward to simulate and understand, and it is known that many nonlinearities have approximately quadratic behavior. The quadratic term in this equation is scaled so that seismic upconversion is not visible beyond the background seismic level until the interferometer displacement goes beyond about $10^{-14}\,\text{m}$, or about an order of magnitude greater than the control requirements. That requirement was chosen to minimize known nonlinear terms in the interferometer output [26].

This upconversion behavior is investigated for varying amplitudes of a seismic disturbance. The constant nominal background seismic disturbance level is modeled by applying filtered white noise to the ground input of both pendulums. The white noise has an amplitude of $1\,\text{m}/\sqrt{\text{Hz}}$ and is filtered by the transfer function in Eq. (8.3), which models the Advanced LIGO pre-isolation seismic goal [7].

$$\frac{5.0532 \times 10^{-10}(s+6.283)^4(s+1.257)}{(s+1.257)^6(s+0.06283)} \tag{8.3}$$

The variable seismic disturbance is applied only to the pendulum with modal damping and is modeled by filtering similar white noise through a lowpass filter with a cutoff at 4 Hz. The lowpass filter consists of two parts. The first is a sixth order elliptical lowpass with 1 dB of passband ripple and 60 dB of stopband attenuation. It is implemented with the MATLAB® function call $ellip(6, 1, 60, 2 * pi * 4, 'low', 's')$. The second part is a simple lowpass filter with 4 complex poles at 4 Hz and damping ratios of $\frac{1}{\sqrt{2}}$. The maximum 0 Hz gain is set to 5×10^{-10}.

Sensor noise is assumed to have a linear influence on the interferometer's performance and is modeled in MATLAB® with an amplitude of $10^{-10}\,\text{m}/\sqrt{\text{Hz}}$ for each

OSEM sensor. Its amplification is bounded by the minimum and maximum levels of modal damping. Here the minimum level is chosen to achieve a damping ratio of 5×10^{-4}, or quality factor of 1000, which is about the level of natural damping and small enough that sensor noise is negligible. The maximum level is chosen as the theoretical maximum damping ratio of 0.203 determined by Section 6.4, or an approximate quality factor of 2.5.

Figure 8-2 plots the seismic disturbance and sensor noise inputs to the pendulums. The solid line is the nominal low amplitude seismic disturbance, the dashed line is the maximum additional disturbance summed with the nominal level in the simulation. This increase represents a maximum increase of two orders of magnitude between 1 Hz and 4 Hz, which are typical anthropogenic disturbance frequencies overlapping with the pendulum resonance frequencies. The dotted line is a model of the measured stationary noise of the stage 1 sensor. Reference [51] presents an investigation into the variation of seismic amplitudes at the Livingston Observatory.

Figure 8-3 illustrates the bounds on the contributions to the simulated interferometer output due to the disturbance and noise shown in Figure 8-2. The blue solid line represents the contribution from the nominal background level of the seismic disturbance, where upconversion is not visible. The dotted blue line represents the disturbance contribution when the maximum variable disturbance is applied with minimum damping. Large amounts of upconversion are visible here. A linear response would reconstruct the previous line at these frequencies. The red solid line represents the sensor noise amplification due to the minimum level of damping. The red dotted line represents the sensor noise amplification with the highest possible level of the damping. The thick solid black line is the design sensitivity for the Advanced LIGO interferometers operating in mode 1a as described by [52]. This mode of operation is chosen since it represents the best low frequency performance. It is calculated using LIGO's Gravitational Wave Interferometer Noise Calculator (GWINC) [53]. Note that the amplified seismic and sensor noises go beyond this design curve at the lowest frequencies. The result of this point has important consequences for which GW sources are influenced by adaptive modal damping.

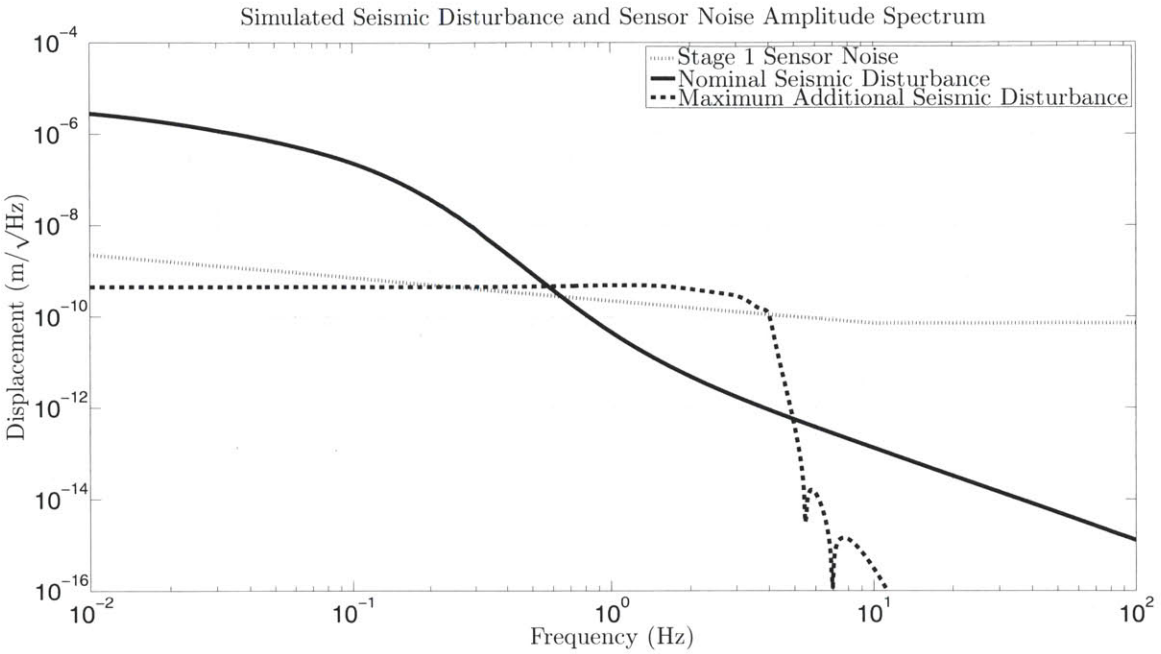

Figure 8-2: Seismic disturbance and damping sensor noise included in the modal damping simulation.

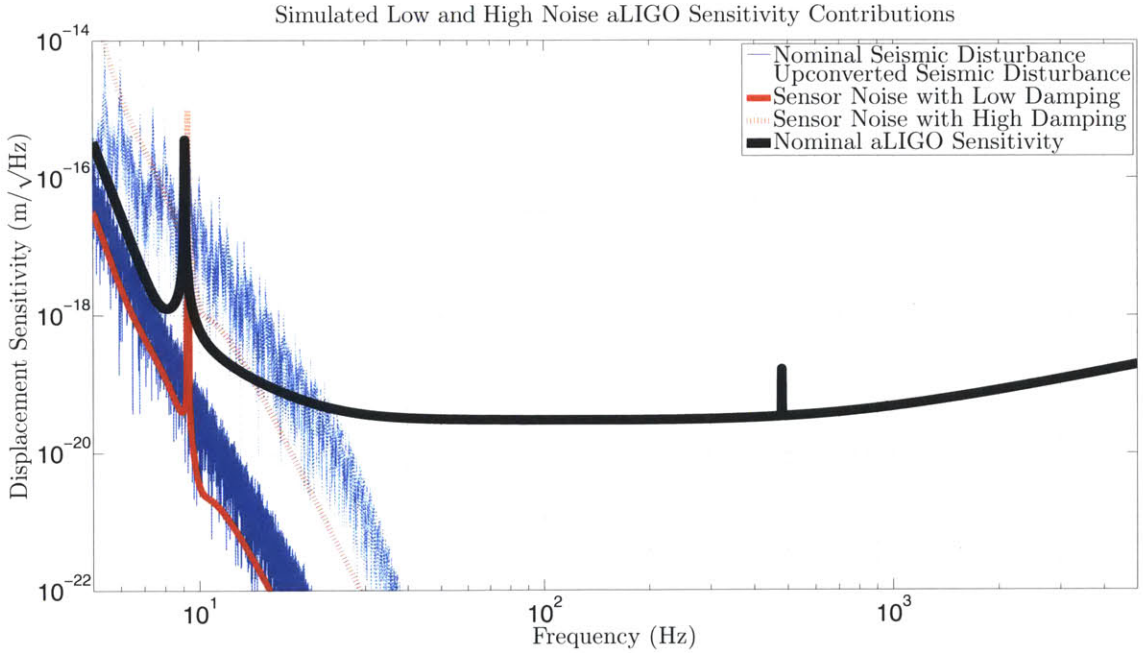

Figure 8-3: Simulated example of relevant low and high noise terms relative to Advanced LIGO's nominal design sensitivity. The high upconverted seismic disturbance line is generated from large disturbances between 0 Hz and 4 Hz when damping is small. The upconversion results from a simulated quadratic nonlinearity term in the behavior of the interferometer with large displacements.

8.2 Advanced LIGO Sensitivity to GW Sources

This section presents the basic theory on how LIGO measures its sensitivity to the GW source classes of binary inspirals and the stochastic background. The sensitivities to these two classes are most likely to be influenced by adaptive modal damping. The metrics presented here will be used to quantify the simulation results in Sections 8.3 and 8.4. The sensitivities to the source classes of pulsars and bursts are unlikely to be influenced because they dominate at frequencies well beyond the reach of adaptive modal damping. Consequently, they are not considered here.

8.2.1 Binary Inspiral

The sensitivity of the Advanced LIGO interferometers to binary inspirals is typically reported as a distance called the 'inspiral range'. This is the distance from the Earth to which the interferometer can detect an inspiral event with a certainty of approximately five standard deviations[1]. The inspiral range, r, can be calculated with the following equations [54].

$$r^2 = \frac{5c^{1/3}\mathcal{M}^{5/3}\theta^2}{96\pi^{4/3}\rho^2}\int_0^{f_{ISCO}} \frac{df}{f^{7/3}h^2} \tag{8.4}$$

$$\mathcal{M} = \frac{G}{c^2}\frac{(M_1 M_2)^{3/5}}{(M_1+M_2)^{1/5}} \tag{8.5}$$

$$f_{ISCO} = \frac{c^3}{6^{1.5}\pi G(M_1+M_2)} \tag{8.6}$$

M_1 and M_2 are the masses of the two inspiraling objects, $\mathcal{M}$ is a quantity known as the 'chirp mass' (in units of meters), c is the speed of light, G the gravitational constant, and ρ is the minimum desired signal to noise ratio. Typically, $\rho = 8$ for the certainty of approximately five standard deviations.

The GW frequency, f_{ISCO}, occurs when the orbital radius reaches the minimum value known as the innermost stable circular orbit (ISCO). At this radius the inspiral

[1]The exact certainty for a given inspiral range depends on a variety of factors, however five standard deviations is approximately true assuming Gaussian noise.

transitions to a merger phase where the objects begin their final plunge toward each other. Consequently, the inspiral phase has a maximum GW frequency at this radius defined by (8.6) [55].

Note that Eqs. (8.4) to (8.6) do not take into account the redshift. The most distant inspirals visible to Advanced LIGO ($\approx 3\,\mathrm{Gpc}$) do experience redshifted GWs by the time they reach Earth. However, the redshift is small enough that the error introduced to the inspiral range is likely to be less than the error introduced from the Advanced LIGO signal calibration. Consequently, the redshift is ignored.

It is expected that Advanced LIGO will be sensitive to inspiraling objects in the mass range from approximately 1 to 300 solar masses. LIGO will only be able to detect these objects in the final moments of their merger. The low end of this mass range will be visible on the order of 10 minutes and the high end on the order of 1 second.

Though the metric of inspiral range is useful to study the detectability of waves from inspiral sources, there is another complementary metric to consider. This metric is the number of cycles to merger. The observations of these cycles, or orbits, of the binary system permit the estimation of its properties. The more cycles observed, the better the deduction of the system properties. The lowest frequencies of the inspiral signal visible in the detector contain the vast majority of these cycles.

The expression determining the number of cycles to merger, N_{cyc}, at time t is given by Eq. (8.7). $f(t)$ represents the frequency of the wave at time t.

$$N_{cyc}(t) = \frac{\pi^{-8/3}}{32}\left(\frac{\mathcal{M}}{c}\right)^{-5/3}\left[f^{-5/3}(t) - f_{ISCO}^{-5/3}\right] \tag{8.7}$$

8.2.2 Stochastic Background

The sensitivity of the Advanced LIGO observatories to the stochastic background is typically reported as the minimum visible energy density of gravitational radiation in the universe. This value is reported as the dimensionless quantity Ω_{gw} and is given by Eq. (8.8) for the network of Advanced LIGO observatories. Ω_{gw} represents the

90% confidence level of detection.

$$\Omega_{gw} = \frac{1.65}{\rho_{critical}} \left[\frac{32G^2}{25c^4\pi^2} T \int_0^\infty S^2(f) df \right]^{-0.5} \tag{8.8}$$

$$\rho_{critical} = \frac{3c^2 H_0^2}{8\pi G} \tag{8.9}$$

$$S(f) = \frac{\gamma(f)}{f^3 P(f)} \tag{8.10}$$

$$\gamma(f) = -0.124842 j_0(\alpha) - 2.90014 * j1(\alpha) + 3.00837 * j2(\alpha) \tag{8.11}$$

$$\alpha = \frac{2\pi d}{c} f, \quad d = 3 \times 10^6 \text{ meters} \tag{8.12}$$

$\rho_{critical}$ represents the energy required to close the universe and is used here to normalize Ω_{gw} to a dimensionless quantity. T represents the length of observation time. Longer observations result in smaller Ω_{gw}, or better sensitivity. $S(f)$ is the sensitivity spectral density of the observatories to the stochastic background. $P(f)$ is the noise power spectral density of the observatories, assumed to be the same for all of them. $\gamma(f)$ is known as the overlap reduction function, which mathematically merges the network of observatories into a single equivalent detector. The function is purely dependent on the separation and relative orientation of the observatories. Eqs. (8.11) and (8.12) represent the overlap reduction function between the Hanford, WA and Livingston, LA observatories. For this case, Eq. (8.11) is a sum of three Bessel functions. The distance between the observatories is given by the value d. More details on this theory are given in reference [20].

8.3 Binary Inspiral Results

To get a preliminary idea of how adaptive modal damping might influence the detection of inspirals, the effects of upconversion and sensor noise are first tested independently. First, with damping turned down to the minimum level, the variable seismic disturbance is tested at a set of amplitudes from 0% to 100% of the maximum (Figure 8-4). Then, the effect of sensor noise is tested for a set of damping levels between the

minimum and maximum with no variable seismic disturbance (Figure 8-5).

Figure 8-4a plots the estimated inspiral range at each seismic amplitude as a function of binary mass. The curves have positive slope at low masses because the amplitudes of the waves increase with the mass. As the mass increases further, the termination frequency, f_{ISCO}, decreases below the sensitive region of the interferometer. This drop in frequency results in a negative slope despite the increase in wave amplitude. Overall, the inspiral range degrades as upconversion begins to dominate the low frequencies of the sensitivity curve. The loss in range is most pronounced for high mass inspirals because they exist primarily at these low frequencies. Low mass inspirals achieve most of their range between 30 Hz and 1000 Hz, beyond the influence of upconversion.

Figure 8-4b plots the corresponding reduction in the inspiral volume (r^3) at each mass value. The inspiral volume is more directly meaningful than the range because the number of potential inspirals scales linearly with the volume, which is proportional to the cube of the range. These curves are generated by cubing the ratio of the curves in Figure 8-4a with the maximum at each mass value. Because of the cubic relation a modest decrease in inspiral range can lead to a significant decrease in volume.

Figure 8-5 is structured in the same way as Figure 8-4, but depends on varying damping sensor noise rather than seismic disturbances. The sensor noise amplification scales linearly with the damping level, causing results that look very similar to those from increasing seismic amplitudes. The loss in inspiral range is not as dramatic since the sensor noise is never amplified as much as the maximum upconverted disturbance in this case, but the loss in visible volume is still significant.

The goal for the adaptive modal damping design is to create cost functions that will balance the effects observed in Figures 8-4 and 8-5. The desirable result will be enough damping to reduce the seismic upconversion while limiting the amount of sensor noise amplification. To explore this trade-off, Figure 8-6 shows the results of a test where the inspiral volume for two 150 solar mass inspiraling black holes was plotted for varying amounts of damping and seismic disturbances. For simplicity, equal damping is applied to the four x modes here. The optimal results where the

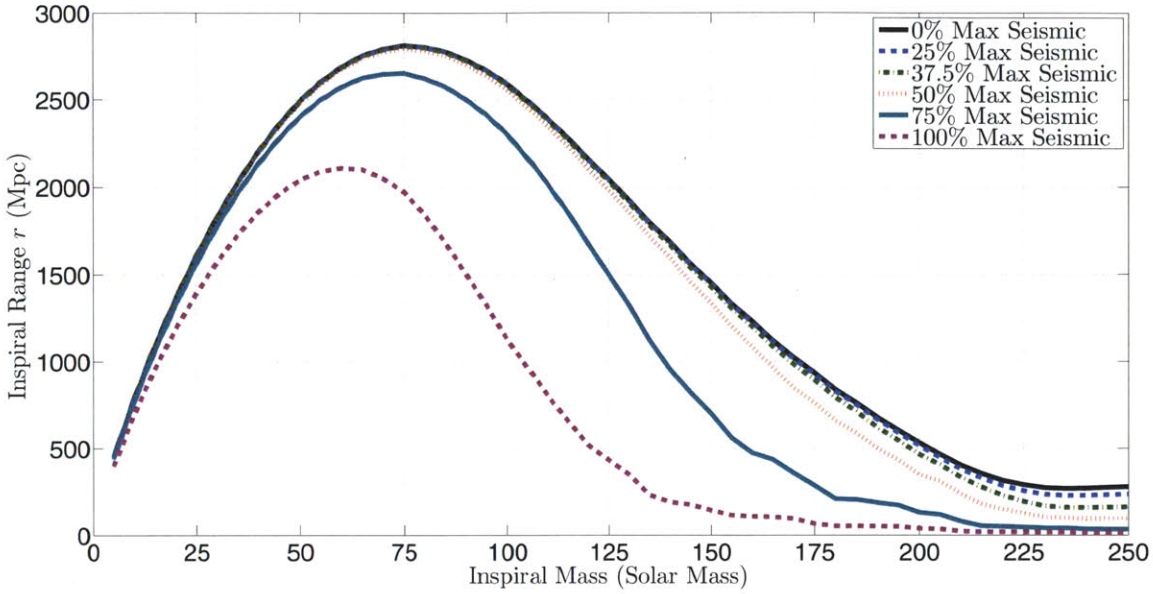

(a) Inspiral range vs. seismic disturbance.

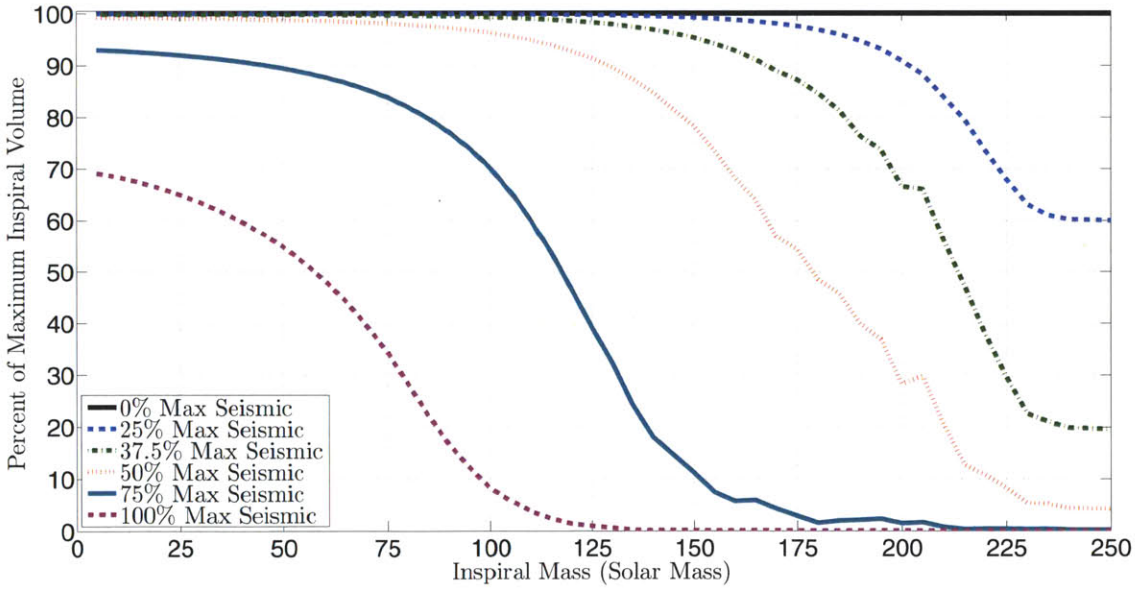

(b) Inspiral volume vs. seismic disturbance.

Figure 8-4: Effect of seismic disturbance on inspiral sensitivity as a function of inspiral mass. The horizontal axis lists the mass of each object in a symmetric binary system.

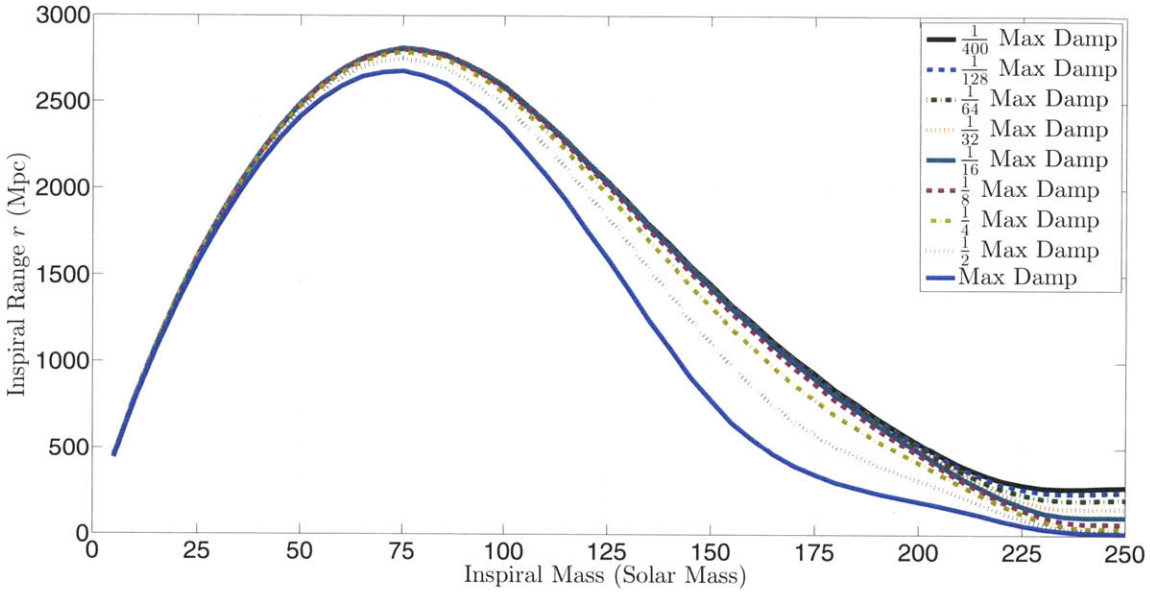

(a) Inspiral range vs. damping senor noise.

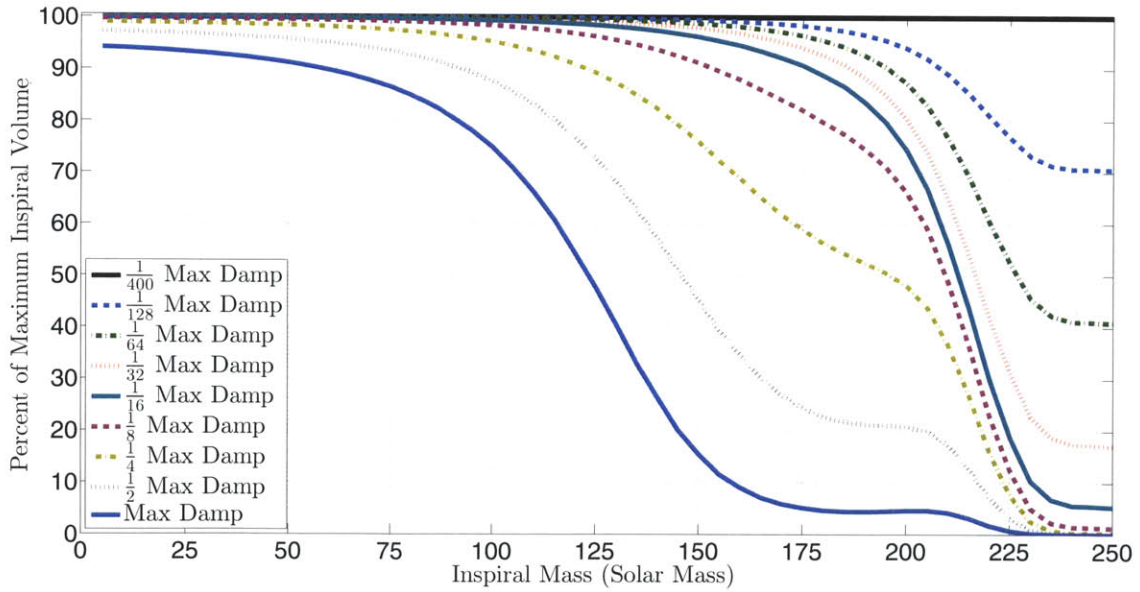

(b) Inspiral volume vs. damping sensor noise.

Figure 8-5: Effect of damping noise on inspiral sensitivity as a function of inspiral mass. The horizontal axis lists the mass of each object in a symmetric binary system.

curves are maximized occur when the sensor noise amplification in the sensitivity curve (e.g. Figure 8-3) equals the magnitude of seismic upconversion. The frequency band where this matters, approximately 5 Hz to 20 Hz, is narrow enough that a good overlap is possible even if the slopes are not exactly the same. Since optimal damping depends on this overlap, it is independent of binary mass. In fact, optimal damping for any source is given by this overlap.

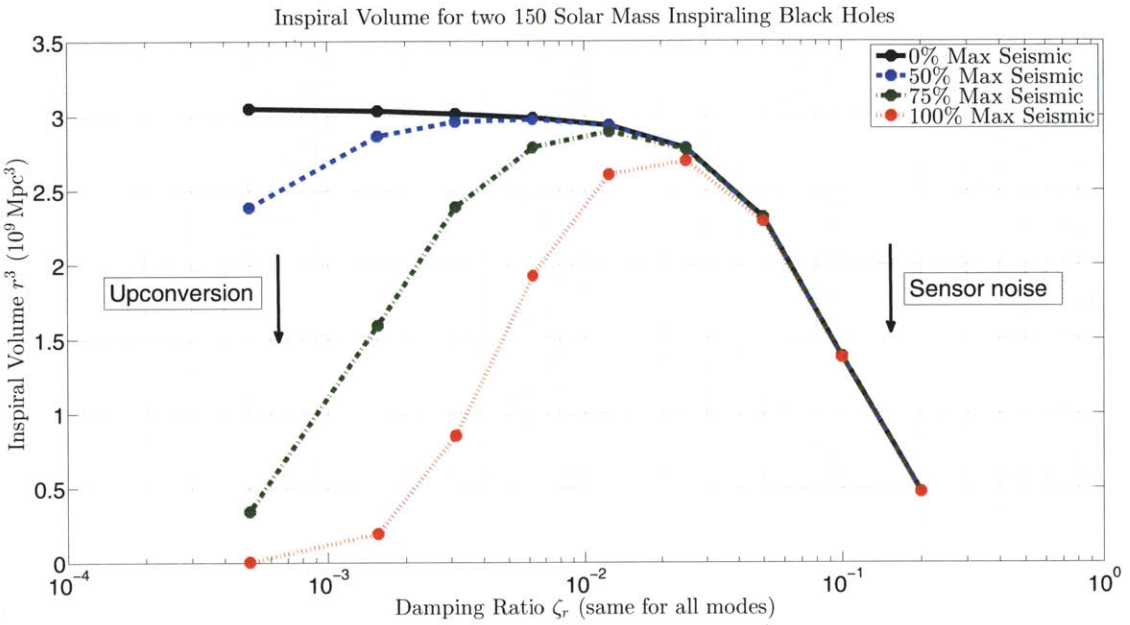

Figure 8-6: The inspiral volume for two 150 solar mass inspiraling black holes for varying seismic amplitudes and damping levels.

In Figure 8-6 the damping applied to all the modes is the same. However, improvements are possible by invoking adaptive modal damping cost functions that incorporate relevant information. These cost functions are realized in the form of Eqs. (7.4) and (7.7) by replacing $k_{0,i}$ with $\zeta_{0,i}$ and k_i with $\zeta_{r,i}$. The ζ_0 and $M_{0,i}$ values are scaled here for each of the four x modes i. The values chosen are selected to improve upon the results in Figure 8-6 and are listed in Table 8.1.

Table 8.1: Adaptive modal damping cost function parameters.

Parameter	**Mode 1**	**Mode 2**	**Mode 3**	**Mode 4**
ζ_0	5×10^{-4}	5×10^{-4}	5×10^{-4}	5×10^{-4}
M_0	2×10^{-8}	1×10^{-9}	8×10^{-10}	5×10^{-10}

Various information was used to select these cost function parameters. The higher frequency modes contribute more sensor noise amplification at the relevant frequencies. Additionally, mode 3 contributes the most to the displacement between the test masses. As a result, the M_0 values are chosen so that the mode 3 damping value is always higher than the optimal of Figure 8-6, and mode 4 is always slightly less. Mode 2 is chosen to have approximately the same damping as mode 4. It contributes less sensor noise, but similarly less displacement as well. The amplitude of mode 1 hardly increases at all because it is below the frequencies of this particular disturbance. As a result, its M_0 is made large enough so that its damping is insensitive to this disturbance, but small enough that it would be sensitive to other disturbances at the mode 1 frequency. All the ζ_0 damping values are chosen as the simulation's minimum damping ratio of 5×10^{-4} in order to ensure sensor noise is negligible for the un-amplified seismic state.

The optimal damping gains determined by these cost functions are listed for a selected set of seismic amplitudes in Table 8.2. The listed 150 solar mass inspiral volume is always greater than or equal to the maximum values in Figure 8-6, with the difference increasing with the seismic amplitude. Note that with this optimal damping the inspiral volume decreases very slowly with increasing seismic amplitude. At 100% of the maximum disturbance, only 8.4% of the volume is lost. As Table 8.3 shows, choosing the maximum damping level instead would lose 84.5%. Maintaining the minimum damping level loses 99.9%. Figure 8-7 illustrates the last column of these tables in a bar graph form.

Figure 8-8 plots the sensitivity spectrum given by the case of optimal damping

Table 8.2: Optimal damping values determined by the selected cost functions. The last column lists the inspiral volume for two 150 solar mass black holes. The percentage of the inspiral volume to the best case volume is given in parenthesis.

Seismic	$\zeta_{r,1}$	$\zeta_{r,2}$	$\zeta_{r,3}$	$\zeta_{r,4}$	r^3 ($\mathrm{Mpc}^3 \times 10^6$)
0% max	0.00054	0.00055	0.00051	0.00050	3051 (100%)
50% max	0.00059	0.0051	0.016	0.0042	2982 (97.7%)
75% max	0.00073	0.0089	0.029	0.0090	2915 (95.5%)
100% max	0.00091	0.014	0.046	0.016	2796 (91.6%)

Table 8.3: Inspiral results for the maximum disturbance with the maximum and minimum damping levels. The last column lists the inspiral volume for two 150 solar mass black holes. The percentage of the inspiral volume to the best case volume is given in parenthesis.

Seismic	$\zeta_{r,1}$	$\zeta_{r,2}$	$\zeta_{r,3}$	$\zeta_{r,4}$	r^3 ($\text{Mpc}^3 \times 10^6$)
100% max	0.2030	0.2030	0.2030	0.2030	473 (15.6%)
100% max	0.0005	0.00050	0.00050	0.00050	3.05 (0.1%)

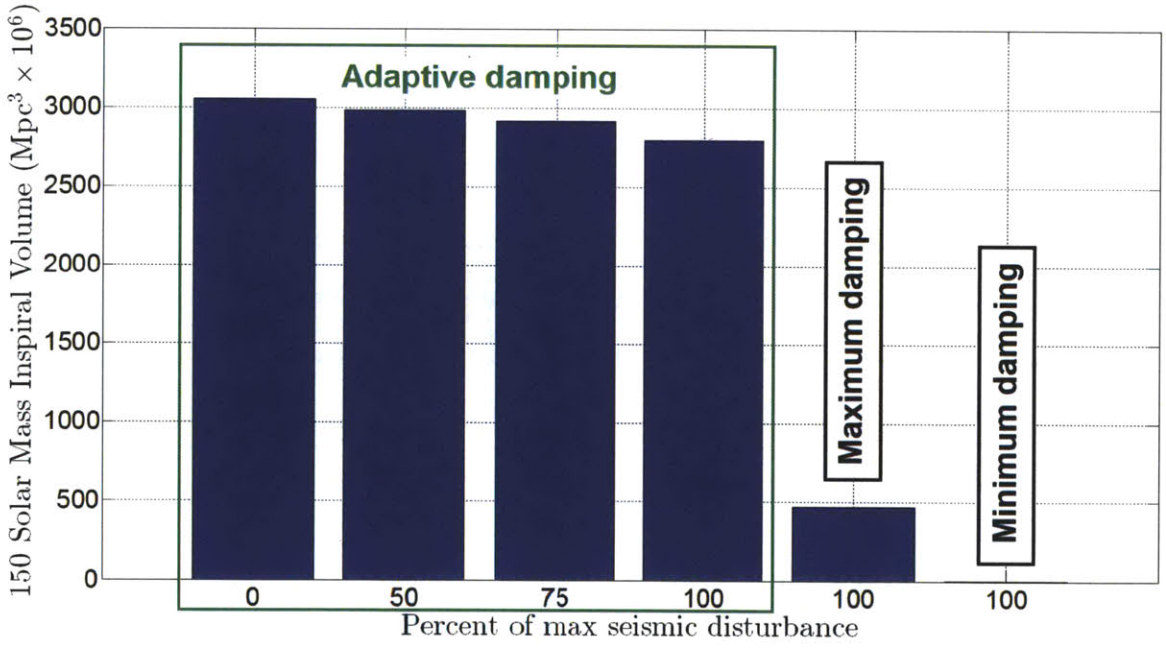

Figure 8-7: Summary of inspiral volume results with and without adaptive modal damping. This data is listed in the last column of Tables 8.2 and 8.3.

with the maximum seismic amplitude (last row of Table 8.2). Note, that the amount of upconversion is significantly less than the maximum level shown in Figure 8-3 despite equal seismic disturbances. Note, also that this is achieved with less than the maximum sensor noise amplification. The overlap between the sensor noise and upconverted seismic noise indicates the realization of optimal damping.

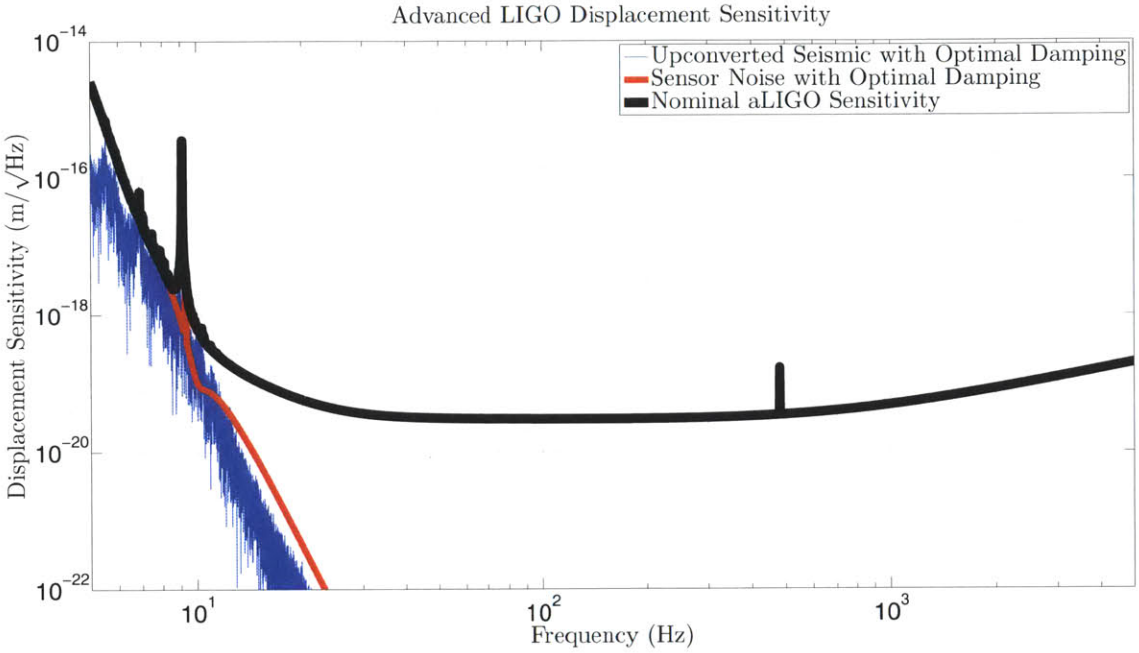

Figure 8-8: Simulated interferometer sensitivity with the optimal modal damping gains for the maximum seismic disturbance.

Since the majority of the observed cycles of an inspiral occur at the lowest frequencies, binaries of relatively low mass whose inspiral range is hardly influenced by adaptive modal damping may still benefit from this method. Figure 8-9 illustrates the number of visible cycles for two 1.4 solar mass neutron stars at 100 Mpc from the Earth. Two cases are considered: the dotted black line shows the sensitivity with the maximum seismic amplitude with minimum damping and the solid black line shows maximum seismic amplitude with the optimal damping. The dashed magenta line shows the inspiral signal at each frequency. The influence from adaptive modal damping on the inspiral range for this pair is negligible. However, The influence on observable cycles is non-negligible. With minimum damping, about 5000 cycles are visible. With optimal damping, about 16000 cycles are visible, over three times

more. Without adaptive modal damping Advanced LIGO would detect this source, but would determine its properties with much less precision.

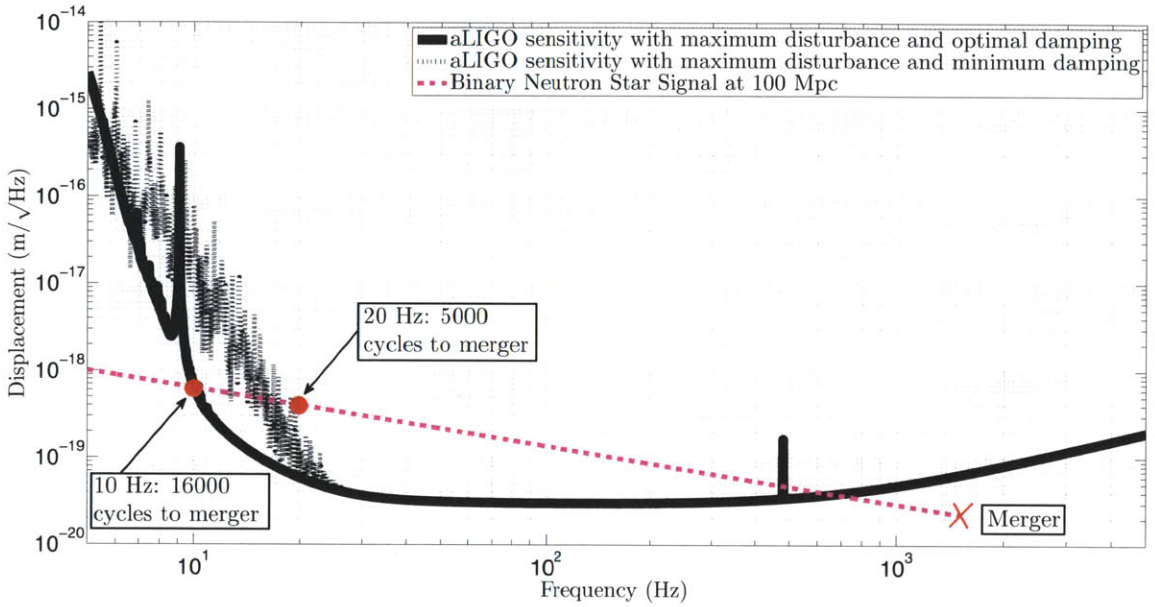

Figure 8-9: The solid black line is the Advanced LIGO sensitivity with the maximum seismic disturbance and optimal damping. The dotted black line is the sensitivity with the maximum disturbance and minimum damping. The dashed magenta line (courtesy Scott Hughes) is the characteristic signal produced by two inspiraling neutron stars of 1.4 solar masses each at a distance of 100 Mpc from the Earth.

8.4 Stochastic Background Results

This section presents the influence of adaptive modal damping on the sensitivity of the simulation to the stochastic background. The same cost functions that are optimal for binary inspirals are optimal here as well. These results are plotted in Figure 8-10 and summarized in Table 8.4.

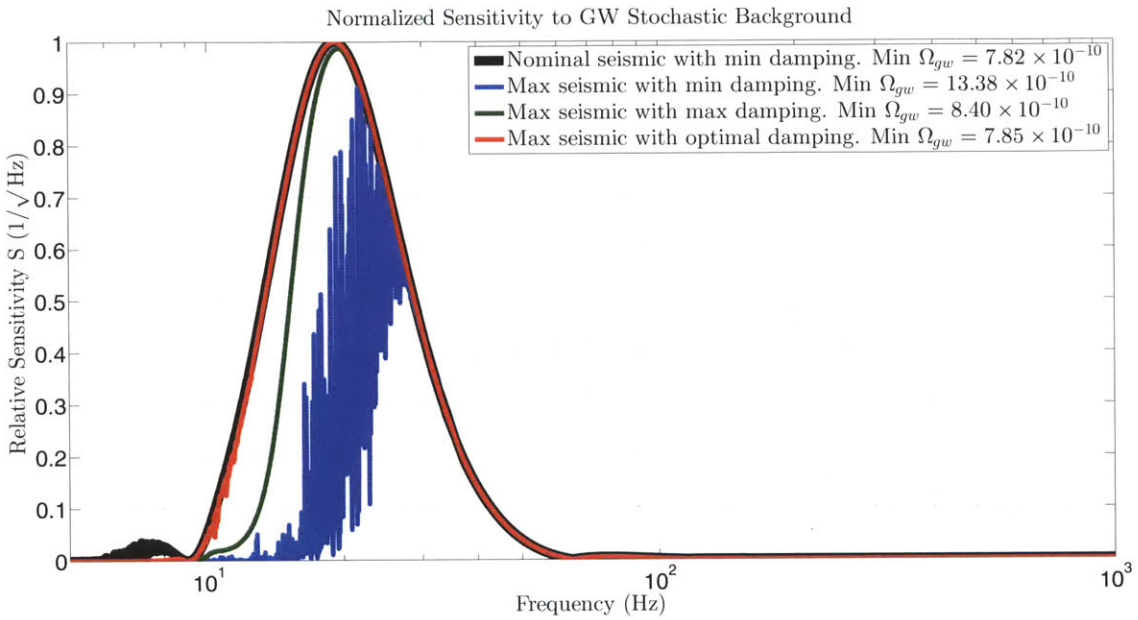

Figure 8-10: The relative sensitivity of the Advanced LIGO detector to stochastic gravitational waves at all frequencies. The listed energy densities, Ω_{gw}, are averaged over one year of data.

Figure 8-10 plots the stochastic sensitivity spectral density $S(f)$ given by Eq. (8.10). These curves show at what frequencies the system is most sensitive to the stochastic background under four different cases. The thick black line corresponds with the top row of the table and refers to the ideal case of minimum seismic amplitude and minimum damping where $\Omega_{gw} = 7.82 \times 10^{-10}$. The blue line corresponds with the second row and represents the case with the maximum seismic amplitude and minimum damping where $\Omega_{gw} = 13.38 \times 10^{-10}$. The green line is the third row and is the case where the seismic level is still maximum but now with maximum damping. Here $\Omega_{gw} = 8.40 \times 10^{-10}$. The red line is the last row where optimal damping is used to suppress the maximum seismic disturbance resulting in $\Omega_{gw} = 7.85 \times 10^{-10}$. In all

cases the curves in Figure 8-10 have been normalized by the same constant so that the largest value in the figure is 1.

Recall that better stochastic sensitivity is given by smaller Ω_{gw}, in contrast to inspiral sensitivity which improves with increasing range. Consequently, we can see that optimal damping at the worst seismic levels virtually reconstructs the best case scenario with the minimum seismic and damping levels. The difference in Ω_{gw} between the first and last rows in the table is only 0.4%. If the minimum damping level is maintained at the greatest seismic level, the sensitivity degrades by 41.6%. These results show that adaptive modal damping can preserve the sensitivity to the stochastic background under all seismic levels in this simulation.

Note though that the sensitivity with the maximum damping level is only worse than optimal damping by 6.5%. Thus, from the point of view of the stochastic background, adaptive modal damping is only slightly better than running LIGO with maximum damping on the pendulums. The reason for these results is that the stochastic sensitivity is focused in a relatively narrow frequency band around 20 Hz. This frequency is on the edge of where seismic upconversion and sensor noise amplification have influence. If one is interested specifically in the background radiation below 20 Hz, then the benefits of adaptive modal damping becomes more obvious.

Table 8.4: Advanced LIGO stochastic sensitivity results from the simulation. The last row is the case of optimal damping with the maximum seismic amplitude.

Seismic	$\zeta_{r,1}$	$\zeta_{r,2}$	$\zeta_{r,3}$	$\zeta_{r,4}$	$\Omega_{gw} \times 10^{10}$
0% max	0.00050	0.00050	0.00050	0.00050	7.82
100% max	0.00050	0.00050	0.00050	0.00050	13.38
100% max	0.20000	0.20000	0.20000	0.20000	8.40
100% max	0.00091	0.014	0.046	0.016	7.85

8.5 Conclusion

This chapter expands upon the adaptive modal damping technique developed in Chapter 7 by examining the benefit to the detection of binary inspirals and the stochastic background. The simulations presented here show that the adaptation

has the power to successfully select modal damping gains for the Advanced LIGO quadruple pendulum resonances that are optimal to the detection of multiple sources of gravitational waves. These damping values balance the influences of damping sensor noise amplification and nonlinear seismic upconversion. The time varying nature of the damping values keep the interferometer operating near optimal sensitivity under a large range of operating conditions, where static damping would only be optimal for a limited set of conditions.

For detecting binary inspirals the optimization is most effective when the combined mass is greater than 150 solar masses. High mass binaries produce high amplitude waves, but exist within LIGO's sensitive frequency band only at the extreme lowest end where the trade-off between upconversion and sensor noise amplification is expected to dominate. The detection of lower mass binaries is less effected because their sensitivity is concentrated at frequencies beyond the reach of adaptive modal damping. However, the problem of estimating the system's properties is still greatly influenced by adaptive modal damping. This estimation depends on the number of observed orbits a binary system undergoes before merging. The greatest number of these orbits occurs at the lowest frequencies. Thus, even if a binary is detected it's properties may not be well estimated unless the low frequency sensitivity is maximized.

Adaptive modal damping also benefits the observation of stochastic background radiation. The influence is less dramatic than the influence on high mass binary inspirals. Stochastic sensitivity is focused at frequencies at the extreme high end of the band where the trade-off between upconversion and sensor noise amplification dominates.

Other sources of gravitational waves potentially visible to Advanced LIGO occur at frequencies beyond the influence of adaptive modal damping. Pulsars radiate GWs too weakly by the time their spins decay to relevant frequencies. Supernovae are similarly expected to emit most of their radiation at higher frequencies.

Although the adaption in this chapter is applied to the damping of the quadruple pendulum mechanical resonances, there are many other possible areas of the

Advanced LIGO interferometers where adaptation could enhance their sensitivity in similar ways. These areas include damping of double and triple pendulums, test mass steering control, and the control of the seismic pre-isolation tables. In fact, the overall contribution from sensor noise is likely to be higher in practice since only the quadruple pendulum modes along the cavity axis are considered here. In all cases the benefits to the sensitivity are likely to occur at similarly low frequencies.

Future work should also study how the optimal adaptive solution depends on the length and angular control of the 4km Fabry-Perot cavities. The dependence will depend strongly on which nonlinearities are dominant. For example, stronger control of these cavities may suppress optical nonlinear behavior but enhance actuator nonlinear behavior.

The results in this chapter are limited to the scope of the simulations investigated. In practice the optimization of the cost functions, and the utility of adaptive modal damping will depend strongly on the nonlinear nature of the as-built Advanced LIGO interferometers.

Appendix A

Quadruple Pendulum Model

The quadruple pendulum x state space equations are listed here. $\mathbf{0}$ represents a zero matrix and $\mathbf{I}$ represents the identity matrix. All units are SI. The control inputs $\vec{u}$ are in units of Newtons and are ordered from stage 1 to 4. The seismic disturbance w is in units of meters and is applied to the pendulum's suspension point. The four outputs are the stage 1 through 4 displacements in units of meters.

$$\begin{bmatrix} \dot{\vec{x}} \\ \ddot{\vec{x}} \end{bmatrix} = \mathbf{A} \begin{bmatrix} \vec{x} \\ \dot{\vec{x}} \end{bmatrix} + \mathbf{B} \begin{bmatrix} w \\ \vec{u} \end{bmatrix}, \quad \vec{y} = \mathbf{C} \begin{bmatrix} \vec{x} \\ \dot{\vec{x}} \end{bmatrix} + \mathbf{D}\vec{u} \tag{A.1}$$

$$\vec{x} = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \end{bmatrix}^T \tag{A.2}$$

$$\vec{u} = \begin{bmatrix} u_1 & u_2 & u_3 & u_4 \end{bmatrix}^T \tag{A.3}$$

$$\mathbf{A} = \left[\begin{array}{cccc:c} & & \mathbf{0}_{4\times4} & & \mathbf{I}_{4\times4} \\ \hdashline -297.3 & 163.5 & 0 & 0 & \\ 162.9 & -267.2 & 104.2 & 0 & \mathbf{0}_{4\times4} \\ 0 & 57.8 & -74.2 & 16.4 & \\ 0 & 0 & 16.4 & -16.4 & \end{array}\right] \tag{A.4}$$

$$\mathbf{B} = \left[\begin{array}{ccccc} & & \mathbf{0}_{4\times5} & & \\ \hdashline 131.4 & 0.046 & 0 & 0 & 0 \\ 0 & 0 & 0.045 & 0 & 0 \\ 0 & 0 & 0 & 0.025 & 0 \\ 0 & 0 & 0 & 0 & 0.25 \end{array}\right] \tag{A.5}$$

$$\mathbf{C} = \left[\mathbf{I}_{4\times4} | \mathbf{0}_{4\times4} \right] \tag{A.6}$$

$$\mathbf{D} = \mathbf{0}_{4\times5} \tag{A.7}$$

Appendix B

MATLAB® Code

This code generates all the components of a modal damping feedback loop for the x dynamics of an Advanced LIGO quadruple pendulum. The top box is the master code that calls all the other boxes. The code calls a file `quad_x_model.mat`. This file is meant to contain the state space system given in Appendix A.

```
% Run this master code to generate the components of a modal damping
% loop for the x dynamics of an Advanced LIGO quadruple pendulum.

% system structure that contains all the plant information
% required for generating the modal damping.
    sys = generate_system;
% The estimator. estimator_opt generates optimization plots.
    [optimalR costmin] = estimator_opt(sys); % optimzation step
    estimator = generate_vel_estimator(sys,optimalR);
% The control
    control = generate_modal_vel_damp(sys);
% The coordinate transformation between modal and Cartesain
% coordinates
    [phi modes] = ModalDecomp2(sys.plant); % mode shapes and freqs
    Modal2Cartesion = phi'\eye(sys.ndof); % transformation matrix
% The full closed loop damped system
    closed_loop = generate_vel_closed_loop(sys,optimalR);
```

```
function sys=generate_system()
% Generates the modal damping quad pendulum model. The plant,
% sys.plant, has 4 outputs for the 4 stage displacements, ordered top
% to bottom. It has 5 inputs. The first is the seismic displacement
% input to stage 1. The following are the force inputs to the 4
% stages, ordered top to bottom. All units are SI.

load quad_x_model % loads state space model pendxmod
natural_damp = -0.001; % install some small natural damping

sys.plant = pendxmod; % model state space matrix
sys.plant.a(5:8,5:8) = natural_damp*eye(4); % add natural damping
sys.ndof=4;     % number of plant degrees of freedom
sys.w_size=1;   % number of seismic noise inputs
sys.v_size=1;   % number of sensor noise inputs
sys.u_size=4;   % Total number of control inputs.
% matrix transform that selects top mass actuator
sys.A=[1 0 0 0;0 0 0 0;0 0 0 0; 0 0 0 0];
% matrix transform that selects top mass sensor
sys.S=[1 0 0 0];
```

```
function estimator = generate_vel_estimator(sys,R)
% This function generates the modal damping state estimator. The
% estimator is a state space variable (estimator.ss) with as many
% outputs as there are modes (sys.ndof), and as many inputs as there
% are controllable stages and sensor signals (sys.ndof+sys.v_size).
% The first inputs are the control forces at each stage, starting at
% the top. The last inputs are the top mass sensor signals. The
% outputs are the modal velocity signals in order of mode frequency.
%
% The estimator is designed using the Linear Quadratic Regulator (LQR)
% method. The diagonal elements of the Q matrix are set by the
% amplitude of the modal velocity impulse responses. This emphasizes
% more accurate state estimation on the modes that contribute most to
```

```
% the overall suspension velocity. An alternative design to emphasize
% the modes that couple most strongly to the test mass would be to
% scale the aforementioned amplitudes by how much they couple to the
% test mass.
%
% Since the Q matrix is hardcoded here, the only user option in the
% estimator design is to chose the LQR R matrix. This matrix weights
% how much we trust the noisy sensor signals. Smaller R means better
% damping but more sensor noise amplification. R must be a positive
% scalar value.
%
% The first input to this function is the sys variable generated by
% the function generate_system. The second input is the user chosen
% value for R.
%
% LQR minimizes the following cost function:
%   J = [integral from 0 seconds to infinity]( q'*Q*q + z'*R*z)*dt
% where q is the modal estimation error, and z is a variable
% proportional to the estimator's sensor noise amplification.

v_size = sys.v_size; % number of sensor signals

% Converting model into a modal form
    [phi modes] = ModalDecomp2(sys.plant); % mode shapes and freqs.
% modal form of the state space A matrix
    Am = sys.plant.a;
    Am(sys.ndof+1:sys.ndof*2,1:sys.ndof)=...
    phi\Am(sys.ndof+1:sys.ndof*2,1:sys.ndof)*phi;
% C matrix that selects available sensor signals
    Ckn = [eye(v_size) zeros(v_size,2*sys.ndof-v_size)];
% modal form of Ckn
    Cknm = [Ckn(:,1:sys.ndof)*phi zeros(v_size,sys.ndof)];

% Calculating the diagonal Q matrix. Q weights the cost of error in
% the modal state estimation. Bigger Q values mean we want smaller
% estimation error.
```

```
% Pendulum inertia matrix:
    M = sys.plant.b(sys.ndof+1:end,1+sys.w_size:end)\eye(sys.ndof);
    Mm = diag(diag(phi'*M*phi)); % Modal pendulum inertia matrix.
    Q = zeros(2*sys.ndof); % initialization of Q for loop below.
% this loop calculates each diagonal element of Q one by one.
for ii=1:sys.ndof
    % These elements correspond to modal velocities.
    Q(ii+sys.ndof,ii+sys.ndof) = 1/Mm(ii,ii);
end
    % The LQR cost function is specified in terms of squared states.
    % This opperation makes Q consistent with the cost function.
    Q = Q.^2;

% Estimator feedback matrix. This is the optimal LQR estimator design
% determined by the choice of Q and R
    Lm = lqr(Am',Cknm',Q,R)'; % estimator feedback matrix.

% In the following estimator structure, only the field estimator.ss
% is needed. The other fields of merely informative to the user.
    Bm = sys.plant.b(:,sys.w_size+1:end);
    Bm(sys.ndof+1:2*sys.ndof,:) = phi\Bm(sys.ndof+1:2*sys.ndof,:);
    Bm = [Bm Lm];
    Cm = [zeros(4) eye(4)];
    estimator.L = Lm; % estimator feedback matrix
    estimator.B = Bm; % estimator b matrix
    estimator.A = Am-Lm*Cknm; % estimator a matrix
    estimator.C = Cm; % estimator c matrix
    estimator.ss = ss(estimator.A,estimator.B,estimator.C,...
    zeros(sys.ndof,sys.u_size+v_size));
```

```
function control = generate_modal_vel_damp(sys)
% This function compiles the modal damping feedback control. The input
% is the sys variable generated by generate_system. The output is a
% structure control, where the matrix field control.filters lists the
% modal damping filters for each mode on the diagonal.
```

```
6  %
7  % The damping filters are scaled so that the closed loop damping
8  % ratios are approximated the value multiplying each filter. For
9  % example, a damping ratio of 0.1 for mode 1 is obtained
10 % (approximately) by closing the mode 1 loop with
11 % 0.1*control.filters(1,1).
12
13 [phi modes] = ModalDecomp2(sys.plant); % mode shapes and frequencies
14  % Inertia matrix
15     M = sys.plant.b(sys.ndof+1:end,1+sys.v_size:end)\eye(sys.ndof);
16     Mm = phi'*M*phi; % Modal inertia matrix
17     Mm = diag(diag(Mm)); % Ensuring the matrix is purely diagonal.
18 % modal to euler coordinate transformation for actuation forces
19     modal2euler = phi'\eye(sys.ndof);
20
21 % compiling the modal damping filters
22 for ii=1:length(modes)
23    % basic modal feedback filter
24    control_filter=generate_one_dof_vel_filter3(modes(ii)*1,80);
25    % Applying some additional 10 Hz lowpass
26    control_filter = control_filter*bump(10,2,0.1);
27    % Scale the filter to be unity at the modal frequency.
28    control_filter = control_filter/...
29        abs(freqresp(control_filter,2*pi*modes(ii)));
30    % scale factor applied to the filter so that the user chosen
31    % modal feedback gain is approximately equal to the closed loop
32    % damping ratio for that mode.
33    dampscale = 4*pi*modes(ii)*Mm(ii,ii)/(phi(1,ii)*modal2euler(1,ii));
34    control_filter = dampscale*control_filter;
35    filter_matrix(ii,ii)=control_filter; %#ok<AGROW>
36 end
37 control.filters=filter_matrix;
38
39 function filter = generate_one_dof_vel_filter3(freq1,ph)
40 % Basic shape of the modal damping feedback filter. It is
41 % parameterized the modal frequency freq1, and the angle, ph, of
```

```
42 % the two complex low pass poles in the complex plane.
43
44 % complex low pass poles at 2.5 times the modal frequency
45 pair1 = 2.5*(2*pi*freq1)*exp([1i -1i]*pi*ph/180);
46 filter=zpk([],-pair1,1); % low pass filter
47
48 function gain_boost=bump(f0,Q,atten)
49 % f0 is the center freq (Hz), Q and atten control the width and height
50 %
51 % Brian Lantz Jan 23, 2003
52
53 top=1/Q;
54 bottom=atten/Q;
55
56 dhf=[1/(2*pi*f0)^2 2*top/(2*pi*f0) 1];
57 nhf=[1/(2*pi*f0)^2 2*bottom/(2*pi*f0) 1];
58 num=nhf;
59 den=dhf;
60 gain_boost = tf(num,den);
```

```
1  function [phi modes] = ModalDecomp2(plant)
2  % The function receives a state space system. Only the A matrix part
3  % of the state space is used here. There are two outputs, phi and
4  % modes. The first, phi, is the mode shape matrix. Each column
5  % corresponds to a mode frequency. The columns are sorted in order of
6  % ascending mode frequency. The rows correspond to the system's
7  % degrees of freedom. The second output is modes. This is a vector
8  % listing the mode frequencies in ascending order in units of Hz.
9  % Each value corresponds the respective column of phi.
10 %
11 % Example: The vertical (bounce) model of the monolithic quad.
12 % [phi modes] = ModalDecomp2(vertmod)
13 % phi =
14 %
15 %     0.1618   -0.3355    0.6926    0.6993
```

```
16 %      0.2840   -0.5244    0.5764   -0.7062
17 %      0.4496   -0.6247   -0.4306    0.1103
18 %      0.8313    0.4714    0.0511   -0.0042
19 %
20 % modes =
21 %
22 %      0.4428    0.9964    2.0061    3.4157
23 %
24 % The first column of phi is the mode shape of the 0.4428 Hz mode.
25 % Each row (unit length) represents the relative amplitude of motion
26 % for each stage.  The first row is stage 1, the second stage 2,
27 % then stage 3, and finally stage 4 the test mass. The numbers here
28 % indicate that each stage is moving in phase with the others because
29 % they are all the same sign, and each stage moves slightly more than
30 % the stage directly above. The second column represents the 0.9964 Hz
31 % mode. At this frequency the top three stages move in the opposite
32 % direction from stage 4. The fourth mode is largely the top two
33 % stages moving in opposite directions. The lower two stages
34 % hardly participate.
35 %
36 % Notes:
37 % The function assumes the A matrix has the
38 % form of
39 % A = [ zeros(ndof)    eye(ndof)
40 %         -M\K      -b*eye(ndof) ]
41 % Where ndof is the number of degrees of freedom, M is the mass
42 % matrix, K is the stiffness matrix and b is some quantity of
43 % velocity damping. All LIGO quad and triple suspension models have
44 % this form and can be used  with this function. The damping value is
45 % typically close to zero  (e.g. 0.001), if not exactly zero.
46
47 % number of suspension degrees of freedom
48     ndof = size(plant.a,1)/2;
49 % negative of the lower left corner of the A state space matrix
50     invMK = -plant.a(ndof+1:2*ndof,1:ndof);
51 % raw modes shapes and frequencies
```

```
52      [phi_temp modes] = eig(invMK);
53  % mode frequencies are converted from units of radians to Hz
54      modes(1,:) = sqrt(diag(modes))/(2*pi);
55  % modes are sorted
56      modes(2:end,:) = [];
57      [modes ind] = sort(modes);
58      phi = phi_temp(:,ind);
```

```
1  function [optimalR costmin] = estimator_opt(sys)
2  % This function finds the optimal R value for the estimator design.
3  % It finds this value by considering the closed loop settling time
4  % and sensor noise amplification.
5
6  noise_max = 10^-20; % m/sqrt(Hz): max 10 hz sensor noise amp allowed
7  f = 10;             % frequency to test LIGO sensitivity
8  Aw = 1.2508e-013;   % 10 Hz seismic noise
9  Av = 1.0012e-010;   % 10 Hz OSEM sensor noise
10 Av = Av/sqrt(2);    % accounting for two sensors in use
11 R_min = -10;        % minimum R value: 10^R_min
12 R_max = -4;         % maximum R value: 10^R_max
13 R_resolution = 50;  % Number of R values to test
14 R = logspace(R_min,R_max,R_resolution); % vector of R values to test
15
16 time = zeros(1,length(R)); % initialization of settling time vector
17 noise_amplitude = zeros(1,length(R)); % amplitude from sensor noise
18 for ii = 1:length(R)
19     % generate closed loop damping
20         closed_loop = generate_vel_closed_loop(sys,R(ii));
21     % calculate closed loop test mass settling time
22         time(ii) = calculate_settling_time(closed_loop(4,2));
23     %  calculate closed loop sensor noise amplification
24         noise_amplitude(ii)=Av.*(squeeze(abs(freqresp(...
25             closed_loop(4,3),f*2*pi))));
26 end
27 time_cost = (time/10).^2; % settling time cost
```

```
noise_cost = (noise_amplitude/noise_max).^2; % sensor noise cost
total_cost = time_cost + noise_cost; % total cost
[costmin min_index] = min(total_cost); % minimum cost
optimalR = R(min_index); % optimal R value

% Generate optimization plots
textsize = 25;
linesize = 5;
figure
subplot(211)
semilogx(R,time,[R(1) R(end)],10*[1 1],'m','Linewidth',linesize)
grid on
set(gca,'Fontsize',textsize,'XLim',[R(1) R(end)],'YLim',[0 99])
legend('Performance','Requirement',2)
ylabel('Time (s)','Fontsize',textsize)
title('x Stage 4 Settling Settling time vs R','Fontsize',textsize)
subplot(212)
semilogx(R,noise_amplitude,[R(1) R(end)],noise_max*[1 1],'m',...
    'Linewidth',linesize)
grid on
set(gca,'Fontsize',textsize,'XLim',[R(1) R(end)])
xlabel('R','Fontsize',textsize)
ylabel('Amplitude m/sqrt(Hz)','Fontsize',textsize)
title('x Sensor Noise Contribution to Bottom Mass at 10 Hz vs R',...
    'Fontsize',textsize)

figure
loglog(R,time_cost,R,noise_cost,R,total_cost,'Linewidth',linesize)
grid on
set(gca,'Fontsize',textsize,'XLim',[R(1) R(end)],...
    'YLim',[1e-5 1e3],'YTick',logspace(-5,3,9))
legend('Settling Cost, T_s^2','Sensor Noise Cost, max(N^2)',...
    'Total Cost, J_R',3);
xlabel('R')
ylabel('Cost Values')
title('x Estimator Cost vs. R')
```

```
64
65
66 function [time1 time2]=calculate_settling_time(system)
67 % This function returns the 1/e, time1, and 10%, time2, settling
68 % times of system, system.
69
70 final_time = 1000; % test time up to this many seconds
71 time_res = 1e-2;   % time resolution in seconds
72 error1 = exp(-1);  % 36.79% settling error
73 error2 = 0.1;      % 10% settling error
74
75 [y,t] = impulse(system,0:time_res:final_time); % impulse response
76
77 % Find 1/e settling time
78 ymax=max(abs(y));
79 ind1=find(abs(y)>=abs(ymax*error1));
80 ind1=max(ind1);
81 if ind1==length(t)
82     time1=final_time;
83 else
84     ind1 = ind1+1;
85     time1 = t(ind1);
86 end
87
88 % Find 10% settling time
89 ind2=find(abs(y)>=abs(ymax*error2));
90 ind2=max(ind2);
91 if ind2==length(t)
92     time2=final_time;
93 else
94     ind2 = ind2+1;
95     time2 = t(ind2);
96 end
```

```
function closed_loop = generate_vel_closed_loop(sys,R)
% Generate the closed loop modally damped pendulum. The order of
% the closed_loop inputs and outputs follow sys.plant, except there
% is an extra final input for the sensor noise.

[phi modes] = ModalDecomp(sys); % mode shapes and freqs
modal2cartesian = phi'\eye(sys.ndof); % modal to cart. transormation
control = generate_modal_vel_damp(sys); % modal feedback filters
estimator = generate_vel_estimator(sys,R); % estimator

% damping gains that damp to 1/e in 10 seconds for each mode
    gains = 1./(2*pi*modes(1:4)*9.0);
% mode 4 hardly couples to test mass, so use less gain
    gains(4)=gains(4)/5;

% build closed loop system
m2r_imp = sys.A*modal2cartesian;
filter_modal2cartesian = m2r_imp*diag(gains)*control.filters;
plant_est_cont=append(sys.plant,estimator.ss,filter_modal2cartesian);
Q=[-ones(3*sys.u_size+sys.w_size,1) zeros(3*sys.u_size+sys.w_size,1)];
for ii = 1:sys.u_size
    Q(ii,:) = [sys.w_size+ii -(2*sys.u_size+ii)]; % plant inputs
    % Estimator control feedback inputs
        Q(sys.u_size+ii,:)=[sys.w_size+sys.u_size+ii ...
            -(2*sys.u_size+ii)];
    if ii <= sys.v_size
        % Estimator sensor inputs
            Q(2*sys.u_size+ii,:) = [sys.w_size+2*sys.u_size+ii ii];
    end
    % filter inputs
        Q(sys.v_size+2*sys.u_size+ii,:) = ...
            [sys.w_size+sys.v_size+2*sys.u_size+ii,sys.u_size+ii];
end
% seismic inputs, coil inputs, and sensor inputs to estimator
    input = [1:sys.w_size+sys.v_size sys.w_size+2*sys.u_size+...
```

```
36          1:sys.w_size+2*sys.u_size+sys.v_size];
37 output = 1:sys.u_size; % same outputs as original plant
38
39 closed_loop = connect(plant_est_cont,Q,input,output);
```

Appendix C

Interferometer Feedback Filters

For the simulations presented in Chapter 8 the interferometer is controlled by applying feedback forces to all four stages of the pendulum that is not modally damped. The feedback filter transfer function for each stage is labeled C_i, where i refers to the number of the stage. s is the Laplace variable. Each filter receives the change in relative test mass displacement as the error signal. The actuators at stages 1 through 4 saturate at 449.6 mN, 12.7 mN, 0.48 mN, and 0.095 mN respectively.

$$C_1 = \frac{6.2443 \times 10^{13}(s + 8.145)^3}{s^2(s + 32.77)^3} \tag{C.1}$$

$$C_2 = \frac{8.1718 \times 10^9(s + 10.41)(s + 3.142)}{s(s + 60.68)} \tag{C.2}$$

$$C_3 = \frac{5.9507 \times 10^9(s + 46.85)}{(s + 273)} \tag{C.3}$$

$$C_4 = \frac{1.8507 \times 10^7(s + 182.2)}{(s + 1062)} \tag{C.4}$$

Bibliography

[1] J. Hough, S. Rowan, and B. S. Sathyaprakash, "The search for gravitational waves," *Journal of Physics B Atomic Molecular Physics*, vol. 38, pp. 497–+, May 2005.

[2] P. R. Saulson, "Gravitational Waves: From Astrophysics to Optics," LIGO, PowerPoint Presentation G050252, 2005.

[3] G. M. Harry and the LIGO Scientific Collaboration, "Advanced LIGO: the next generation of gravitational wave detectors," *Classical and Quantum Gravity*, vol. 27, no. 8, pp. 084 006–+, Apr. 2010.

[4] B. Shapiro, "Modal Control with State Estimation for Advanced LIGO Quadruple Suspensions," Master's thesis, Massachusetts Institute of Technology, Sep. 2007.

[5] N. A. Robertson, B. Abbott, R. Abbott, R. Adhikari, G. S. Allen, H. Armandula, S. M. Aston, A. Baglino, M. Barton, B. Bland, R. Bork, J. Bogenstahl, G. Cagnoli, C. Campbell, C. A. Cantley, K. Carter, D. Cook, D. Coyne, D. R. Crooks, E. J. Daw, D. B. DeBra, E. Elliffe, J. Faludi, P. Fritschel, A. Ganguli, J. A. Giaime, S. Gossler, A. Grant, J. Greenhalgh, M. Hammond, J. Hanson, C. Hardham, G. M. Harry, A. Heptonstall, J. Heefner, J. Hough, D. Hoyland, W. Hua, L. Jones, R. Jones, J. E. Kern, J. LaCour, B. T. Lantz, K. Lilienkamp, N. Lockerbie, H. Lück, M. MacInnis, K. Mailand, K. Mason, R. Mittleman, S. A. Nayfeh, J. Nichol, D. J. Ottaway, H. Overmier, M. Perreur-Lloyd, J. Phinney, M. V. Plissi, W. Rankin, D. I. Robertson, J. Romie, S. Rowan, R. Scheffler, D. H. Shoemaker, P. Sarin, P. H. Sneddon, C. C. Speake, O. Spjeld, G. Stapfer, K. A. Strain, C. I. Torrie, G. Traylor, J. van Niekerk, A. Vecchio, S. Wen, P. Willems, I. Wilmut, H. Ward, M. Zucker, and L. Zuo, "Seismic isolation and suspension systems for Advanced LIGO," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, J. Hough & G. H. Sanders, Ed., vol. 5500, Sep. 2004, pp. 81–91.

[6] J. Kissel, "Advanced LIGO Active Seismic Isolation," LIGO, PowerPoint Presentation G1100431-v1, Apr. 2011.

[7] P. Fritschel, D. Coyne, J. Giaime, B. Lantz, and D. Shoemaker, "Seismic Isolation Subsystem Design Requirements Document," LIGO, LIGO Engineering Document E990303, 2001.

[8] L. Carbone, S. Aston, R. Cutler, A. Freise, J. Greenhalgh, D. Heefner, J. Hoyland, N. Lockerbie, D. Lodhia, N. Robertson, C. Speake, K. Strain, and A. Vecchio, "Sensors and Actuators for the Advanced LIGO Mirror Suspensions," LIGO, LIGO Publication P1100208, 2011.

[9] J. Heefner, "UK Top Driver Pre-Production Prototype Test Plan," LIGO, Internal Technical Document T080014-v2, Apr. 2009.

[10] ——, "UK UIM Driver Pre-Production Prototype Test Plan," LIGO, Internal Technical Document T080021-v2, Dec. 2008.

[11] ——, "UK PUM Driver Pre-Production Prototype Test Plan," LIGO, Internal Technical Document T090164-v1, Apr. 2009.

[12] J. Miller, "On Non-Gaussian Beams and Optomechanical Parametric Instabilities in Interferometric Gravitational Wave Detectors," Ph.D. dissertation, University of Glasgow, Feb. 2010.

[13] M. Barton, N. Robertson, P. Fritschel, D. Shoemaker, and P. Willems, "Cavity Optics Suspension Subsystem Design Requirements Document," LIGO, Internal Technical Document T010007-v2, Oct. 2009.

[14] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*, 1977.

[15] J. Hartle, *Gravity: An Introduction to Einstein's General Relativity.* Addison Wesley, 2003.

[16] K. Thorne, "The Scientific Case for Advanced LIGO Interferometers," LIGO, LIGO Publication P000024-A, Jan. 2001.

[17] P. R. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors*, Saulson, P. R., Ed., 1994.

[18] K. S. Thorne, "Sources of gravitational waves and prospects for their detection." in *Recent Advances in General Relativity. Essays in honor of Ted Newman*, A. I. Janis & J. R. Porter, Ed., 1992, pp. 196–229.

[19] A. A. Wells, N. A. Gehrels, N. E. White, S. D. Barthelmy, J. A. Nousek, D. N. Burrows, P. W. A. Roming, K. O. Mason, G. Chincarini, and P. Giommi, "The SWIFT Gamma-Ray Burst Observatory," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, G. Hasinger & M. J. L. Turner, Ed., vol. 5488, Oct. 2004, pp. 403–414.

[20] B. Allen, "The Stochastic Gravity-Wave Background: Sources and Detection," in *Relativistic Gravitation and Gravitational Radiation*, J.-A. Marck & J.-P. Lasota, Ed., 1997, p. 373.

[21] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gursel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, and K. S. Thorne, "LIGO - The Laser Interferometer Gravitational-Wave Observatory," *Science*, vol. 256, pp. 325–333, Apr. 1992.

[22] B. C. Barish and R. Weiss, "LIGO and the detection of gravitational waves." *Physics Today*, vol. 52, pp. 44–50, Oct. 1999.

[23] R. Adhikari, "Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae," Ph.D. dissertation, Massachusetts Institute of Technology, United States – Massachusetts, Aug. 2004.

[24] P. Kwee, L. Barsotti, M. Evans, M. Zucker, P. Fritschel, and N. Mavalvala, "Proposal for a Filter Cavity Experiment at LASTI," LIGO, Internal Technical Document T1200024-v3, 2012.

[25] J. Kissel, "Calibrating and Improving the Sensitivity of the LIGO Detectors," Ph.D. dissertation, Louisiana State University, United States – Louisiana, Dec. 2010.

[26] R. Adhikari, S. Ballmer, and P. Fritschel, "Interferometer Sensing and Control Design Requirements," LIGO, LIGO Technical Document T070236, 2008.

[27] J. C. Driggers, M. Evans, K. Pepper, and R. Adhikari, "Active noise cancellation in a suspended interferometer," *Review of Scientific Instruments*, vol. 83, no. 2, p. 024501, Feb. 2012.

[28] L. Ruet, "Active Control and Sensor Noise Filtering Duality Application to Advanced LIGO Suspensions," Ph.D. dissertation, Institut National des Sciences Appliques de Lyon (INSA Lyon), Jan. 2007.

[29] K. A. Strain and B. N. Shapiro, "Damping and Local Control of Mirror Suspensions for Laser Interferometric Gravitational Wave Detectors," *Review of Scientific Instruments*, vol. 83, no. 4, p. 044501, 2012.

[30] L. Zuo, "Element and System Design for Active and Passive Vibration Isolation," Ph.D. dissertation, Massachusetts Institute of Technology, United States – Massachusetts, Feb. 2005.

[31] L. Zuo, J.-J. Slotine, and S. Nayfeh, "Model reaching adaptive control for vibration isolation," *Control Systems Technology, IEEE Transactions on*, vol. 13, no. 4, pp. 611 – 617, july 2005.

[32] L. Zuo and J.-J. E. Slotine, "Robust vibration isolation via frequency-shaped sliding control and modal decomposition," *Journal of Sound and Vibration*, vol. 285, no. 45, pp. 1123 – 1149, 2005. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0022460X04007588

[33] C. Torrie, "Development of Suspensions for the Geo 600 Gravitational Wave Detector," Ph.D. dissertation, University of Glasgow, Jan. 1999.

[34] N. A. Robertson, G. Cagnoli, D. R. M. Crooks, E. Elliffe, J. E. Faller, P. Fritschel, S. Goßler, A. Grant, A. Heptonstall, J. Hough, H. Lück, R. Mittleman, M. Perreur-Lloyd, M. V. Plissi, S. Rowan, D. H. Shoemaker, P. H. Sneddon, K. A. Strain, C. I. Torrie, H. Ward, and P. Willems, "Quadruple Suspension Design for Advanced LIGO," *Classical and Quantum Gravity*, vol. 19, pp. 4043–4058, Aug. 2002.

[35] N. Robertson, M. Barton, G. Cagnoli, C. Cantley, D. Coyne, D. Crooks, E. Elliffe, P. Fritschel, S. Gossler, A. Grant, A. Heptonstall, J. Hough, J. Jones, H. Lueck, R. Mittleman, M. Perreur-Lloyd, M. Plissi, D. Robertson, J. Romie, S. Rowan, D. Shoemaker, P. Sneddon, K. Strain, C. Torrie, H. Ward, and P. Willems, "Advanced LIGO Suspension System Conceptual Design," LIGO, Internal Technical Document T010103.

[36] A. V. Cumming, A. S. Bell, L. Barsotti, M. A. Barton, G. Cagnoli, D. Cook, L. Cunningham, M. Evans, G. D. Hammond, G. M. Harry, A. Heptonstall, J. Hough, R. Jones, R. Kumar, R. Mittleman, N. A. Robertson, S. Rowan, B. Shapiro, K. A. Strain, K. Tokmakov, C. Torrie, and A. A. van Veggel, "Design and Development of the Advanced LIGO Monolithic Fused Silica Suspension," *Classical and Quantum Gravity*, vol. 29, no. 3, p. 035003, Feb. 2012.

[37] S. Aston, "Advanced LIGO BOSEM Noise Measurement Report," LIGO, Internal Technical Document T0900496-v4, Mar. 2011.

[38] S. Aston and D. Hoyland, "Noise Prototype OSEM Design Document and Test Report," LIGO, Internal Technical Document T050111-02-K, Jan. 2008.

[39] M. Barton, "Models of the Advanced LIGO Suspensions in Mathematica^{TM}," LIGO, Internal Technical Document T020205-02D, Jun. 2006.

[40] M. Barton and N. Robertson, "Wire Attachment Points and Flexure Corrections," LIGO, Internal Technical Document T080096-00-K, May 2008.

[41] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*, 3rd ed. Addison-Wesley, San Francisco, 2002.

[42] D. Degenring, C. Froemel, G. Dikta, and R. Takors, "Sensitivity Analysis for the Reduction of Complex Metabolism Models," *Journal of Process Control*, vol. 14, pp. 729–745, 2004.

[43] R. Li, M. Henson, and J. Kurtz, "Selection of Model Parameters for Off-Line Parameter Estimation," *IEEE Transactions on Control Systems Technology*, vol. 12, pp. 402–412, May 2004.

[44] L. Ljung, Ed., *System Identification: Theory for the User*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1999.

[45] B. Shapiro, "Quad Pendulum Model Fitting Code," LIGO, Internal Technical Document T1100163, 2011.

[46] U. Dariusz, *Optimal Measurement Methods for Distributed Parameter System Identification.* CRC Press, Boca Raton, Florida, 2005.

[47] B. Shapiro, N. Mavalvala, and K. Youcef-Toumi, "Actuator Sizing of a Quadruple Pendulum for Advanced Gravitational Wave Detectors," *Proceedings of the 2011 American Control Conference, ACC 2011.*

[48] B. Shapiro, "Fitting the Quad Noise Prototype Model to Measured Data," LIGO, Internal Technical Document T1000458, Aug. 2010.

[49] G. Franklin, J. Powell, and M. Workman, *Digital Control of Dynamic Systems*, ser. Addison-Wesley world student series. Addison-Wesley, 1998. [Online]. Available: http://books.google.ie/books?id=z6UeAQAAIAAJ

[50] M. Evans, "Lock Acquisition in Resonant Optical Interferometers," Ph.D. dissertation, California Institute of Technology, United States – California, Dec. 2001.

[51] P. Fritschel and S. Waldman, "Reference Seismic Data for LLO," LIGO, Internal Technical Document T0900312-v1, Jun. 2009.

[52] P. Fritschel, "Advanced LIGO Systems Design," LIGO, Internal Technical Document T010075-v2, Sep. 2009.

[53] https://awiki.ligo-wa.caltech.edu/aLIGO/GWINC.

[54] P. Sutton, "S3 Performance of the LIGO Interferometers as Measured by Sense-Monitor," LIGO, LIGO Technical Document T030276, 2003.

[55] D. Brown, "Searching For Gravitational Radiation from Binary Black Hole Machos in the Galactic Halo," Ph.D. dissertation, University of Wisconsin-Milwaukee, United States – Wisconsin, Dec. 2004.

Public internal LIGO documents may be found at
https://dcc.ligo.org/cgi-bin/DocDB/DocumentDatabase/.