QUIET HYDRAULIC ACTUATORS FOR LIGO

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Corwin Hardham
May 2005

LIGO-P050031-00-R

© Copyright by Corwin Hardham 2005
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dan DeBra
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

J. Christian Gerdes

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dave Beach

Approved for the University Committee on Graduate Studies.

Abstract

A quiet hydraulic actuator has been designed and built for use in low-frequency isolation and alignment systems to stabilize the optics in gravitational wave detectors. By comparison to traditional electromagnetic and mechanical actuators, the hydraulic actuator is particularly well suited for this application due to its relatively high force, large stroke, low noise, and low frequency capabilities.

This type of actuator can be adapted to a variety of specifications. A matrix of design trades is described which characterizes an actuator for use in low-frequency alignment and isolation. However, these trades can also be used to scale the actuator design to meet other performance requirements.

Three generations of actuators have been developed. The final design incorporates a passive hydraulic damping mechanism that suppresses an undesirable resonance peak in the open loop response. The absence of the peak is shown to greatly facilitate feedback controller design and robustness.

Eight of these actuators were installed on a platform at LASTI – the LIGO test facility at MIT. A diagonalized MIMO controller was implemented on the platform incorporating both sensor correction and sensor blending techniques. Isolation performance has been demonstrated in translational directions from 70 mHz to 7 Hz.

Large scale installation of the actuators has recently been completed at LIGO-Livingston; one of two LIGO observatories in the U.S. Prior to installation at Livingston, the observatory was inoperable during the daytime because of severe cultural noise. However, with the aid of quiet hydraulically actuated platforms, the observatory may now operate at all times of day.

Acknowledgements

to Dan DeBra for red wine, long bike rides and revealing beauty in design and mathematics

to Brian Lantz for being the fearless leader

to Wensheng Hua for showing me how to see the world differently

to the Byer/Fejer group for great support and a wonderful working environment

to Mike Henessey and Jim Perales for building a great lab

to the MIT LIGO Lab for graciously hosting the HEPI project and the author on multiple visits

to the LIGO collaboration and the NSF for generous support and an inspiring project.

to my dearest Emily

R.I.P. Chi Theta Chi

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The observation of gravitational waves will open a new window on the universe. Several countries have engaged in the construction of observatories to pursue this purpose. In the United States, there are two installations of the Laser Interferometer Gravitational-wave Observatory (LIGO) that have reached the initial phase of commissioning. However, in order to succeed in measuring gravitational waves, these observatories must attain extraordinary levels of displacement sensitivity. Terrestrial seismic noise is one of many disturbances that must be overcome before LIGO can attain this sensitivity. In this work, a new actuator is developed to aid in the active suppression of seismic noise and offer long-term alignment for the LIGO observatory.

1.1 Gravitational Waves

The existence of gravitational waves was first predicted by Einstein in 1916 to provide a causal explanation to the gravitational force exerted by an accelerating mass. By expressing gravitational force with the wave equation, it ceases to act instantaneously, as suggested earlier by Newton, and instead travels at the speed of light.

At this time, gravitational waves have yet to be observed, but there has been an indirect observation made by Hulse & Taylor [1], [2]. Through careful study of the orbital decay in a neutron star binary system, Hulse & Taylor found the decay rate to be in excellent agreement with the predicted energy lost to gravitational radiation.

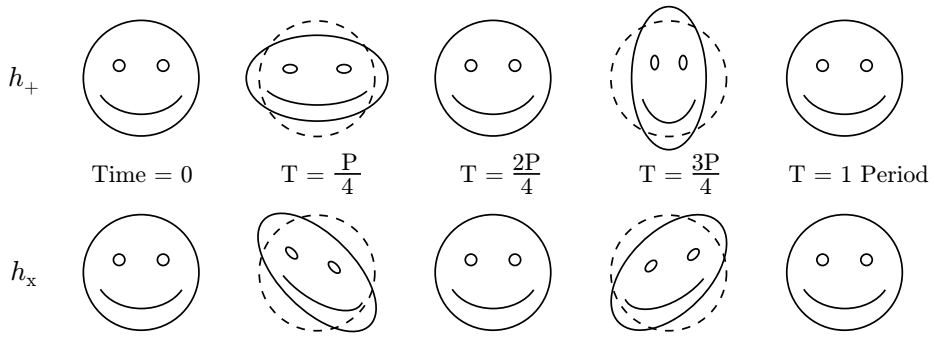

Figure 1.1: The effect of a gravitational wave passing perpendicularly through an object. Note the plus h_+ and cross h_x polarizations. Courtesy of Brian Lantz.

Gravitational waves are differential planar strain waves, meaning that an object subjected to a gravitational wave is alternatingly stretched in one axis while compressed in the orthogonal axis (figure 1.1). Similar to an electromagnetic wave, there is both a plus, h_+, and cross, h_x, polarization.

The strain produced by an gravitational wave is tiny: $h \sim r_{s1}r_{s2}/(r_oR)$ where r_{s1} and r_{s2} are the Schwarzschild radii of the masses involved ($r_s = 2GM/c^2$). The remaining variables are G, is the gravitational constant; R, the distance to the source; M, the mass of each star and c, the speed of light. The ratio of G^2/c^4 is so small that the only measurable sources of gravitational waves are produced by masses on the order of a solar mass.

1.1.1 Gravitational Wave Sources

Despite the need for large masses, there are a variety of sources that may be powerful enough to be detected. A commonly discussed source is the coalescence of two compact objects such as neutron stars or black holes. For an estimate of the strain that these sources would produce on Earth, consider a pair of 1.4 solar mass ($\sim 3\times10^{30}$ kg) neutron stars located in one of the closest galaxies (the Virgo Cluster, for example) at a distance R of approximately 15 megaparsec or 4.5×10^{23} meters. Moments before impact these stars may orbit each other at frequencies approaching 400 Hz. The resulting strain on Earth will be approximately 10^{-21} [14].

1.2 Gravitational Wave Detection

In spite of the extraordinarily small strain on earth, several methods have been proposed to detect gravitational radiation for astrophysical observation. These include *bar detectors* which consist of a large suspended mass whose longitudinal flexible mode is at a frequency of about 1 kHz for which there are anticipated gravitational radiation sources. However, in more recent times, most research effort has been directed toward laser interferometric detectors, and at this time, several countries have recently commissioned detectors of this type.

Rainer Wies first proposed a practical interferometric detection scheme in the 1970's [3]. However, the first to embark on the path toward building a interferometric detector was a British-German group known as GEO [4]. This group is responsible for the GEO600 detector in Hanover, Germany. Subsequent to this, several countries have constructed detectors: the Japanese built a detector with impressive sensitivity for its size called TAMA [5], there is an Italian/French effort known as Virgo [6], and the Australian group, ACIGA, built a detector in Western Australia [16], [7]. LIGO is both larger and more sensitive than any other existing detector, but for each of these, the mechanism for detection remains fundamentally the same.

1.2.1 Laser Interferometric Detection

Laser interferometry is attractive for gravitational wave detection because of its inherently high displacement sensitivity and the capability to project light over large distances. The high displacement sensitivity is necessary because of the weak coupling of gravitational waves to masses on Earth, and since the detection of gravitational waves is based on measuring a strain, the signal is amplified by a large baseline.

The Michelson configuration is typical for all detectors, and the fundamental aspects of the interferometer arrangement differ little from the 1887 original [17]. The important distinction for gravitational wave detection is that the mirrors are not connected to a rigid structure. Instead, each mirror is freely suspended to respond to gravitational wave effects.

In the event of a gravitational wave, the space between the freely suspended

mirrors, or test masses, is stretched much like the smiley face in figure 1.1. The interferometer is sensitive to this distortion and outputs the resulting displacement signal. It is important that the arms of the interferometer be arranged nearly orthogonally so that the detector is sensitive to all orientations of gravitational wave.

Unfortunately, for terrestrially rooted suspensions[1], a freely suspended mass is an impossible ideality. No realizable suspension system can support a mass as though it were floating in inertial space. Practical systems approach this ideal through the application of multiple stages of suspension for isolation.

1.2.2 LIGO

There are two installations of LIGO: one in Hanford, Washington and another in Livingston Parish, Louisiana. Both of these observatories are now operational. The current configuration of each observatory is commonly known as Initial LIGO with the expectation of an Advanced LIGO configuration by approximately 2010. The work discussed in this thesis was originally intended for Advanced LIGO, but due to an unexpected excess of ground motion at the Livingston observatory (LLO), the hydraulic actuators discussed here have been commissioned as part of Initial LIGO at Livingston.

The LIGO observatories consist of two 4 km long beam tubes arranged orthogonally to one another (figure 1.2). Each beam tube contains one arm of a Michleson interferometer with a Fabry-Perot resonant cavity. The end mirrors of the Fabry-Perot cavity are contained in Beam Splitter Chambers (BSC) at either end of each 4 km long beam tube. The BSC at the Corner Station houses the beam splitter and the surrounding Horizontal Access Modules (HAM) contain a variety of support optics for the main interferometer.

Initial LIGO is intended to sense gravitational waves at frequencies between 50 and 7000 Hz from sources within 15-20 megaparsecs. The goal of Advanced LIGO is to increase the sensitivity of the instrument to distances approaching 200 megaparsecs over a frequency range of 10 to 10,000 Hz. This translates into a factor of $\sim$10

[1] An alternative to the terrestrial observatories is a space based observatory such as the proposed LISA mission. The masses in the LISA satellites will clearly be immune from seismic disturbances.[8]

Figure 1.2: The Laser Interferometric Gravitational-wave Observatory (LIGO). The beam tubes are 4 kilometers long and contain the two axes of the LIGO interferometer. The test masses are suspended within the Beam Splitter Chambers (BSC) while the Horizontal Access Module (HAM) chambers contain a variety of support optics.

reduction in the strain-equivalent noise floor (figure 1.3) which requires nearly every aspect of the detector to be improved or replaced with the notable exception of the vacuum envelope.

Among the many noise sources that threaten the successful operation of the LIGO detector, ground motion induced vibrations can disrupt the operation of the interferometer and add noise at the low end of the gravitational wave detection band. The sources and magnitude of seismic disturbances vary with frequency (figure 1.3). Overall, the root-mean-square (rms) of the ambient ground motion at each site is approximately 1 μm. Much of the spectral contribution to this rms motion comes from the so called microseismic peak in the 0.1-0.3 Hz band. The microseismic peak results from coastal ocean water waves exciting surface waves along the Earth's crust. Another notable disturbance source is human activity which contributes largely between 1 and 10 Hz. This is of particularly apparent at LLO, where commercial logging in the surrounding forest causes a factor of $\sim$10 increase in motion during the daytime. At very low frequencies, the surface of the Earth undergoes a tidal motion on the order of 200 μm peak to peak [9] caused by attraction to the sun and the moon. Seasonal temperature variations may also introduce annual length variations as large as 1 mm.

In the presence of these disturbances, suspensions for LIGO must provide both alignment and isolation. Alignment control is important in all degrees-of-freedom (DOF) to aim the interferometer beam at the center of the test mass, and particularly along the beam line, to control the length between the two ends of each optical cavity. To achieve lock and limit noise in the output from the interferometer, the rms fluctuations in length between test masses at the ends of each optical cavity must be limited to 10^{-13} m for Initial LIGO, and 10^{-14} m for Advanced LIGO (within the detection band). Isolation is necessary to reduce test mass motion so as to enable gravitational wave detection.

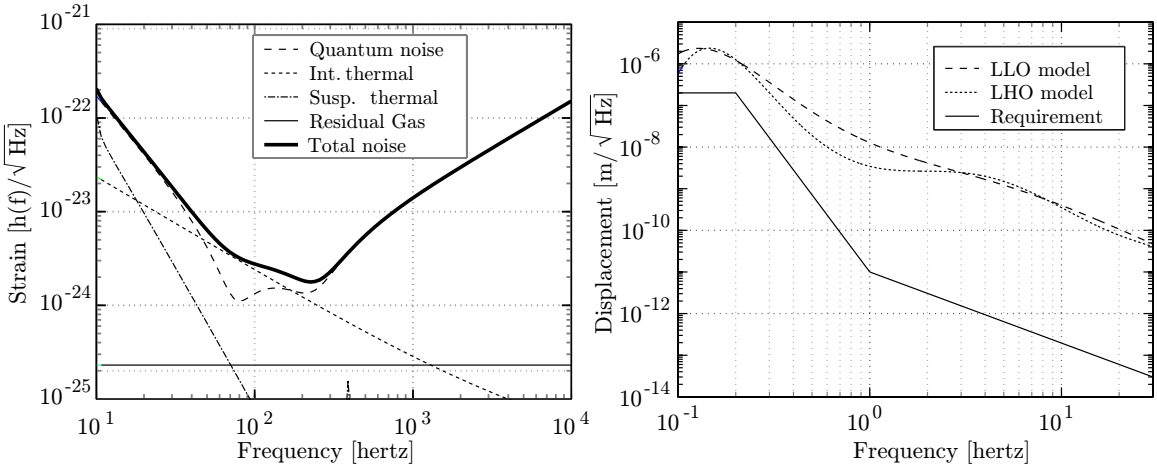

Figure 1.3: Current estimates of the contributions of various detector noise sources to the Advanced LIGO strain (h) sensitivity level (left), and the resulting requirement on displacement noise for the second stage of the two stage active isolation system (right). Models of the ground motion measured at LLO and LHO are included (right). Note that the ground motion around 1 Hz at the Livingston site (LLO) is much larger than that at Hanford (LHO). [10] [12]

1.2.3 Suspension Systems in LIGO

The suspension systems in LIGO (figure 1.4) include both passive and active systems distributed over multiple stages. The overall design intent is to employ small throw, low noise actuators near the test mass, and once these actuators approach the limits of their range, reallocate the control effort to larger, noisier actuators closer to the ground. This minimizes coupling of actuator noise to the test mass, and leads to a design where the actuators are well matched for their application.

Initial LIGO

The suspension system for initial LIGO consists of a single, two-wire pendulum hanging from a passive isolation stack. The passive isolation stack provides isolation in the vertical direction and rests on the support table, which is held externally to the vacuum chamber by the crossbeams and eventually the piers (figure 1.5). The isolation system that interfaces the pier tops to the corners of the crossbeams is the External

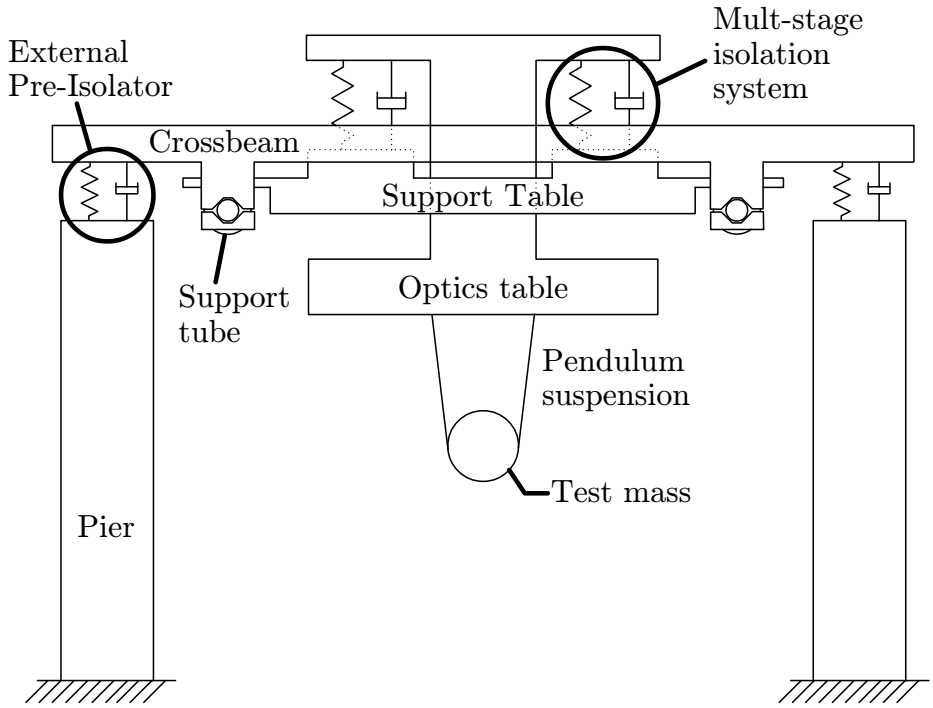

Figure 1.4: A schematic of the suspension system used in LIGO. The suspension system from the ground up starts with the piers which support the External Pre-Isolator (EPI) system. The EPI system supports the crossbeams which pierce the vacuum chamber (not shown) and carry both the multi-stage isolation system and the pendulum suspension.

Pre-Isolator (EPI) which is the stage of interest for this work.

The position of the test mass at the bottom of the pendulum is controlled by four electromagnetic actuators. The control authority of these actuators is limited and when the motion of the mass begins to exceed their capability, control effort is diverted to the EPI system.

The EPI system, as originally conceived for Initial LIGO, is capable of horizontal motions (X and Rz), and is actuated by piezoelectric stacks. The range of motion of the piezoelectric stacks is limited to ± 90 μm, and since the piezoelectric material is stiff, it is difficult to accomplish vibration isolation at frequencies as low as 1 Hz. Nevertheless, the piezoelectrically actuated EPI system (PEPI) is capable of operating with the test mass actuators to acquire and maintain lock of the interferometer for

limited periods of time.

The PEPI system at LLO has been replaced with a Hydraulically actuated External Pre-Isolator (HEPI) because the range of motion of PEPI is not sufficient to track the Earth tides and the extreme cultural noise observed at the LLO site. The increased cultural noise at LLO can be attributed to the commercial logging activities in the forest surrounding the observatory and a nearby freight railway track.

Advanced LIGO

The proposed suspension design for Advanced LIGO [12] is considerably more sophisticated than Initial LIGO. The test mass is the lowest mass in a quadruple pendulum that hangs from a two-stage, electromagnetically actuated isolation platform. The two stage system rests on the support table which is supported by the HEPI system.

Similar to Initial LIGO, the HEPI system is included in the proposed Advanced LIGO suspension to attenuate large amplitude disturbances so that systems within the vacuum tank only need to operate about their centered position. This is reduces noise, and the range of motion required of stages located inside the vacuum chamber where the electromagnetic actuators are ill suited to holding large, long-period offsets.

Specifications for HEPI

For both generations of suspension, the specifications for the HEPI system are fundamentally the same. The HEPI system must be capable of translating ± 1 mm in X, Y and Z while providing a factor of 10 of isolation from frequencies as low as the microseismic peak at $\sim 1/6$ Hz to frequencies where cultural noise is significant, 1-5 Hz. To reach these objectives, the HEPI system must be at least a factor of 10 quieter than the ground and capable of providing adequate bandwidth to meet the isolation requirements.

Stiff vs. Soft Suspensions

A common classification for suspended stages in LIGO and elsewhere in the literature, is either stiff or soft, based on the natural frequency of the suspended platform with

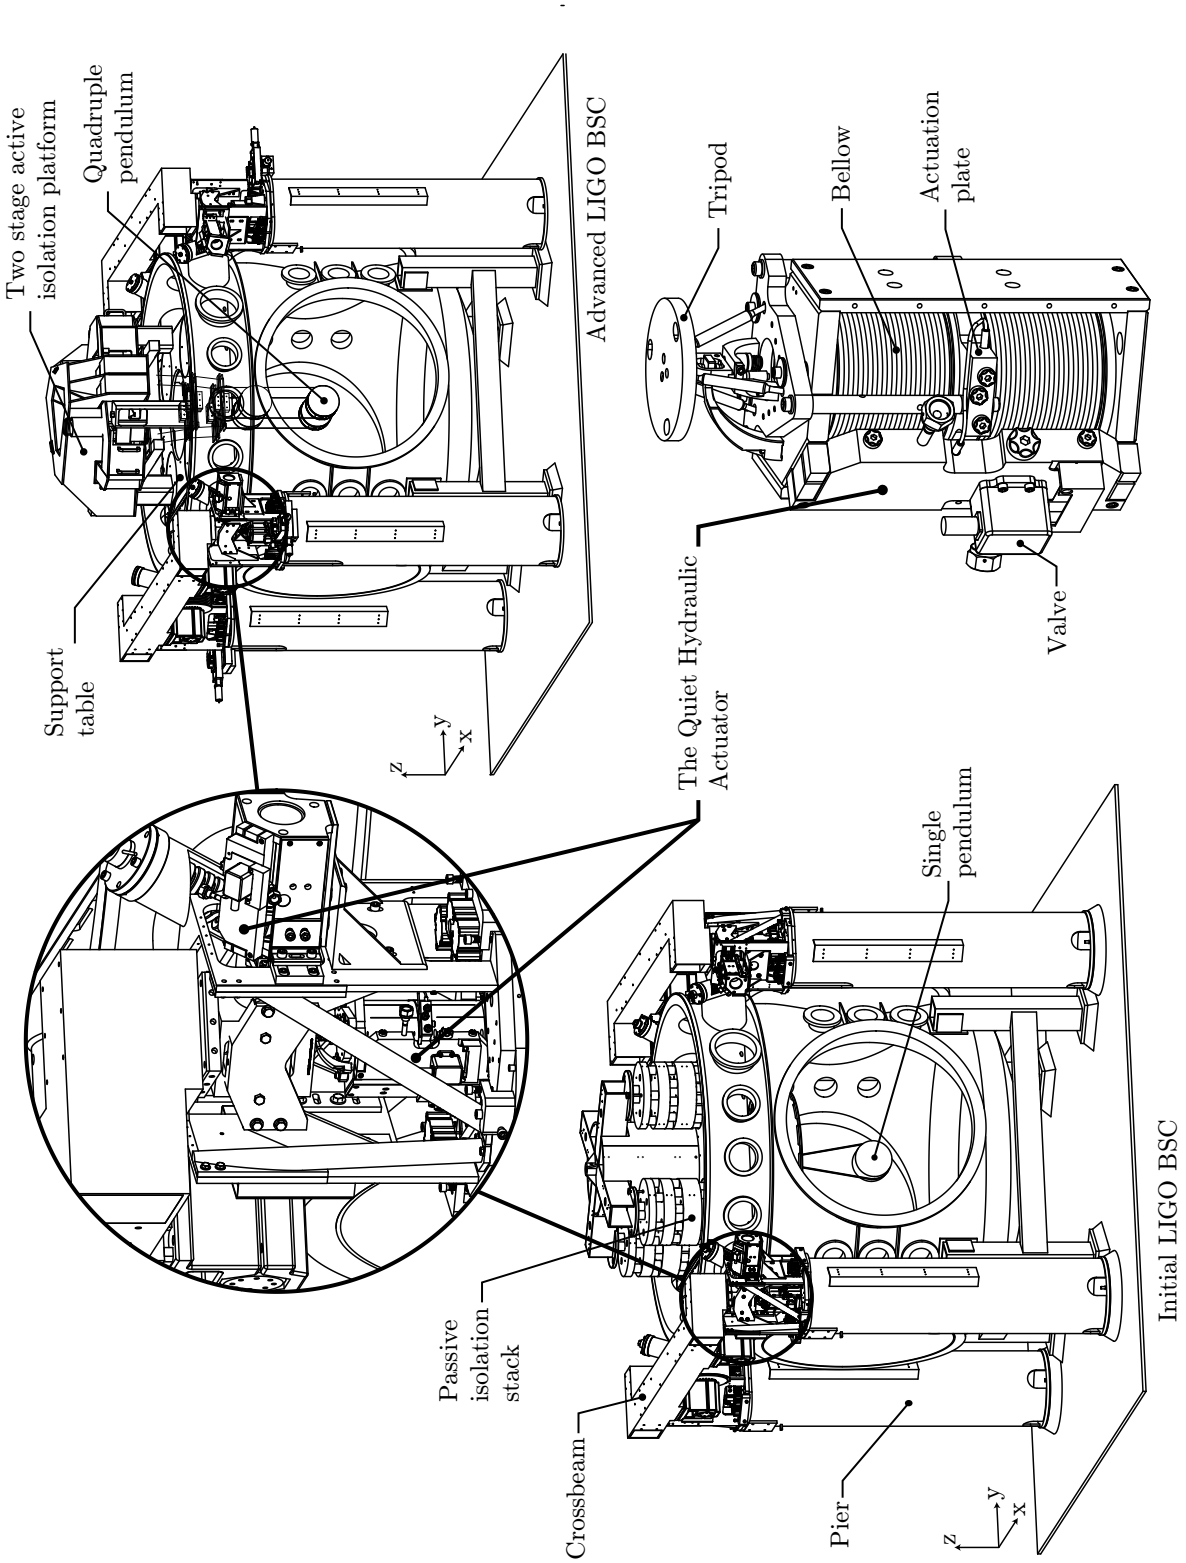

Figure 1.5: The suspension system for Initial LIGO, and the proposed suspension design for Advanced LIGO as applied to the BSC tank. Both the Initial and Advanced suspension systems include an external stage.

respect to where the attenuation is needed. A soft suspension attempts to attenuate motion passively by setting the natural frequency below the frequencies where the isolation is required. Conversely, the natural frequency of a stiff system is set higher than the isolation band, and low-frequency isolation is attained through active control.

The decision between a stiff or soft system lies in the choice of natural frequency. The natural frequency determines the stiffness of the offload springs that support the weight of the suspended platform, and the choice of natural frequency dictates much about the design of the stage. The advantage of a soft system is that it requires less actuator authority because the offload spring stiffness is lower.

The major disadvantage of a soft system is long term stability. The position of a softly suspended platform is very sensitive to temperature variations due to changes in spring stiffness and length (Appendix C.2). Since the static sag of even a 1 Hz suspension is 250 mm, a small fluctuation in spring stiffness can cause major changes in position. Furthermore, for the large payloads as found in LIGO, a *soft* spring is either excessively large or exhibits stress levels that pose concerns of safety and noise caused by creep events.

Consequently, the EPI system is stiffly suspended at 8 Hz, and the two-stage active isolation system proposed for Advanced LIGO also features natural frequencies at or above 2 Hz.

1.3 The Quiet Hydraulic Actuator

There are several examples in the literature [20],[18] that describe systems capable of meeting the proposed alignment and isolation objectives set forth for HEPI. The unique challenge posed by the HEPI system is to develop an actuator capable of meeting these performance objectives, but with the formidable low-frequency, range and low-noise capabilities requisite to the LIGO application.

The EPI payload weighs approximately 2000 kg, and with an 8 Hz natural frequency, the actuator must deliver 1250 N of force in order to displace 1 mm. The actuator must be capable of maintaining these high force levels indefinitely while not exceeding a displacement noise level at least a factor 10 less than the ground.

	Force	Velocity	Stiffness	Displacement	Stiction	Hysteresis	Mech. noise
Quiet	high	low	med	med	low	low	low
Hydraulic	$\oplus$	$\ominus$	$\ominus$	$\ominus$	$\oplus$	$\oplus$	$\oplus$
Mechanical	high	low	high	high	high	low	high
	$\oplus$	$\ominus$	$\oplus$	$\oplus$	$\otimes$	$\oplus$	$\otimes$
Electromagnetic	high	high	low	high	low	low	low
	$\oplus$	$\oplus$	$\otimes$	$\oplus$	$\oplus$	$\oplus$	$\oplus$
Piezoelectric/	med	high	high	low	low	high	low
Magnetostrictive	$\ominus$	$\oplus$	$\oplus$	$\otimes$	$\oplus$	$\otimes$	$\oplus$

Table 1.1: Candidate actuator technologies for the EPI system. A 'high' above a '$\oplus$' indicates that the corresponding quality is prominent and desirable. A '$\ominus$' indicates the quality is acceptable and a '$\otimes$' designates an attribute as unsuitable. For example, the high force of the hydraulic actuator is desirable while the medium stiffness is merely acceptable.

1.3.1 Candidate Actuators

There are several types of actuators that potentially could meet the requirements of the EPI platform. These potential candidates are listed in table 1.1. The mechanical category includes various ball screw and lever arrangements. These are unacceptable because of the noise inherent in rolling interfaces. Some mechanical flexure/non-contact electromagnetic combinations were considered, but the resulting package was excessively large. The absence of stiffness found in direct actuation with a linear electromagnetic actuator results in a poorly behaved system (such an arrangement was attempted as discussed in Appendix C.1). It is also challenging to obtain an electromagnetic actuator with sufficient force and range of motion, especially with the required lateral clearance to accommodate a multi-DOF installation. Piezoelectric materials fail because of their limited range and high hysteresis. A piezoelectric stack large enough to meet the range requirements of EPI would be very sensitive to shear failure from transverse forces.

The quiet hydraulic actuator is the only candidate that is acceptable in every category. The success of hydraulics in this application is made possible by the application of the principles which define quiet hydraulics.

1.3.2 Quiet Hydraulics

The field of quiet hydraulics was pioneered by the work of Dan DeBra of Stanford University (formerly Lawrence Livermore National Labs) and Taco Viersma of Delft University. In general, quiet hydraulic systems exhibit much of the high force character of conventional hydraulics but with significantly lower noise at the cost of decreased bandwidth.

The low noise performance can be attributed to the defining attributes of quiet hydraulics: laminar flow and the absence of frictional interfaces. Conversely, most conventional hydraulic systems operate in the turbulent flow regime. The pressure fluctuations implicit in turbulent flow couple directly into piston position causing displacement noise, and often, an audible hum. Quiet hydraulic systems forfeit the speed of conventional hydraulics to attain low noise levels by limiting flow velocities to remain laminar.

The quiet hydraulic actuator avoids frictional interfaces by incorporating a pair of flexible bellows instead of the conventional piston and cylinder arrangement (figure 1.6). Flow in and out of the bellows is modulated by a network of variable hydraulic resistances that resemble a Wheatstone bridge. The variable resistances are nozzle/flapper impedances that are embedded in a commercially manufactured hydraulic servo valve.

1.4 Outline of the Thesis

Commercially available servo valves are intended for use in turbulent systems, and thus require modification for use in the quiet hydraulic actuator. Chapter 2 addresses the design of laminar flow nozzles and other modifications to the stock servo valves.

The design of the quiet hydraulic actuator is motivated by mathematical models

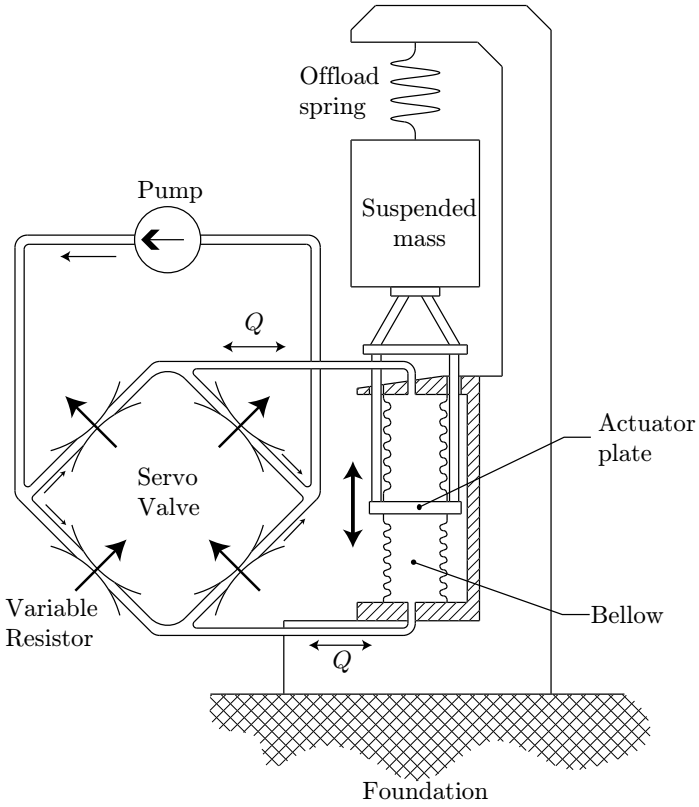

Figure 1.6: A schematic of the hydraulic actuator. A pump supplies pressurized fluid to the servo valve which modulates flow, Q, to the bellows on either side of the actuator plate. The actuator plate displaces according to the differential pressure on either side of it and applies force to the mass suspended by the offload spring.

described in Chapter 3. Included in Chapter 3 is a design trades chart that shows how various parameters of the actuator may be adjusted to meet alternative specifications.

The mechanical design of the actuator and the platforms that support the actuator is developed in Chapter 4. Based on the discovery in Chapter 4 of a undesirable resonance in the bellows, Chapter 5 discusses the redesign of the bellows and the development of a novel hydraulic bypass network.

A discussion of the various control techniques applied to hydraulically actuated platforms is given in Chapter 6, and Chapter 7 provides ideas for future installations and actuator designs.

1.5 Prior Art

The use of bellows as a hydraulic or pneumatic cylinder is an old, perhaps ancient, concept. The arcane bellows used to move air or fluids in pumps to stoke fires or irrigate fields are not functionally dissimilar from the basic geometry presented here. More coincident arrangements appeared in applications like vehicle braking in the early 1900's (for example, see patent no. 1744241 by Pirrle). In more recent times, the differential bellow arrangement of the quiet hydraulic actuator has become a common solution for obtaining sub-micron level precision with hydraulics [21].

However, it is difficult to find examples of this differential bellows architecture developed specifically for feedback control in a multi-DOF system. Hence, the actuator design discussed here is unique not in its underlying functionality, but rather, in the specifics that make it suitable for active alignment and isolation.

The suitability for control is realized through careful design of the bellows and the bypass network developed for this application. Slocum [21] mentions need to consider the design of the bellows for a bellows based actuator, but does not mention the specific design feature that is found to be important here: the *breathing* stiffness of the bellows (discussed in Chapter 5). Similarly, the bypass network, a laminar flow hydraulic resistor and capacitor has been demonstrated as a effective means to damp undesirable modes by DeBra [22] for passive isolation systems. However, the application of this passive network into an active differential bellows system is unique.

1.6 Research Contributions

The primary contribution of the work presented in this thesis is the development of a quiet hydraulic actuator proven to be suitable for multi-DOF control. This result is the culmination of sustained development both in the design of the actuator and the surrounding systems.

The final actuator prototype used in LIGO is the last of three prototypes developed over the course of this research. The final design incorporates two novel features to a quiet hydraulic, two-bellow configuration. Firstly is the low-inductance, passive

hydraulic bypass network that suppresses the bellows' *breathing resonance* observed in the first actuator prototype. Secondly, the final actuator prototype incorporates custom designed bellows with a convolution shape that maximizes the ratio of breathing stiffness to axial compliance. Thirdly, to operate in the high quality clean room environment of LIGO, the actuator also includes an embedded bleed network to ensure that hydraulic fluid will remain sealed in the actuator during the removal of entrapped air.

Fourthly, several new control techniques were developed for the hydraulic actuator installed in the LIGO observatory. These include the application of sensor correction and sensor blending generalized to modal control of an eight actuator platform. Finally, the performance requirements for LIGO have been experimentally demonstrated.

Chapter 2

The Hydraulic Flapper Valve

The hydraulic actuator converts differential flow into velocity of the suspended mass. Similarly, the purpose of the hydraulic flapper, or servo, valve is to convert constant flow provided by a pump into a controllable source of differential flow to drive the actuator.

Hydraulic servo valves are commercially available from a variety of manufacturers, but in their stock configuration, these valves are intended to operate in the turbulent flow regime. In this experiment, turbulence must be avoided to reach the low noise criteria for LIGO, and while the valves may be operated at low flow rates, the stock nozzles inside of the valves are designed to stimulate turbulence. In this chapter, modifications to the valve nozzles are discussed along with techniques to mitigate the nonlinearity associated with operating the valve in the laminar regime.

The servo valve is a hydraulic analog of a Wheatstone bridge. The value of the variable resistances is dependent on the electrical control current applied to the servo valve. By applying current to the valve, it is possible to unbalance the hydraulic bridge and create differential pressure between the corners of the bridge (C1 and C2 in figure 2.1).

A servo valve is typically comprised of a steel flapper controlled by an electromagnetic torque motor operating between two nozzles. The position of the flapper is dependent on the electrical current input into the valve torque motor. The hydraulic resistance of the flapper valve is determined by both the resistance of the nozzles and

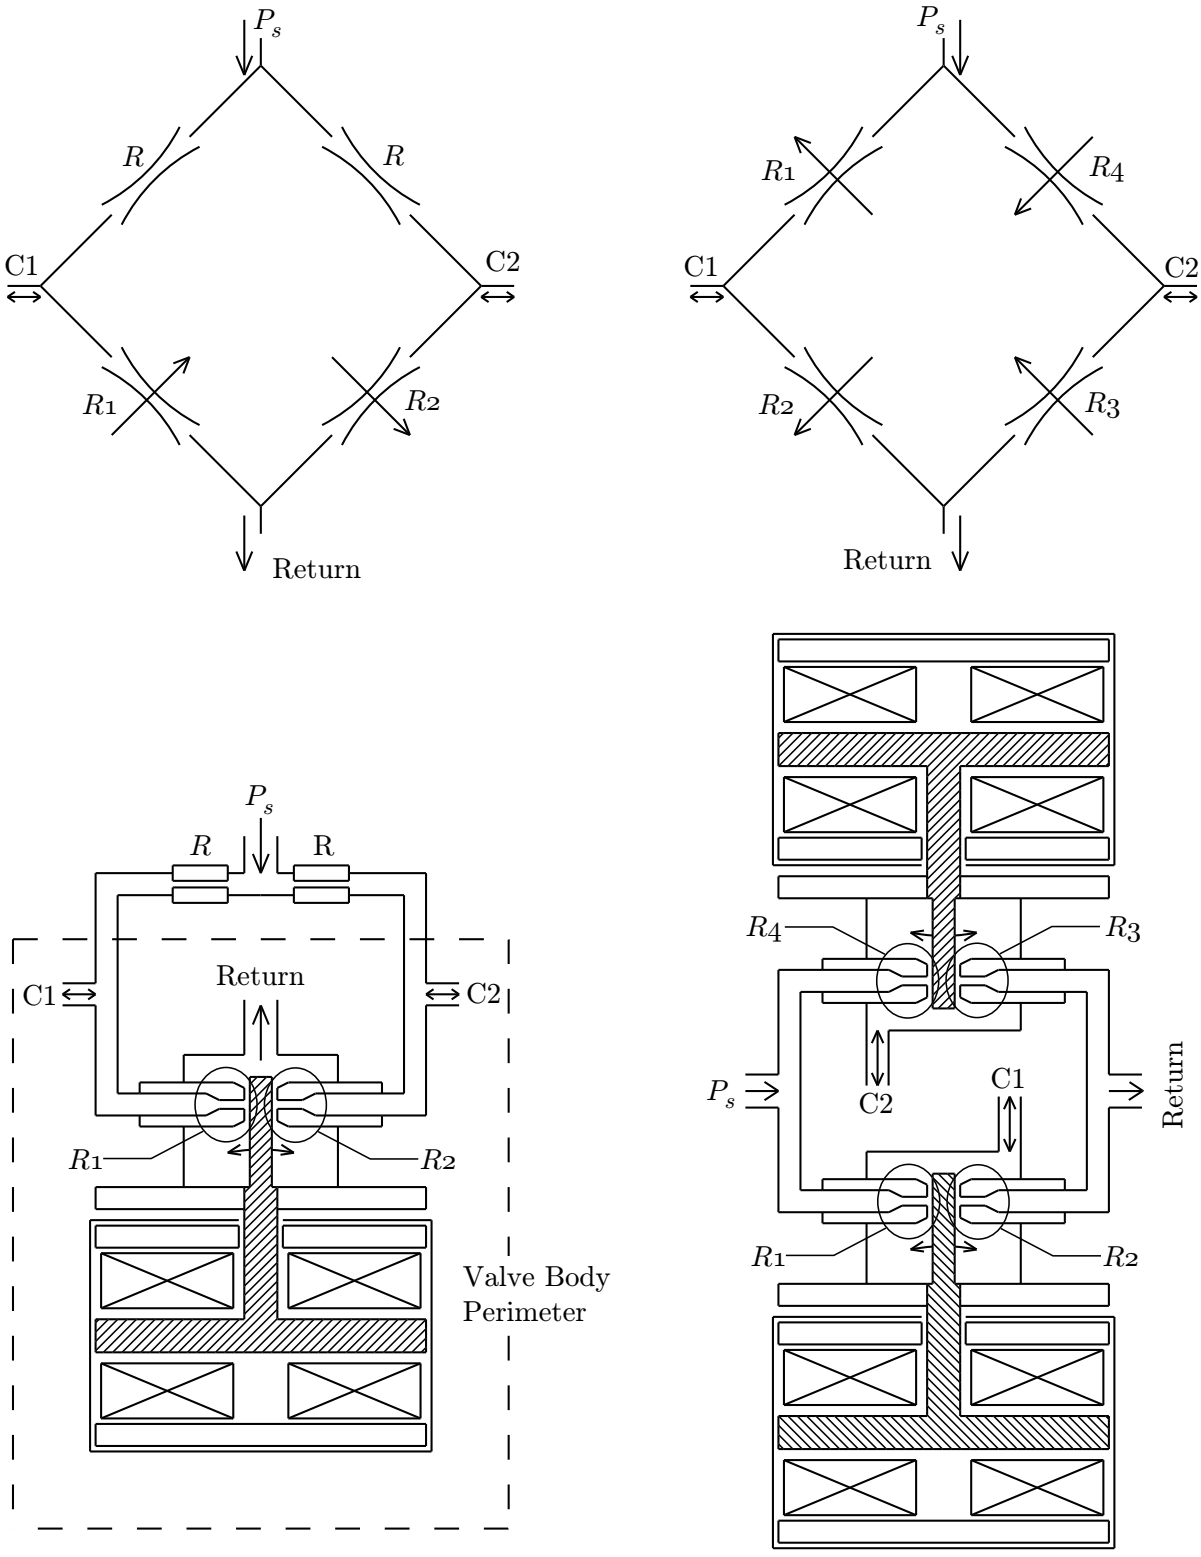

Figure 2.1: The HSC Valve and the Parker DYP2S Valve

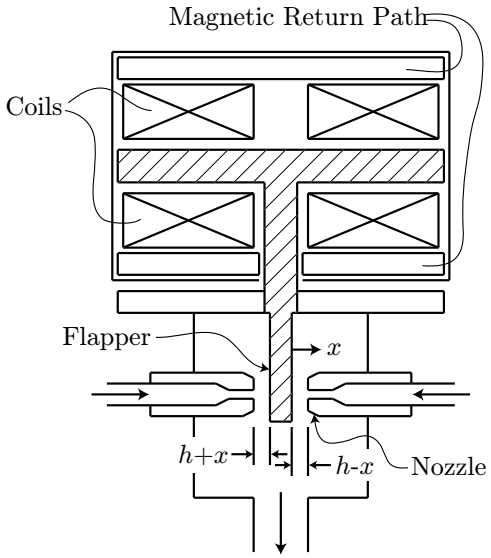

Figure 2.2: A Typical Flapper Valve Layout

the variable resistance produced by the height of the gap between the flapper and the nozzle face (h in figure 2.2). A typical nominal gap between the nozzle and the flapper is 0.2 mm.

There are several different flapper valve configurations. In this experiment two configurations were tested. The first actuator prototype used the single flapper HSC (Hydraulic Servo-Valve Corporation) manufactured valve. Since the HSC valve has only one flapper, the first actuator prototype includes a manifold with two fixed resistors to complete the bridge. Alternatively, later models of the actuator used the Parker manufactured DYP2S servo valve which has two separate flapper/torque motor pairs and four nozzles. With the DYP2S valve, the bridge is internal to the servo valve case. There are subtle differences in the application and operation between these two configurations, but both are capable of successfully controlling the actuator. The choice as to which valve to use was made mostly on the basis of commercial availability.

2.1 Hysteresis

The upper portion of the flapper is within the magnetic path of the torque motor (figure 2.2). The coils change the dipole orientation in the material they surround in order to create angular motion at the tip of the flapper (near the nozzles). As a result, the torque motor operates with a 2% hysteresis. This small hysteresis is suppressed by the feedback loops used to control the actuator. Hence, the valve hysteresis has not been observed to disrupt the linearity of the actuator. However, the actuator can be observed to relax to a new equilibrium position when the feedback loop is switched off, and this may partially attributed to this hysteresis.

2.2 Nozzle Design

The stock commercial valves are designed to operate in the turbulent regime. To this end, the stock nozzle design features a sharp edge intended to help trip turbulence. For this actuator, turbulent flow must be avoided and the obvious solution is to soften the edges of the nozzle. However, this must be done with some care so as to avoid creating areas of localized accelerated flow. Furthermore, with the reshaped nozzle, the nozzle face and flapper behave like a parallel plate impedance. This behavior is nonlinear in the laminar flow regime, but can be made more linear by adjusting the fixed resistance of the nozzle. The nozzle design for this application is then driven by two objectives: maintain laminar flow and minimize nonlinearity.

2.2.1 Nozzle Exit Shape

[19] A sketch of the stock nozzle design is shown Figure 2.3 along with two alternatives. The first alternative is the intuitive solution of simply replacing the sharp edge with a fillet and flattening the face of the nozzle. A more elaborate solution is to use a hyperbolic shape as shown on the far right of Figure 2.3. The hyperbolic shape is applied in an effort to maintain constant area and therefore constant flow velocity throughout the nozzle and flapper interface.

The hyperbolic shape is a consequence of the desire to minimize flow velocity

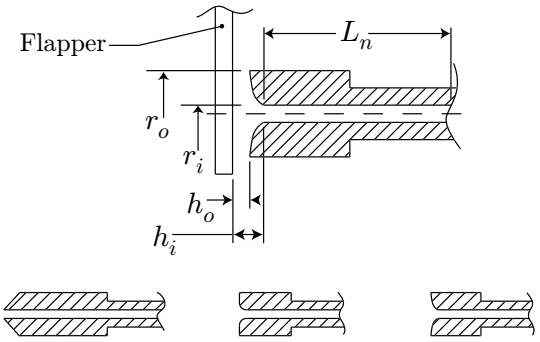

Figure 2.3: Nozzle dimensions and three possible nozzle shapes. From left to right: the stock nozzle, a flat faced nozzle with a fillet, and a nozzle with hyperbolic curvature.

fluctuations at the flapper/nozzle interface. This implies that the cross sectional area of the fluid path should remain constant as fluid travels down the circular nozzle bore and radially outward between the flapper and the nozzle face. Hence, the area of the cylinder defined by the radius of the outward flowing fluid (from r_i to r_o) and the distance between the flapper and the nozzle face (from h_i to h_o) must remain constant and equal to the cross-sectional area of the nozzle bore.

$$\pi r_i^2 = 2\pi r h(r) \tag{2.1}$$

$$h(r) = \frac{r_i^2}{2r} \qquad \text{for } r_i \leq r \leq r_o \tag{2.2}$$

At first, this may appear to be an ideal solution, but it is unreasonable to assume that fully developed laminar flow will prevail throughout this region. Fluid dynamics texts [27] suggest that unsteady (not fully developed) flow may continue where the ratio of entrance length to height is less than 1000. In the case of the flapper and nozzle, that ratio $((r_o - r_i)/h_{avg})$ is much closer to unity.

In consideration of the unstable flow in this region and the difficulty of accurately manufacturing a hyperbolic profile (at this miniature scale), the nozzle design features a simple fillet (the center case in figure 2.3). A valve with these nozzles was subsequently tested on a small flow bench, and in the operating range expected of the actuator, no significant unstable flow was observed. By virtue of the empirical

results of this test and many that followed, the nozzle design with a simple fillet was adopted for all future valves.

2.2.2 The Variable Nozzle/Flapper Resistance

There are two components to the resistance in the flapper valve: the fixed resistance from the bore inside the nozzle, R_n, and the variable resistance created by the flapper and nozzle face, R_f. In the laminar regime, the variable resistance exhibits an inverse cubic nonlinearity that is unacceptable at large flapper deflections. However, it is possible to minimize this nonlinearity by careful choice of nozzle impedance.

By the notation set forth in figure 2.3, the fixed resistance of the nozzle is:

$$R_n = R_{nozzle\,bore} = \frac{8\mu L_n}{\pi r_i^4} \tag{2.3}$$

where μ is the viscosity of the working fluid and L_n is the length of the nozzle bore. Alternatively, the nozzle/flapper interface can be approximated by a parallel plate impedance. The impedance of parallel plates to laminar flow is [27]:

$$R_{parallel\,plate} = \frac{12\mu \cdot Length}{Width \cdot Height^3} \tag{2.4}$$

where the $Height$ is the distance between the flapper and the nozzle face (h in figure 2.2), and the $Width$ and $Length$ can be loosely associated to the diameter of the nozzle face. A more accurate representation results from deriving the parallel plate resistance for radial flow constrained between two plates. The $Width$ becomes the circumference of the outward advancing flow perimeter, $2\pi r$, and the $Length$ is the distance along the radial path, r:

$$R_{radial\,plate} = \int_{r_i}^{r_o} \frac{12\mu}{2\pi r h^3} dr \tag{2.5}$$

$$R_{radial\,plate} = \frac{12\mu}{2\pi\ h^3} \ln\left(\frac{r_o}{r_i}\right) \tag{2.6}$$

The flapper resistance may be separated into two constituents: a normalized change

in height, $h/(h \pm x)$, and a nominal fixed resistance, R_o, for the impedance of the flapper centered condition:

$$R_f = R_o \left(\frac{h}{h \pm x} \right)^3 \tag{2.7}$$

$$R_o = \frac{12\mu}{2\pi\ h^3} \ln \left(\frac{r_o}{r_i} \right) \tag{2.8}$$

where h is the nominal height when the flapper is centered. The second term in equation 2.7 is further normalized with respect to the change in height:

$$R_f = R_o \left(\frac{1}{1 \pm \epsilon} \right)^3 \tag{2.9}$$

where $\epsilon = x/h$. The total resistance of the flapper/nozzle interface is the combination of the fixed resistance of the nozzle bore and the variable resistance of the flapper/nozzle interface:

$$R_{flapper/nozzle} = R_n + R_f \tag{2.10}$$

The purpose of deriving this model of the flapper/nozzle resistance is to mitigate the nonlinearity in R_f by specifying the nozzle resistance, R_n. This decision can be made empirically by plotting the differential pressure over a range of flapper positions while varying the ratio of the flapper/nozzle interface resistance, R_o, to the nozzle resistance, R_n. Based on inspection of figure 2.4, the most linear plot of differential pressure results from a nozzle resistance that is equal to the nominal flapper resistance, R_o (flapper centered). This can also be shown analytically by matching the slope of the differential pressure function at the flapper centered point and at some point of reasonable deflection. For the fully differential bridge in the DYP2S, the pressures at

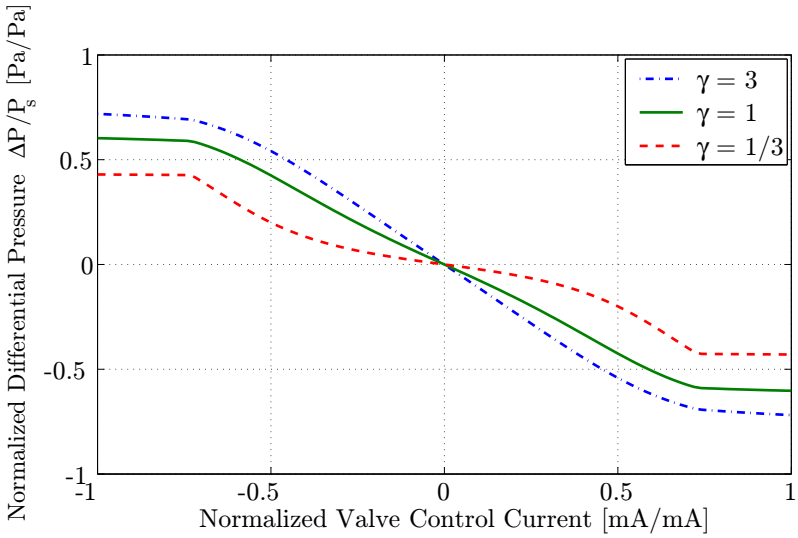

Figure 2.4: Normalized differential pressure (across port C1 and C2, figure 2.1) as a function of normalized valve control current. $\gamma = R_o/R_n$

the control ports are:

$$P_{c1} = \frac{R2}{R1+R2}P_s \tag{2.11}$$

$$P_{c2} = \frac{R3}{R3+R4}P_s \tag{2.12}$$

$$R1 = R_n + R_o\left(\frac{1}{1+\epsilon}\right)^3 \tag{2.13}$$

$$R2 = R_n + R_o\left(\frac{1}{1-\epsilon}\right)^3 \tag{2.14}$$

where P_{c1} and P_{c2} are the pressures at the control ports and P_s is the supply pressure. Due to the differential nature of the valve, $R1 = R3$ and $R2 = R4$. The differential pressure across the control ports:

$$\Delta P(\epsilon) = P_{c2} - P_{c1} = \frac{\gamma\left(\left(\frac{1}{1+\epsilon}\right)^3 - \left(\frac{1}{1-\epsilon}\right)^3\right)}{2+\gamma\left(\left(\frac{1}{1+\epsilon}\right)^3 + \left(\frac{1}{1-\epsilon}\right)^3\right)} \tag{2.15}$$

where $\gamma = R_o/R_n$.

Setting the slope of the differential pressure with the flapper centered equal to the slope when the flapper is displaced at $\epsilon = .7$:

$$\Delta P'(\epsilon) \mid_{\epsilon=0} = \Delta P'(\epsilon) \mid_{\epsilon=.7} \tag{2.16}$$

and solving for γ yields the expected $\gamma \cong 1$.

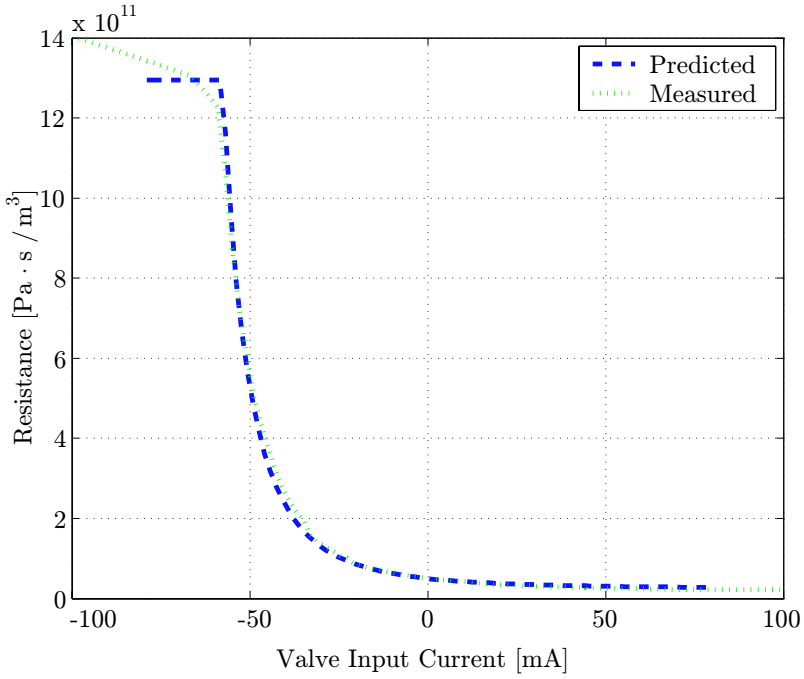

Figure 2.5: Measured and predicted resistance of one resistance in the DYP2S valve (e.g. one of the four resistances $R1$, $R2$, $R3$, and $R4$).

A flow bench was constructed to determine the resistance of valves with the new nozzles. The flow bench can also provide resistance measurements of each of the individual flapper nozzle pairs within the DYP2S valve (e.g. resistances $R1$, $R2$, $R3$, and $R4$). One such measurement is shown in figure 2.5 along with the predicted resistance from equations 2.14 or 2.14. The predicted behavior of the valve closely matches the measured resistance except for large negative inputs. This discrepancy is attributed to the inability of the actual valve to completely seal off the nozzle. As the flapper touches down on the nozzle face, the resistance increases sharply, but

the nozzle is not completely sealed. Higher values of control current continue to close off the gap increasing the resistance slowly. Alternatively, the model assumes a maximum possible resistance that cannot be surpassed with increasing (negative) control current, and as a result, the plot of the model levels off at large values of control current.

More information regarding the operation of the DYP2S valve can be found in Appendix A.

Chapter 3

Actuator Design Synthesis

The design of the hydraulic actuator was motivated by several one-dimensional mathematical models. In this chapter, these models will be developed and discussed with respect to the functionality of the actuator. These relationships eventually lead to a design trades chart that links significant design parameters to operational characteristics of the hydraulic actuator.

3.1 A Static Model

A first step in understanding the fundamentals of the actuator is to investigate a static model. The static model is useful for sizing the actuator and providing estimates of noise limitations imposed by the fluid supply pump and the servo valve. In this model, the hydraulic fluid is assumed incompressible and all structural compliance in the system is ignored. The equations here are derived with respect to Figure 3.1.

The derivation begins with a force balance between the actuator and the foundation:

$$(P_1 - P_2)A_b = K_{off}(z_f - z_a) \tag{3.1}$$

where A_b is the area of the bellows. In the previous chapter the differential pressure, $P_1 - P_2$, was defined as a function of the variable resistance created by the nozzle and

flapper valve. The variable resistance was defined as:

$$R_f = R_o \cdot \left(\frac{1}{1 \pm \epsilon}\right)^3 \tag{3.2}$$

where $\epsilon = x/h$. This can be approximated using a Taylor series including terms only up to first order as:

$$R_f = R_o \cdot (1 \mp 3\epsilon) \tag{3.3}$$

With this approximation, the equation of differential pressure is greatly simplified. For the HSC manufactured valve, differential pressure is:

$$\Delta P(\epsilon) = \frac{R_1}{R + R_1} - \frac{R_2}{R + R_2} = \frac{-6R_o\epsilon R}{(3R_o\epsilon + R + R_o)\,(3R_o\epsilon - R - R_o)} \tag{3.4}$$

where R is the fixed resistance built into the manifold (only for the HSC configuration). In equation 3.4, $3R_o\epsilon$ is much smaller than R or R_o and can be ignored. This can then be substituted into equation 3.1 to solve for the ratio of static force to normalized flapper position ϵ.

$$\frac{F}{\epsilon} = \underbrace{6A_bP_s}_{(1)} \underbrace{\frac{R \cdot R_o}{(R + R_o)^2}}_{(2)} \tag{3.5}$$

For any flapper position, it is advantageous to maximize the force, F. This can be done by increasing the area of the bellows or the supply pressure in part (1) of equation 3.5, but both of these choices have associated penalties. Alternatively, the ratio (2) can be maximized by choosing the appropriate value of R. This is the case when R is set equal to R_o (which can be shown by searching for the local maxima). This can also be understood as maximizing the *pressure recovery* or converting as much of the supply pressure into differential pressure. Setting $R/R_o = 1$ further simplifies equation 3.5 to:

$$\frac{F}{\epsilon} = \frac{3}{2}A_bP_s \tag{3.6}$$

The maximum allowable valve drive current noise can be estimated by setting $F = k_{off}\delta_{max}$ where δ_{max} is the maximum allowable displacement due to noise (about

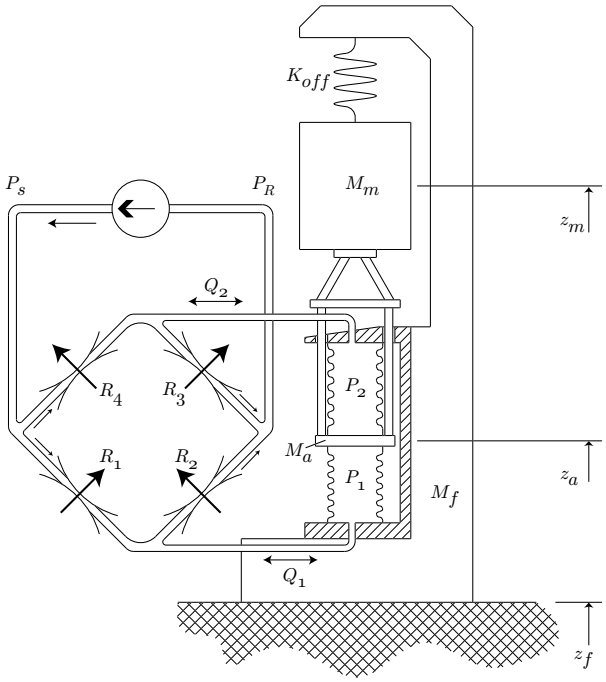

Figure 3.1: The Actuator Layout

1 nm at lower frequencies ∼1 Hz), and k_{off} is the offload spring stiffness. For an offload spring stiffness of $1{\cdot}10^6$ N/m and a bellows area of 0.1 m, the allowable drive current noise can be estimated as:

$$\delta\epsilon = \frac{2k_{off}\delta_{max}}{3A_bP_s} \cong 1 \cdot 10^{-6}\text{mA} \tag{3.7}$$

Similar approximations can be made for the Parker manufactured DYP2S valve. For the DYP2S valve, the differential pressure becomes, after the simplification of equation 3.3 (recall from chapter 1 that $R_2 = R_4$ and $R_1 = R_3$),

$$\Delta P(\epsilon) = P_1 - P_2 = \frac{R_2}{R_1 + R_2} - \frac{R_3}{R_4 + R_3} = \frac{R_1 - R_2}{R_1 + R_2} = 3\epsilon \tag{3.8}$$

The DYP2S is inherently optimized in that the upper two resistors, $R1$ and $R4$, are nominally matched to the lower two resistors, $R2$ and $R3$. It should be noted that due to the fully differential nature of the DYP2S valve, the pressure recovery is twice as

high as that yielded by the HSC valve (compare equation 3.8 to 3.4 with $R/R_o = 1$).

3.2 Fluid Compliance

The simple static model can be refined by accounting for the compressibility in the hydraulic fluid. To illustrate the effect of fluid compressibility, consider one-half of a conventional piston shown in figure 3.2. In this figure there are two ways to decrease the volume. The fluid may be drawn in through the orifice by moving the piston in the horizontal direction ($\dot{x}_p$); this is equivalent to adding fluid to the volume (Q_p). Alternatively, the fluid can be compressed by increasing the pressure P_p by the applied force, F_p.

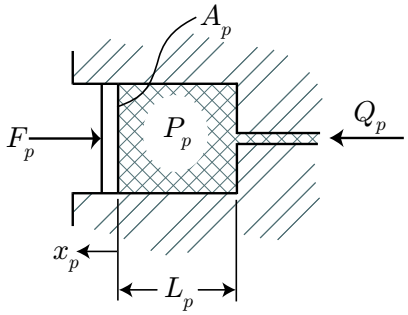

Figure 3.2: A conventional piston.

The mass balance for the fluid in the piston is dependent on the piston position and the flow out of the piston:

$$\sum M_p = \rho A_p L_p + \int \rho Q_p dt \tag{3.9}$$

where M_p is the fluid mass in the piston and A_p is the cross-sectional area of the piston. The change in mass is:

$$\frac{dM_p}{dt} = \underbrace{-A_p L_p \frac{d\rho}{dt}}_{(1)} + \underbrace{\rho A_p \frac{dL_p}{dt}}_{(2)} + \underbrace{\rho Q_p}_{(3)} \tag{3.10}$$

Term (3) in equation 3.10 represents the flow into the piston, term (2) is the flow generated by piston movement (changes in the length, L_p) and term (1) results from the compliance of the fluid under the pressure P_p. Terms (2) and (3) are equivalent since dL_p/dt is a velocity. In order to better understand the relationship between density and pressure, consider a fixed mass, M, of length L and area A:

$$M = \rho AL \tag{3.11}$$

$$\frac{dM}{dt} = 0 = AL\frac{d\rho}{dt} + \rho A\frac{dL}{dt} \tag{3.12}$$

$$AL\frac{d\rho}{dt} = -\rho A\frac{dL}{dt} \tag{3.13}$$

The change in length with time, dL/dt is related to the internal pressure of the piston by the compliance of the fluid. In this one dimensional case, fluid compliance is analogous to axial elastic deformation of a solid bar. The axial spring rate of a bar is EA/L where E is the modulus of elasticity, A is the cross-sectional area and L is the length. The spring rate is equal to the change in force divided by the change in length: $dF/dL = EA/L$. If both sides of this expression are divided by the cross-sectional area, A, the expression becomes $dP/dL = E/L$ where P is the pressure. If this is divided by dt and the modulus of elasticity is replaced by the bulk modulus, the result is $dL/dt = \beta/L \cdot dP/dt$ which can be applied to equation 3.13:

$$AL\frac{d\rho}{dt} = \frac{-\rho\beta A}{L}\frac{dP}{dt} \tag{3.14}$$

$$\frac{d\rho}{dP} = \frac{-\rho}{\beta} \tag{3.15}$$

Equation 3.15 expresses the relationship between density and pressure and may be substituted into equation 3.10:

$$\frac{dM_p}{dt} = \frac{\rho V_p}{\beta}\frac{dP}{dt} + \rho A_p\frac{dL_p}{dt} + \rho Q_p \tag{3.16}$$

where in term (1), the chain rule has been applied and $d\rho/dt$ is expanded to $d\rho/dP_p \cdot$

dP_p/dt. Finally, for constant volume, equation 3.16 becomes:

$$\frac{\dot{P}_p V_p}{\beta} + \dot{x_p} A_p + Q_p = 0 \tag{3.17}$$

where $dL_p/dt = dx_p/dt$.

3.3 A Dynamic Model of the Actuator

A dynamic model may be derived for the actuator that incorporates the compliance of the fluid. This derivation is based on the actuator with the DYP2S valve, but the results are very similar with the HSC valve configuration. The variables in the derivation correspond to figure 3.1. For this preliminary definition, the stiffness of the foundation K_f ,and the stiffness of the connection between the actuation plate M_a and the mass M_m are both assumed infinite. This indicates that the foundation moves with the ground ($z_f = z_g$), and the motion of the actuator plate is equal to that of the suspended mass ($z_a = z_m$).

In this model there are three inputs: the ground motion, z_g; the valve flapper position, ϵ; and disturbance forces, D, applied to the suspended mass. For this first derivation, the primary output of interest is the actuator plate position, z_a, but as the compliances for the actuator connection and foundation are added to the model, the list of states will grow to include the foundation and suspended mass position (equations 3.40 - 3.43).

The derivation begins with force balance:

$$(M_m + M_a)\ddot{z} = (P_1 - P_2)A_b + K_{off}(z_g - z_a) + D \tag{3.18}$$

Equations of volume (similar to section 3.2):

$$\frac{\dot{P}_1 V_b}{\beta} = Q_1 + A(\dot{z_a} - \dot{z_g}) \tag{3.19}$$

$$\frac{\dot{P}_2 V_b}{\beta} = Q_2 - A(\dot{z_a} - \dot{z_g}) \tag{3.20}$$

Nodal equations (P_R is the return pressure):

$$\frac{P_S - P_1}{R_1} - Q_1 - \frac{P_1 - P_R}{R_2} = 0 \tag{3.21}$$

$$\frac{P_S - P_2}{R_4} - Q_2 - \frac{P_2 - P_R}{R_3} = 0 \tag{3.22}$$

Bridge equations:

$$P_S = P_1 + P_2 \tag{3.23}$$

$$\Delta P = P_1 - P_2 \tag{3.24}$$

these may be combined to yield:

$$2P_1 = P_S - \Delta P \text{ and } 2P_2 = P_S + \Delta P \tag{3.25}$$

Assuming that the return pressure is close to zero (gage), P_R can be ignored and the summation of equations 3.19 and 3.21 can be used to eliminate Q_1:

$$\frac{\dot{P}_1 V_b}{\beta} = \frac{P_S - P_1}{R_1} - \frac{P_1}{R_2} + A(\dot{z}_a - \dot{z}_g) \tag{3.26}$$

For the next step, equation 3.25 can be used to replace P_1 with the more convenient ΔP. In addition, the resistors R_1 and R_2 are replaced with their respective linearized models from equation 3.3.

$$\frac{\dot{P}_S - \Delta\dot{P} V_b}{2\beta} = \frac{P_S + \Delta P}{2R_o(1 - 3\epsilon)} - \frac{P_S - \Delta P}{2R_o(1 + 3\epsilon)} + A(\dot{z}_a - \dot{z}_g) \tag{3.27}$$

The two fractions on the right side of equation 3.27 are nonlinear with respect to ϵ, but can be approximated with a Taylor series:

$$\frac{P_S + \Delta P}{2R_o(1 - 3\epsilon)} - \frac{P_S - \Delta P}{2R_o(1 + 3\epsilon)} \rightarrow \frac{\Delta P}{R_o} + \frac{P_S}{R_o}\epsilon + \left(\frac{P_S + \Delta P}{R_o} - \frac{P_S}{R_o}\right)\epsilon^2 \ldots \tag{3.28}$$

In the interests of linearity, only the first two terms of the expansion are maintained.

The derivative of the supply pressure, $\dot{P}_S$ in equation 3.27 may also be omitted because if the supply pressure is properly conditioned, variations in P_S should be small. The final equations of motion take the following form:

$$(M_m + M_a)\ddot{z_a} = \Delta P A_b + K_{off}(z_g - z_a) + D \tag{3.29}$$

$$\Delta\dot{P} = \frac{2\beta}{V_b}\left(A_b(\dot{z_g} - \dot{z_a}) - \frac{\Delta P}{R_o} - \frac{P_S\epsilon}{R_o}\right) \tag{3.30}$$

These two equations can be combined to form a single transfer function relating actuator displacement to valve position (assuming that $z_g = 0$).

$$\frac{Z_A}{\epsilon} = \frac{\dfrac{2\beta P_S A_b}{(M_m + M_a)V_b R_o}}{s^3 + \dfrac{2\beta}{V_b R_o}s^2 + \dfrac{2\beta A_b^2 R_o + K_{off}V_b R_o}{(M_m + M_a)V_b R_o}s + \dfrac{2K_{off}\beta}{(M_m + M_a)V_b R_o}} \tag{3.31}$$

In equation 3.31, Z_A is the Laplace Transform of z_a and s is the Laplace variable. Figure 3.3 illustrates the transfer function of equation 3.31 for typical values.

3.3.1 Pole Frequencies

The third order denominator of the transfer function can be broken into a single real pole and two complex poles. The frequency of the real pole represents the frequency where the force created by the offload spring is equal to the force of the differential pressure applied to the area of the bellows ($K_{off}z_a = \Delta P A_b$). This frequency is named the *channel frequency*. Below the channel frequency, the actuator behaves like a force actuator (a valve command is translated into a force), and above this frequency, a valve command translates to a velocity (evident from the first order negative slope at frequencies above the channel frequency).

There are several ways to determine the channel frequency. One method is to carry out the derivation of the dynamic actuator with equation 3.18 modified to:

$$\Delta P A_b = K_{off} z_a \tag{3.32}$$

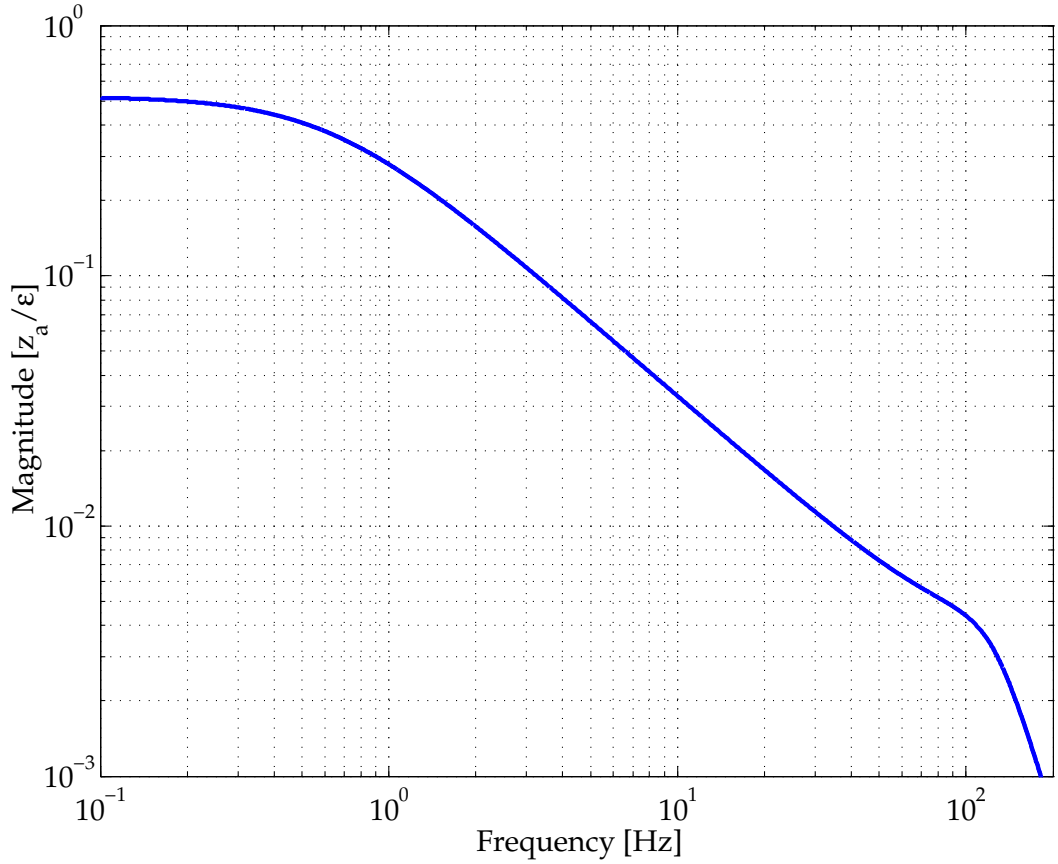

Figure 3.3: The transfer function of equation 3.31.

Due to the relatively low frequency of the channel frequency, all of the inertial terms are omitted. The disturbance forces and the ground motion are also absent in an effort expose the terms of interest. Solving for the transfer function between the actuator plate position, z_a, and valve command, ϵ, yields:

$$\frac{Z_A}{\epsilon} = \frac{constants}{s + \frac{2K_{off}\beta}{\underbrace{2\beta A_b^2 R_o}_{(1)} + \underbrace{K_{off}V_b R_o}_{(2)}}} \tag{3.33}$$

The frequency of the pole can be simplified by looking at the relative magnitude of the terms in the denominator (1) and (2). For typical values, $2\beta A_b^2 R_o \gg K_{off}V_b R_o$ and as a result, $K_{off}V_b R_o$ can be ignored. The resulting value of the channel frequency is:

$$\omega_{ch} = \frac{K_{off}}{A_b^2 R_o} \tag{3.34}$$

The same result may be obtained by examining the denominator of equation 3.31. At low frequencies, the magnitudes of the higher order (s^3 and s^2) coefficients are much smaller than those associated with the s and zeroth order terms. Removing the s^3 and s^2, and solving for the remaining pole frequency yields the same result.

It is possible to isolate the second order system from the denominator of equation 3.31 by carrying out polynomial long division with the known channel frequency pole. After some simplification, the actuator transfer function may be rewritten in

the canonical form.

$$\frac{Z_A}{\epsilon} = \frac{\Theta}{(s+\omega_{ch})\left(s^2+2\zeta\omega_n s+\omega_n^2\right)} \tag{3.35}$$

$$\Theta = \frac{2\beta P_S A_b}{(M_m+M_a)V_b R_o} \tag{3.36}$$

$$\zeta = \frac{1}{R_o A_b}\sqrt{\frac{(M_m+M_a)\beta}{2V_b}} \tag{3.37}$$

$$\omega_n = \sqrt{\frac{2\beta A_b^2+3K_{off}V_b}{(M_m+M_a)V_b}} \tag{3.38}$$

Equation 3.38 can, with only minor loss in precision, be simplified to:

$$\omega_n = \sqrt{\frac{2\beta A_b^2}{(M_m+M_a)V_b}} \tag{3.39}$$

3.3.2 Hydraulic Resonance

It is important to note that the natural frequency of the hydraulic actuator is weakly dependent on the offload spring stiffness and strongly a function of the stiffness of the fluid. When the offload spring is ignored, the resonant frequency is equivalent to the *hydraulic resonance*. Hydraulic resonance is the mass–spring mode associated with the fluid stiffness and the payload mass ($M_m + M_a$ in figure 3.1).

The spring stiffness for a hydraulic piston may be understood by analogy to the axial stiffness of a bar as was done in section 3.2. The axial stiffness of a bar is $K = EA/L = F/\Delta x$. If the numerator and denominator of this relationship are multiplied by area, the relationship is $K = \beta A^2/V = P/\Delta V$ where $\Delta V = A\Delta x$. Using this K in the typical expression for natural frequency, $\sqrt{K/M}$, yields equation 3.39 where $A = A_b$, $V = V_b$ and $M = M_m + M_a$ (there is a remaining factor of 2 since the actuator plate has a bellow on either side of it).

Equation 3.39 may be used to estimate the frequency of hydraulic resonance. Unfortunately, the result will be very optimistic. Two things compromise the stiffness of the fluid. The first is entrapped air in the fluid. This can be reduced by procedural

techniques. Bleed ports and running fluid through the system for a long period before first time operation can reduce the affect of entrapped air such that the bulk modulus may be as high as 1/3 the advertised value. To further reduce gas entrapment, Viersma [32] adds that at no point in the system should the fluid be allowed to fall freely in air as is sometimes seen at the return into the pump station reservoir. The second and more nefarious difficulty results from the *breathing stiffness* of the bellows and will be described in Chapter 5.

3.3.3 Foundation and Connection Stiffness

In the formulation of the dynamic model, it was assumed that the foundation of the actuator was rigidly attached to the ground and the actuator was rigidly attached to the payload. While both of these interfaces are quite stiff, they are not infinitely stiff, and therefore, must be accounted for in the model. In the simple one degree of freedom model, these terms can be introduced with a small modification of equation 3.29 and the inclusion of two more states.

$$M_m \ddot{z_m} = K_c(z_a - z_m) + K_{off}(z_g - z_m) + D \tag{3.40}$$

$$M_a \ddot{z_a} = \Delta P A_b + K_c(z_m - z_a) \tag{3.41}$$

$$M_f \ddot{z_f} = -\Delta P A_b + K_f(z_g - z_f) \tag{3.42}$$

$$\Delta \dot{P} = \frac{2\beta}{V_b}\left(A_b(\dot{z_f} - \dot{z_a}) - \frac{\Delta P}{R_o} - \frac{P_S \epsilon}{R_o}\right) \tag{3.43}$$

3.4 Design Trades Chart

The simple models derived in this section along with the requirements for the hydraulic external pre-isolation system can be combined to yield a design trades chart.

		Specifications		Design					Related Parameter	
		$\delta = 1mm$	$\Delta t = 10sec$	$P_s = 5bar$	$\beta = 2 \cdot 10^3 bar$	$R = 5 \cdot 10^{10} \frac{Pa \cdot sec}{m^3}$	$A_b = .01m^2$	$V_b = 3 \cdot 10^{-4} m^3$	$M_m = 1000kg$	$K_{off} = 4 \cdot 10^6 \frac{N}{m}$
1) Natural Frequency $\omega_n = \sqrt{\frac{2\beta A_b^2}{(M_m + M_a)V_b}}$	$\Uparrow$			$\Uparrow$	$\Uparrow$		$\Uparrow$	$\Downarrow$	$\Downarrow$	$\uparrow$
2) Damping $\zeta = \frac{1}{R_o A_b}\sqrt{\frac{(M_m + M_a)\beta}{2V_b}}$	$\Uparrow$			$\Uparrow$	$\Uparrow$	$\Downarrow$	$\Downarrow$	$\Downarrow$	$\Uparrow$	
3) Bridge Power Dissipation $P_\delta = \frac{P_s^2}{R}$	$\Downarrow$			$\Downarrow$		$\Uparrow$				
4) Acquisition Power $P_{acq} = \frac{K_{off}\delta^2}{\Delta t}$	$\Downarrow$	$\Downarrow$	$\Uparrow$							$\Downarrow$
5) Microseism Power $P_\mu = K_{off}\delta\delta_\mu\omega_\mu$	$\Downarrow$	$\Downarrow$								$\Downarrow$
6) Micorseism vs. Bridge $\frac{P_{acq}}{P_\delta} = \frac{K_{off}\delta\delta_\mu\omega_\mu R}{P_\delta^2}$	$\Downarrow$	$\Downarrow$		$\Uparrow$		$\Downarrow$				$\Downarrow$
7) Micorseism vs. Acquisition $\frac{P_\mu}{P_{acq}} = \frac{\delta_\mu\omega_\mu\Delta t}{\delta}$	$\Downarrow$	$\Uparrow$	$\Downarrow$							

Table 3.1: A design trades chart for the hydraulic actuator. A $\Uparrow$ indicates a strong influence while a $\uparrow$ represents a smaller $\sim 5\%$ affect.

Chapter 4

Actuator and Platform Design

The design of the quiet hydraulic actuator is optimized for ease of control. While control theory, and especially the computational power available have both improved dramatically in recent times, it is the opinion of the author and this research group that systems designed to simplify control will yield higher performance and be more robust than systems where the mechanical design is completed without such forethought.

Similarly, the challenge in designing the platform lies in minimizing cross-coupling and maximizing stiffness to facilitate control.

4.1 The Design Space

A design space may be either sharply confined with only a small range of solutions, or alternatively, broad in the choice of parameters enabling multiple successful solutions. The design of this actuator falls in the latter category. The principal parameters may deviate from those described here and the actuator can still meet the performance requirements set forth in Chapter 1.

Due to the flexibility in the design space, the parameters for this hydraulic system were chosen to be compatible with previous quiet hydraulics research where similar technology was successfully demonstrated on precision machine tools [19]. This was a considerable advantage that accelerated the development of the actuator and greatly

facilitated the commissioning of the prototype actuators and pumping station.

Much of the design of the actuator is dictated by the system it is destined for. The LIGO requirements set forth in chapter 1 specify a payload mass of 2000 kg, a range of motion of ±1 mm, a bandwidth greater than 20 Hz and a noise floor a factor of 10 less than the ground motion (about 10^{-9} m/$\sqrt{\text{Hz.}}$ from 1 to 10 Hz.). The parameters of the actuator are born from these requirements.

4.1.1 Natural Frequency Selection

The first step in the design of an isolation system is the choice of natural frequency. While this decision has implications for the entire system, it is of critical importance to the actuator design because the natural frequency dictates the force requirement by setting the stiffness of the offload springs.

The selection of natural frequency has many associated tradeoffs common to all suspension systems. However, due to the high stiffness and high force qualities of the hydraulic actuator, this system is somewhat unique by comparison to canonical soft or hard mount isolation platforms [20].

A soft suspension (lower natural frequency) is desirable because it affords more passive isolation at lower frequencies. However, the additional passive isolation comes at the cost of system stability. The static sag, z_{sag}, is inversely proportional to the square of the natural frequency: $z_{sag} = g/\omega_n^2$. Hence, lowering the natural frequency strongly affects the static sag (for example, a 1 hz suspension sags 250 mm while a 10 Hz suspension will sag only 2.5 mm). Moreover, since the modulus of elasticity is temperature dependent, the spring constant changes with temperature causing the static sag to change as well (appendix C.2). As a result, the position of the platform will change with temperature in direct proportion to the sag. This can be shown by manipulation of a few common relationships. In the following, δK_{off} and δz_{sag} are

the deviations from the nominal spring constant and sag.

$$z_{sag} = \frac{M_m g}{K_{off}} \tag{4.1}$$

$$z_{sag} + \delta z_{sag} = \frac{M_m g}{K_{off} + \delta K_{off}} \tag{4.2}$$

$$\delta z_{sag} = -\frac{M_m g \delta K_{off}}{K_{off}^2 + \delta K_{off} K_{off}} \tag{4.3}$$

If the cross term, $\delta K_{off} K_{off}$ is ignored,

$$\frac{\delta z_{sag}}{z_{sag}} = -\frac{\delta K_{off}}{K_{off}} \tag{4.4}$$

Thus, the change in sag is proportional to the change in spring constant.

In conclusion, the stability of the system overall, and particularly the stationarity of the horizontal modes (for example, the swing modes), is greatly improved with a stiffer system.

A stiff system is also advantageous for the design of the offload springs. Considering that the size of the spring is approximately proportional to the inverse square of the associated decreasing natural frequency, the change in size of the offload mechanism is significant as the natural frequency is made smaller. Lastly, a stiffer system greatly improves the elements of installation and initial alignment.

Naturally, there are some limitations incurred by adopting a stiff system. A stiff system will reduce the amount of passive isolation and require more force for a given displacement. However, the impedance of the actuator is large up to the frequency of the hydraulic resonance (ω_n in equation 3.38). Thus, for this system, the only benefit to setting the natural frequency significantly below the hydraulic resonance is the reduction of required force for displacement. In the case of hydraulics, this last concern is typically insignificant because the hydraulic actuator is inherently a high force device.

Based on these conclusions, a natural frequency of 8 Hz was selected for the offload suspension system in this project.

4.1.2 Actuator Sizing

Much of the physical size of the actuator is a dictated by the size of the bellows. The principal characteristics of the bellows, the diameter and length, are determined by a set of requirements that are not independent.

The force requirement for the actuator is set by the maximum displacement (1 mm) and the offload spring stiffness (set by the payload natural frequency of 8 Hz). Given these specifications and the differential pressure, it is simple to determine the required bellows area, and thus, the diameter of the bellows. However, the differential pressure is difficult to define because it is obscured by the choice of supply pressure and the amount of *pressure recovery* (the fraction of supply pressure converted into differential pressure) that is possible in the bridge network (the servo valve). Theoretically, if the servo valve was capable of completely closing off a nozzle (causing resistance to be infinite), the differential pressure could reach the supply pressure (for the DYP2S configuration figure 2.1). Practically this is not the case, and it is undesirable to operate the valve to its limits due to the nonlinear behavior associated with such large amplitude inputs (see figure 2.4). Experience with the servo valves has shown that it is reasonable to operate with a pressure recovery of about 30%. These properties of the valve set two requirements for the bellows: as an upper limit, the bellows must be able to survive exposure to the supply pressure, and for typical operation, the cross-sectional area of the bellows must be large enough to meet the force requirement with a differential pressure of $0.3P_s$.

Alternatively, another set of requirements arise from the length of the bellows. The axial stiffness of the bellows is directly proportional to the length of the bellows. It is desirable that the axial stiffness be small by comparison to the stiffness of the offload springs for two reasons. First, the stiffness of the bellows is in parallel to the offload springs, and therefore, a significant bellows stiffness will require equivalently more force for a given displacement. Next, if the stiffness of the bellows is comparable to the offload springs, any change of the bellows stiffness caused by a change in the fluid temperature could impose a time dependance of the fundamental system vibration modes. To avoid these adverse effects, the stiffness of the bellows is, somewhat arbitrarily, limited to $\frac{1}{10}$ of the offload spring stiffness.

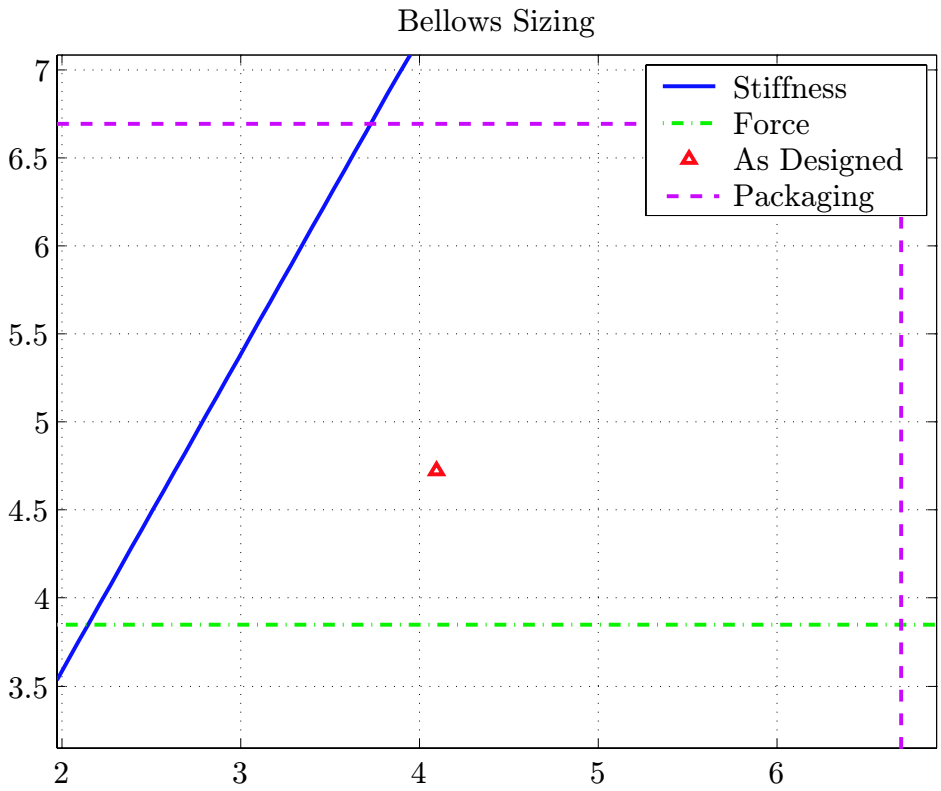

Figure 4.1: The design space for the bellows sizing. *Stiffness* refers to the requirement that the axial bellow stiffness be less than $\frac{1}{10}$ of the offload spring stiffness. *Force* is the minimum diameter required to create enough force for $1\,mm$ displacement (for $P_s = 0.7\,MPa$). *Packaging* provides an indication of the maximum reasonable length, and diameter for the actuator to fit comfortably with the ancillary equipment.

The considerations discussed above are shown graphically in Figure 4.1. Note that this figure is not a thorough description of the space. In the figure, both the supply pressure ($P_s = 0.7$ MPa) and the convolution geometry of the bellows remain fixed. Furthermore, Figure 4.1 does not address the effect on the performance variables defined in table 3.1. Nonetheless, it does show how the current bellows size (diameter of 120 mm and length of 104 mm fits within some of the physical constraints.

4.2 Quiet Hydraulic Actuator Design

The mechanical design of the quiet hydraulic actuator may be broken into three parts: the quiet hydraulic piston, the parallel motion flexures and the tripod flexure.

4.2.1 The Quiet Hydraulic Piston

The quiet hydraulic piston and cylinder is comprised of two bellows arranged around a plate (figure 4.2). This flexible bellows based design makes this configuration immune to the nonlinear, stick-slip behavior that is characteristic of the frictional interfaces found in conventional piston and cylinder arrangements.

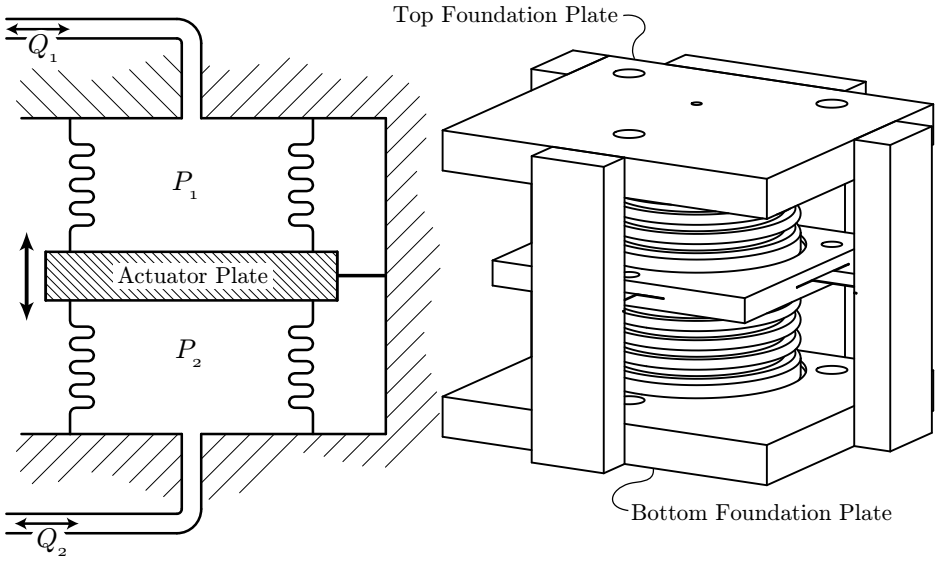

Figure 4.2: The quiet hydraulic piston. Flows Q_1 and Q_2 supply pressures P_1 and P_2 inside of the bellows. The piston is shown schematically on the left and as designed in the first version of the actuator on the right. Note the parallel motion flexures attached to the actuation plate.

4.2.2 Parallel Motion Flexures

The motion of the actuator plate or piston is maintained in the axial direction by two parallel motion flexures. The flexures must be soft in the axial direction to avoid

reducing the range of motion, and conversely, they must be large enough to resist buckling under lateral compressive loads.

In the axial direction, the flexures behave as beams with fixed - fixed end conditions. The flexures can be constructed from thin material allowing for a small bending stiffness. However, the thin section that enables the low bending (axial) stiffness also increases the likelihood of buckling.

There are two sources of compressive loading that can lead to buckling: the weight of the actuator plate when the actuator is installed horizontally and the horizontal force created when the tripod (figure 4.3) is displaced by a neighboring actuator. The tripod is designed to be soft in rotation, and based on this design (and results of quantitative models), the weight of the actuator plate is the predominant force. As this weight is not very large, it is straight-forward to develop flexures that meet both criteria.

The final design of the flexure is 0.5 mm (.020 in.) thick by 25.4 mm (1.0 in.) wide and 38.1 mm (1.5 in.) long. The axial stiffness is $2.41 \cdot 10^4$ N/m which is less than $\frac{1}{10}$ the offload spring stiffness, and the combined expected compressive loadings are approximately a factor of 10 less than the critical buckling load.

4.2.3 The Tripod Flexure

The quiet hydraulic piston as described above is complete and sufficient for single degree of freedom applications, but the applications considered here require that multiple actuators operate in concert. A multi-degree of freedom system dictates that actuators be arranged orthogonally to one another, and this further requires that the design of the actuator tolerate transverse displacements and tilts.

The solution employed here is to combine two rotational joints separated by a distance. The first rotational joint is the parallel motion flexures. The parallel motion flexures allow the actuator plate to be soft in pitch and roll, but stiff in yaw. The bellows accomodate these degrees of freedom. Near the connection to the payload, the tripod (figure 4.3) flexure is the other rotational joint. The tripod flexure is similar to a ball and socket joint in that it is soft in all rotational degrees of freedom,

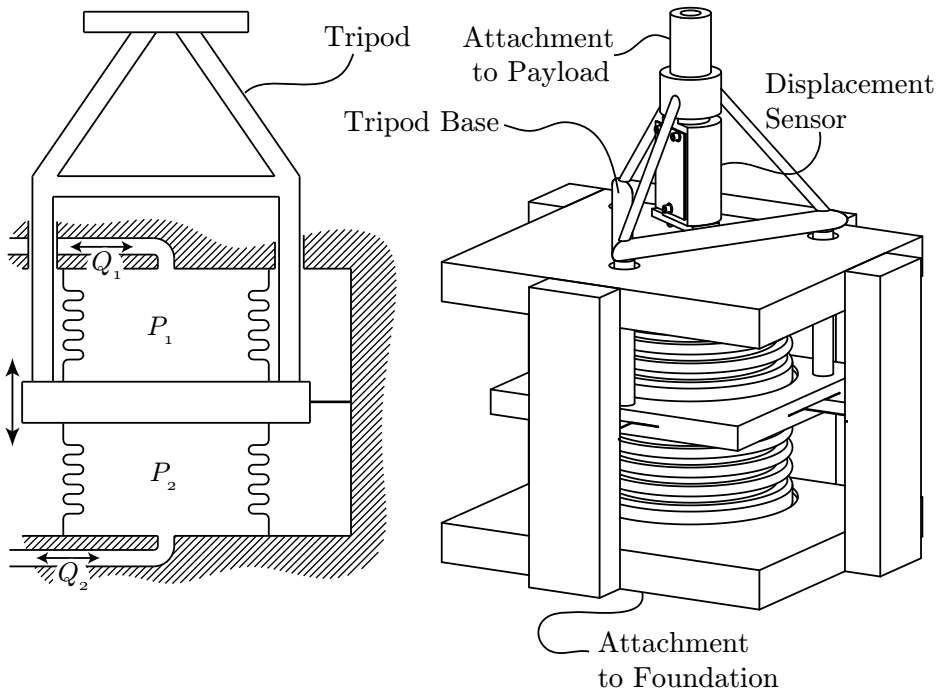

Figure 4.3: The quiet hydraulic piston with tripod flexure. The tripod flexure is added to the actuator to accommodate transverse motions. This enables multiple actuators to operate in concert.

but stiff in the three translations. Transverse motion is made possible because of the coordinated flexibility of both the parallel motion flexures and the tripod (figure 4.4).

4.2.4 First Actuator Prototype Results

This design process culminates in the first actuator prototype (figure 4.3). A horizontal and vertical actuator pair were tested extensively in the Quiet Hydraulic Test Platform (section 4.3.2). The initial open-loop response of the actuator is shown in figure 4.5. This differs notably from the predicted response using the models presented in chapter 3. The obvious difference between the predicted behavior and the measured behavior is the lightly damped resonance at ∼27 Hz. While the model can be made to capture these dynamics by dramatically reducing the fluid bulk modulus, the cause of the decreased resonant frequency is not a shortcoming of the fluid.

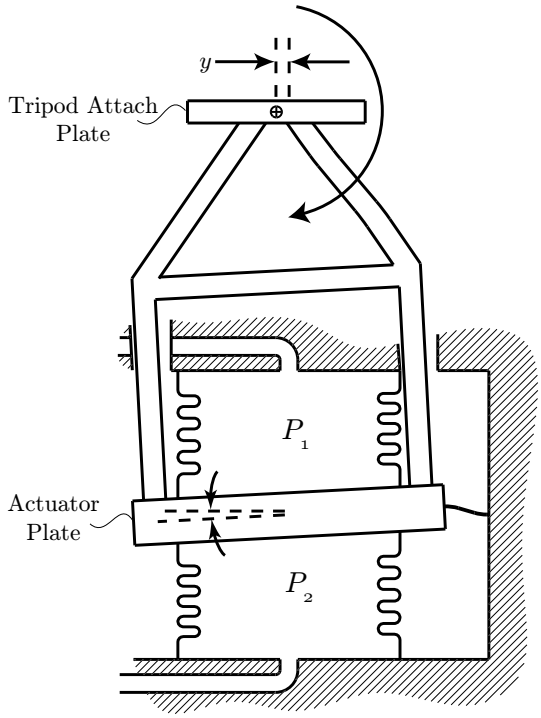

Figure 4.4: Transverse motion compliance with the tripod flexure. In the event of a transverse motion, y, the tripod flexure rotates clockwise (in this example) and the actuator plate tilts counterclockwise. The combination of these two rotations and the separation between them allow for the transverse displacement.

Instead, the reduced stiffness results from an unexpected compliance of the bellows. This *breathing* compliance observed in the bellows is discussed in detail in chapter 5.

4.2.5 The Quiet Hydraulic Actuator Bolted Prototype

The manufacture of a bolted prototype (figure 4.6) followed the initial welded prototype to address the need for flexibility in the actuator configuration. The effort to reduce the bellows breathing resonance necessitated testing different bellow geometries and modifications to components internal to the bellows (namely the bypass network described in chapter 5). The bolted prototype offered the opportunity to completely disassemble the actuator and modify or replace any of the critical components.

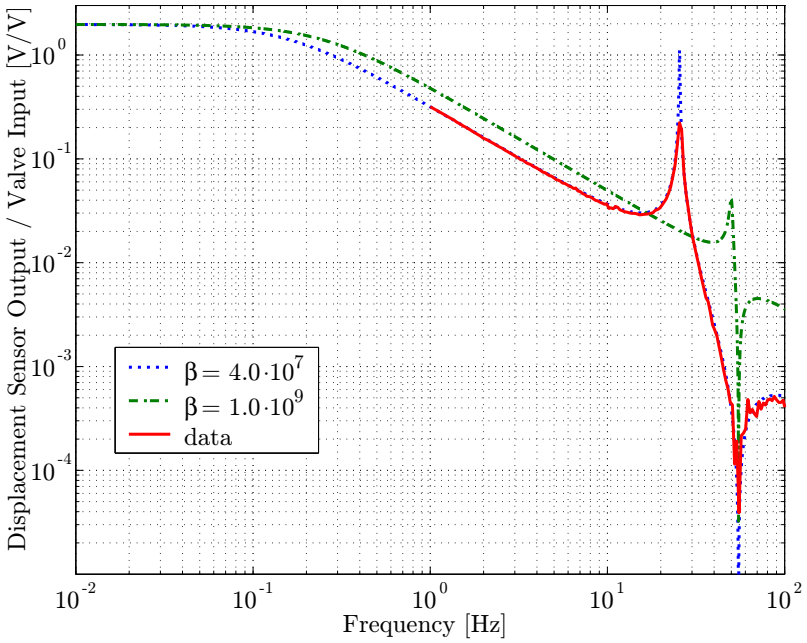

Figure 4.5: An open loop transfer function from vertical valve drive to displacement sensor output. The two traces labeled with bulk modulus values, β [Pa], are generated with the model whereas the data is measured from the hydraulic test platform. Note that the model matches the data only when the bulk modulus is reduced by nearly two orders of magnitude from the expected value.

In addition to the modifications made to suppress the breathing resonance, the bolted prototype also features modified tripod legs and bellows slugs. The tripod legs were modified to increase the axial stiffness of the connection between the actuator and the payload (believed to be partially responsible for the first zero in figure 4.5). The small notches shown on the tripod in figure 4.6 greatly decrease the rotational stiffness of the legs allowing for the nominal diameter to be larger (0.375 in. as opposed to 0.25 in. in the original prototype). This larger nominal diameter provides a much greater axial stiffness since the stiffness along the bar increases as the square of the radius.

The bellows slugs are placed within the bellows to remove excess fluid that might contribute to reducing the hydraulic resonant frequency. While the stiffness of the fluid is not expected to dominate the hydraulic resonance (even with stiffer bellows),

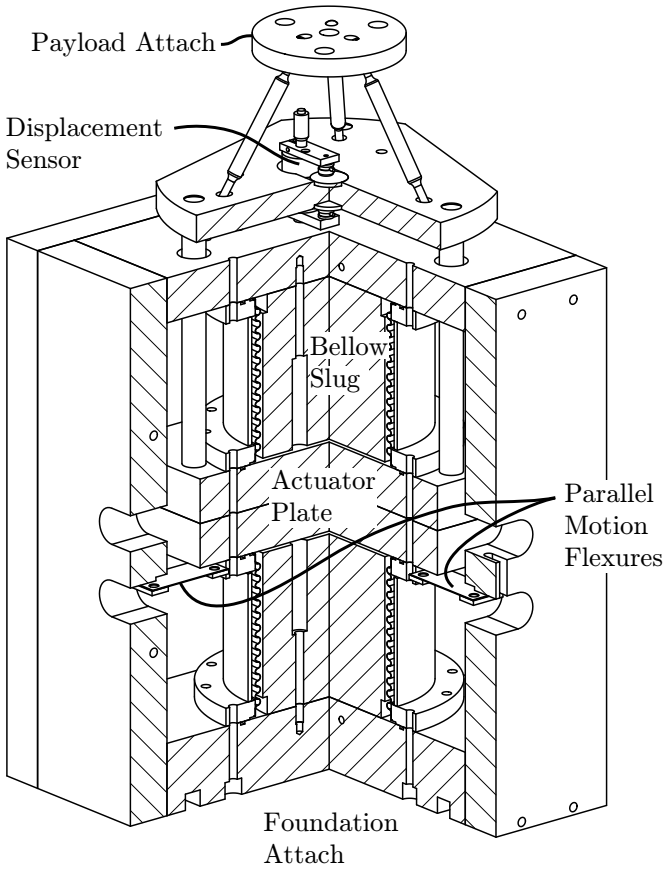

Figure 4.6: A section view of the bolted actuator prototype.

removing fluid from the actuator has potential advantages for weight savings and opening up additional space inside of the actuator.

4.2.6 The Final Quiet Hydraulic Actuator Design

The final actuator design incorporates all of the developments of the previous two prototypes into a tightly integrated package. The most significant addition is the bypass network integrated into the center of the actuation plate, but other features are introduced into this design to make the actuator more suitable for installation into the LIGO clean-room environment.

The concern for clean-room compatibility led to an all welded design that minimizes the number of o-ring seals. While the experience in the Stanford labs has been that o-ring seals are very reliable, the LIGO Lab community is more comfortable with welded connections, and hence, the final design is monolithic at the cost of the serviceability apparent in the bolted prototype.

In all of the previous prototypes, entrapped air is removed or bled from the bellows by opening up the pressurized bellows to the surrounding environment. Clearly, this messy procedure would not be acceptable in the LIGO installation, and in order to circumvent this, pin valves were integrated into the structure of the actuator.

There are five small and one large pin valve on the actuator. Each of the small pin valves allow cavities in the actuator to be open to the return pipe. The small pin valves are oriented such that the actuator may be bled in either the horizontal or vertical orientation. The large pin valve controls the return from the valve. When the large pin valve is closed (tightened), the fluid is forced to circulate into the bellows and out through the five return ports opened by the small pin valves. This configuration allows for fluid to be circulated through the actuator for long periods of time in order to ensure that the entrapped air volume is minimized. After several days of circulation, the large pin valve is opened and the small pin valves are closed. Following this, the actuator is ready for operation.

Instrumentation in the Final Actuator Design

Embedded in the final design are mounts for both a Kaman manufactured DIT-5200 displacement sensor and a Marks Products L4C seismometer. The displacement sensor is mounted on a linear bearing with micrometer control. This facilitates the calibration of the displacement sensor and enables the sensor to be centered after the actuator has been mounted.

The feedback seismometer is intended to be placed inside of the upper bellow and attached to the tripod plate. This arrangement provides a high degree of collocation, but is only permissable in the vertical direction. The seismometer cannot be mounted to the tripod plate in the horizontal orientation because of the tilt experienced by the tripod plate during transverse displacements. As is implied in figure 4.4, the tripod

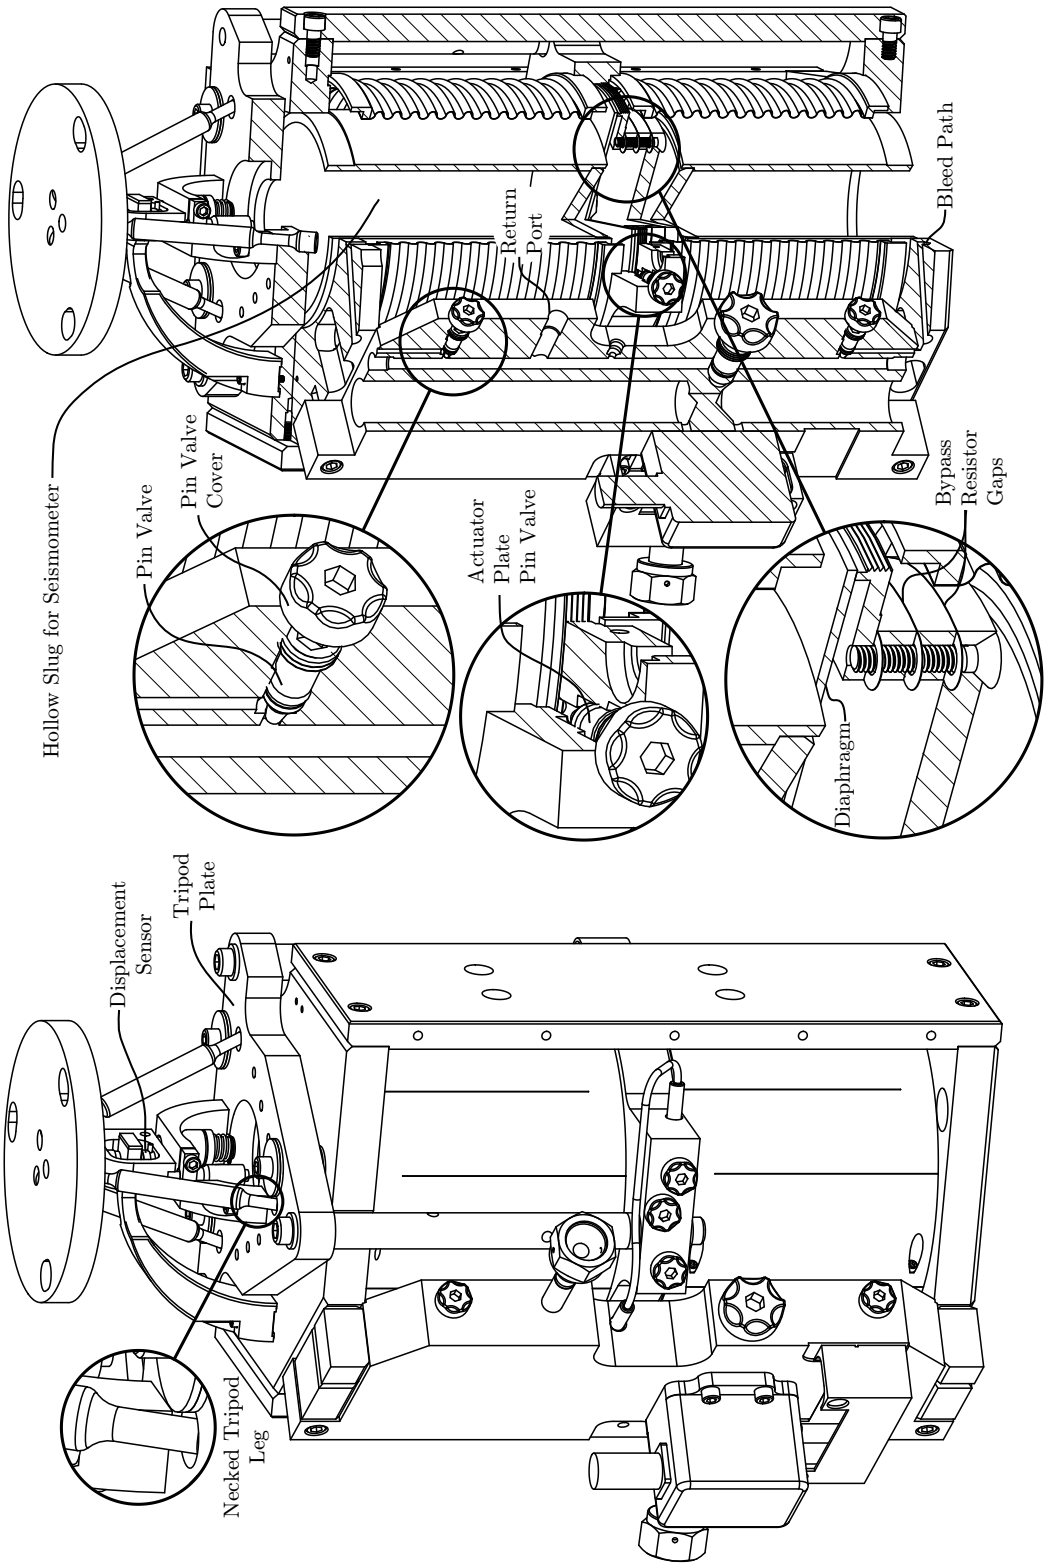

Figure 4.7: The final actuator design. The bypass network is discussed in chapter 5. Note that in this final design, the bellows slugs are hollow, and in the vertical configuration, the L4C seismometer fits inside of the top slug.

plate tilt can approach $1/10^{\text{th}}$ of a degree which is far beyond the acceptable limit for a manageable frequency of the tilt-horizontal coupling zero.

4.3 The Platform

In order for the quiet hydraulic actuator to be used in LIGO, it must be integrated into the structure that supports the LIGO optics. This presents several challenges as the LIGO support structure was not originally intended for use with the quiet hydraulic actuator.

To gain experience with the scale and configurations encountered in LIGO, a quiet hydraulic test stand (figure 4.10) was constructed with two hydraulic actuators. Once performance was verified on the test platform, the suspension and support structure were developed to incorporate the hydraulic actuator around the LIGO BSC chamber. The actuators and supporting structure is often referred to as the External Pre-Isolator (EPI) because it is installed external to the vacuum system. When installed with the hydraulic actuator, the system is called the Hydraulic External Pre-Isolator (HEPI).

4.3.1 Design for Control in Active Platforms

The design of the hydraulic test stand and the EPI incorporate many features intended to make the systems amenable for feedback control. Paramount among these is the practice of collocating sensors and actuators. However, other advantages may be accrued by careful positioning of the actuator and spring attachments, and attention to stiffness in both the suspended elements and the foundation.

Collocation

In the context of feedback control, the metric of collocation between sensors and actuators is expressed as the stiffness of the connection between the sensor and actuator. A sensor and actuator are described as well-collocated if the member which connects them is rigid (a solid bar of metal, for example). However, at some frequency this

connection will begin to exhibit compliance and the sensor will no longer be perfectly in phase with the actuator. Above the frequency where this occurs, it is difficult to maintain stable feedback control as discussed in appendix C.3. Since compliance becomes prominent at higher frequencies, it is most important that the high-frequency feedback sensor (typically a seismometer) be well collocated to the actuator.

Interfacing with the Suspended Mass

The goal in choosing the positions of attachment for the springs and actuators is to minimize cross-coupling between degrees of freedom. For feedback control, coupling between horizontal force to tilt is the most crippling because the horizontal seismometer cannot distinguish between these two motions at low-frequency (appendix D.2).

Tilt-horizontal coupling, as it is commonly labeled, may result from two types of misalignments. If the horizontal actuator is displaced vertically (a distance h_f in figure 4.8) from the center of mass, force applied by the actuator will create a torque, and therefor, a tilt. However, at low-frequencies this coupling is inconsequential by comparison to any misalignment between the horizontal actuator and the *center of stiffness*.

The center of stiffness is the point where the horizontal and vertical spring forces intersect (labeled *cs* in figure 4.8). For typical values of spring stiffness and spring spacing (r_k), the spring forces dominate the dynamics of the suspended mass at low-frequency. As a result, if the actuation force is misaligned (a distance h_k) from the center of stiffness, coupling from horizontal actuation to tilt will be prevalent at low-frequencies.

Hence, at high-frequencies, the misalignment h_f from the center of mass will dominate the coupling to tilt from horizontal actuation, and at low-frequencies, tilt coupling will be dominated by the misalignment h_k from the center of stiffness. The equations of motion for the system shown in figure 4.9 are (in state space form):

$$\begin{pmatrix} \ddot{x} \\ \ddot{\theta} \end{pmatrix} = \begin{pmatrix} \frac{-k_h}{m} & \frac{k_h h_k}{m} \\ \frac{k_h h_k}{I} & \frac{-k_h h_k^2 - k_v r_k^2}{I} \end{pmatrix} \begin{pmatrix} x \\ \theta \end{pmatrix} + \begin{pmatrix} \frac{1}{m} & 0 \\ 0 & \frac{1}{I} \end{pmatrix} \begin{pmatrix} 1 \\ h_f \end{pmatrix} F_h \qquad (4.5)$$

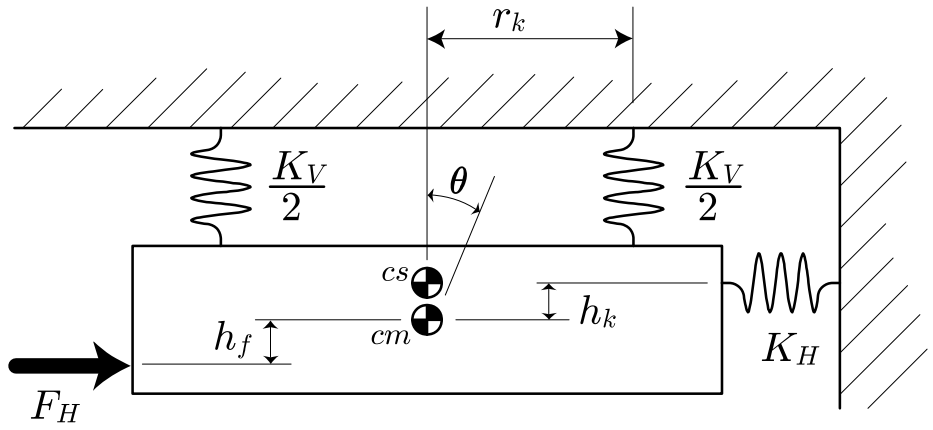

Figure 4.8: A two degree of freedom suspension. In this simple model, it is assumed that the springs possess only axial stiffness (as opposed to a realistic spring which exhibits stiffness in the other five degrees of freedom as well). A vertical and horizontal actuator apply forces F_V and F_H respectively. The center of mass is labeled *cm* and the center of stiffness is indicated by *cs*.

Numerical values for each of the constants in equation 4.5 can be derived from the geometry of the hydraulic test platform. Based on this, the values for some of the constants for the test stand are: $m = 150$ kg, $r_k = 0.3$ m and $I = 13.2$ kg/m^2. The spring constants are inferred by setting the vertical natural frequency to 10 Hz and the horizontal natural frequency to 2 Hz (these values are only approximations to the actual platform). On the hydraulic test platform, the actuator is well aligned with the center of mass, but poorly aligned with the center of stiffness.

The difficulty in aligning with the center of stiffness arises because the position of a spring's effective pivot point can be difficult to predict. A coil spring confounds this problem by typically exhibiting a pivot point that changes with orientation. In the case of a coil spring, the effective pivot point location from off-axis motions will differ when the spring is rotated 90 degrees. In other words, the transverse stiffness of a coil spring is orientation specific simply by virtue of the asymmetric shape of the coils and end shape of the spring. Alternative spring designs such as a wire hanging from a cantilever blade (the length of the wire dictates the horizontal frequency while

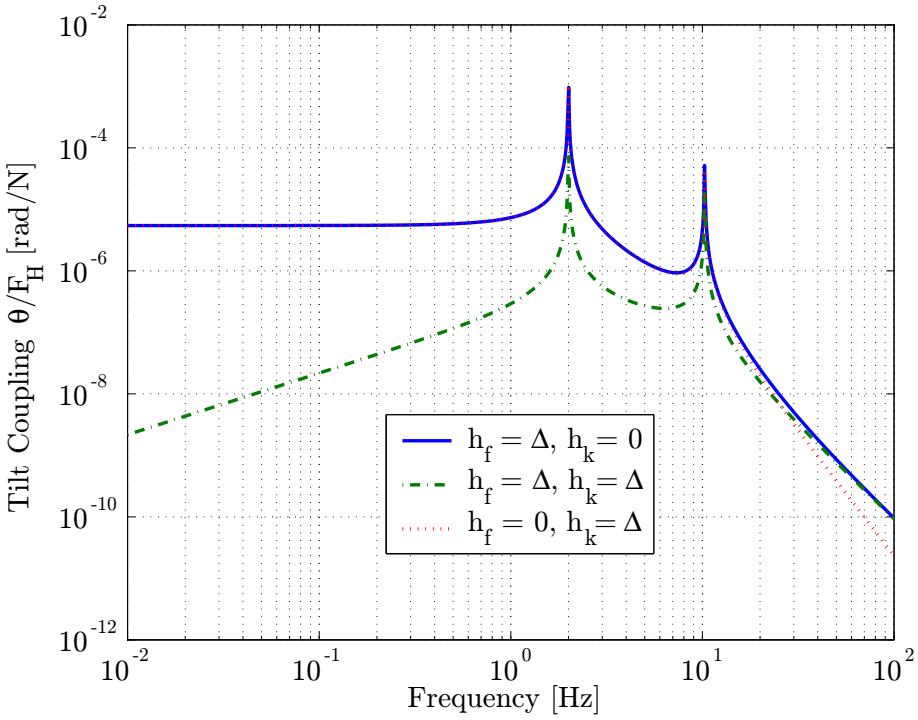

Figure 4.9: A driven transfer function of tilt caused by a horizontal force. This figure is generated based on a model of the simple platform shown in figure 4.8. Here, the tilt generated by horizontal force is minimized if the force is applied to the center of stiffness at low-frequencies ($h_f = h_k$), and at high-frequencies, it is preferable to align the actuator with the center of mass ($h_f = 0$). The ideal case is to align the horizontal stiffness and the actuator with the center of mass ($h_f = h_k = 0$) which reduces the tilt coupling to zero.

the blade stiffness sets the vertical frequency) can be much easier to predict.

In order to investigate the effects of misalignment in the context of the simple model of figure 4.8, an offset, Δ= 30 mm, is introduced. The coupling to tilt from horizontal force is illustrated in figure 4.9 for the cases of the actuator misaligned from the center of mass ($h_f = \Delta$ and $h_k = 0$), the center of stiffness ($h_f = 0$ and $h_k = \Delta$) or both ($h_f = h_k = \Delta$). Note that if there is no misalignment ($h_f = h_k = 0$) then there is no coupling to tilt.

It is unlikely that the optimum, no tilt, case will be observed. Not only because of the difficulty of predicting the pivot point of the springs, but more significantly

because of the qualities of the sensors. At low-frequencies the sensitivity to tilt of seismometers exceeds the available precision for aligning the actuator to the center of stiffness. Indeed virtually any tilt in the platform will be amplified and displayed prominently in the output of the seismometer as the tilt-horizontal coupling zero (appendix D.2).

At high-frequencies, the situation is reversed because the seismometer is operating past its internal natural frequency, and therefore, resists tilt. As implied in figure 4.9, and due to the dynamics of the seismometer, it is reasonable to assume that with standard shop practices all output at high-frequency will be due to horizontal motion.

Alternative means for reducing tilt are to increase the moment of inertia or the suspended mass or the separation of the vertical springs (r_k). While neither of these techniques have the opportunity to completely abolish tilt, the coupling can be reduced by these practices.

Design for Stiffness

Design for stiffness as opposed to strength is the hallmark of mechanical design for control. A design optimized for stiffness typically far exceeds the requirements for strength, and most importantly for control, exhibits high resonant frequencies.

High stiffness is what makes collocation possible, but is also a concern for both the suspended mass and the structure that supports the actuator. Compliance in the suspended mass results in resonant modes that populate the response of the system and complicate controller design. Usually, it is possible to damp these modes providing that they do not interfere with the collocation between the sensor and actuator. Nonetheless, control design in the absence of flexible modes is much easier.

Flexibility in the foundation that supports the actuator can also lead to detrimental effects. From the perspective of feedback control, the most inconvenient of these occurs when the suspended mass and the foundation move in unison causing the displacement sensor output to be zero. This can be understood by imagining the foundation and the suspended mass as discrete rigid bodies connected by springs. The foundation mass is connected to the ground by the foundation spring and the

suspended mass is connected to the foundation by the offload spring. The actuator acts between the suspended mass and the foundation, and the displacement sensor measures the distance between the two masses. At a particular frequency ($\omega = \sqrt{K_f/(M_m + M_f)}$) the two masses move together and the displacement sensor outputs zero.

For these reasons, control design is simplified by stiff structures with flexible modes far outside the desired bandwidth of control.

4.3.2 The Quiet Hydraulic Test Platform

The quiet hydraulic test platform, Figure 4.10, is designed to mimic one corner of the eight actuator installation in LIGO. To this end, the suspended mass is approximately one-quarter the mass of the LIGO installation, and the suspension frequency is tuned to be the same as intended for LIGO.

Quiet Hydraulic Test Platform Instrumentation

The test platform is instrumented with sensors for the purposes of feedback and feedforward control. Since the LIGO requirements specify both alignment and isolation performance, there are both displacement and inertial sensors connected to the suspended mass for feedback control. The primary feedback sensor, a displacement sensor, is integrated into the tripod to be well collocated with the actuation plate. The feedback inertial sensor, an HS-1 manufactured by Geospace Corporation, is mounted near the displacement sensor. For sensor correction (section 6.2.2) control, a Streckheisen corporation manufactured STS2 seismometer is located on the ground near the platform. The STS2 measures accelerations in x, y and z simultaneously down to frequencies of 8 mHz. Embedded within the suspended mass are two S13 seismometers (manufactured by Teledyne / Geotech) with a 1 Hz natural frequency. These seismometers are used as witness sensors to measure the isolation performance of the control system external to the feedback Loop.

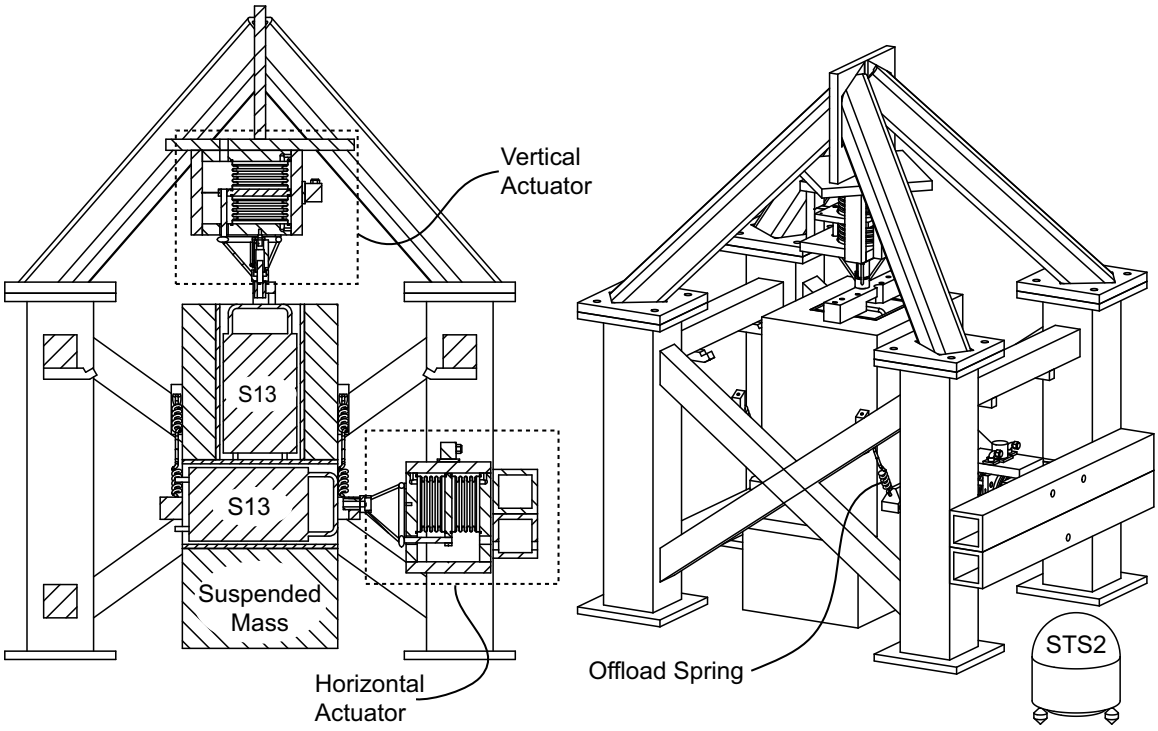

Figure 4.10: The quiet hydraulic test platform. Two actuators are arranged around a suspended mass of similar weight to one-quarter of the working mass in LIGO. In addition to the displacement sensors resident to the actuators, there are two S13 witness seismometers nested in the mass and a STS2 low frequency seismometer on the ground. The feedback seismometers are not shown here.

Design of the Test Platform

The hydraulic test platform embodies many of the ideas discussed in section 4.3.1. However, the design is not free from imperfections which serve as good examples to develop an even greater appreciation for the ideas outlined earlier.

The connection between the vertical actuator and the mass is an example of a stiffness driven design. The actuator is bolted (with a 1.0 inch diameter bolt) to a 2.0 inch x 2.0 inch aluminum beam (figure 4.10). This beam is also bolted to the mass, but it must span the gap created by the housing of the vertical S13 seismometer. The bending stress experienced by the beam when the actuator is in operation is at most $1/10^{\text{th}}$ the yield stress of the material. However, in spite of the seemingly excessive

design, the beam is compliant enough to have negative implications on performance. Figure 4.5 shows a driven transfer function of the displacement sensor output from the first actuator prototype in the test platform. The compliance of the connection bar is the cause of the zero at 55 Hz. While the internal resonances of the connection beam are much higher, the bending stiffness of the beam (pinned at the connection points) is predicted numerically to be $4.0{\cdot}10^7$ N/m which in combination with a 300 kg payload explains the mode at 55 Hz. Hence, even structures that appear to be excessively overbuilt can contribute to low-frequency resonances especially wherever material is used in bending.

The pyramid structure at the top of the platform supports the vertical actuator and spans the distance from the corners (1 m corner to corner). By virtue of the diagonal (as opposed to purely horizontal) load lines and large tubular structure, this aspect of the foundation has a stiffness predicted by finite element methods to be $1.0{\cdot}10^{10}$ N/m. The first modes believed to result from from the foundation are above 200 Hz.

The most significant deviation from the guidelines of section 4.3.1 occurs in the attachment of the actuators to the payload. The axis of the vertical actuator naturally traverses the center of mass, and comes acceptably close to the center of stiffness. Alternatively, while the axis of the horizontal actuator is well aligned with the center of mass, it is well separated from the center of stiffness. The cause of this is that the location of the center of stiffness is, in this case, dominated by the transverse stiffness of the vertical actuator as opposed to merely the coil springs. Since the horizontal actuator is separated from the effective flex point of the vertical actuator by approximately 700 mm, tilt is readily observed from low-frequency horizontal actuation. The horizontal S13 witness seismometer is sensitive to tilt, and tilt motion is evident in the transfer function of S13 output (figure 4.11).

The tilt introduced by this oversight can be suppressed by adding additional horizontal springs. However, before this can be done, the torsional stiffness of the platform and the position of the center of stiffness must be determined. This is done by monitoring changes in the frequency of the tilt-horizontal coupling zero while adding horizontal stiffness to the mass. At the frequency of the tilt-horizontal coupling zero,

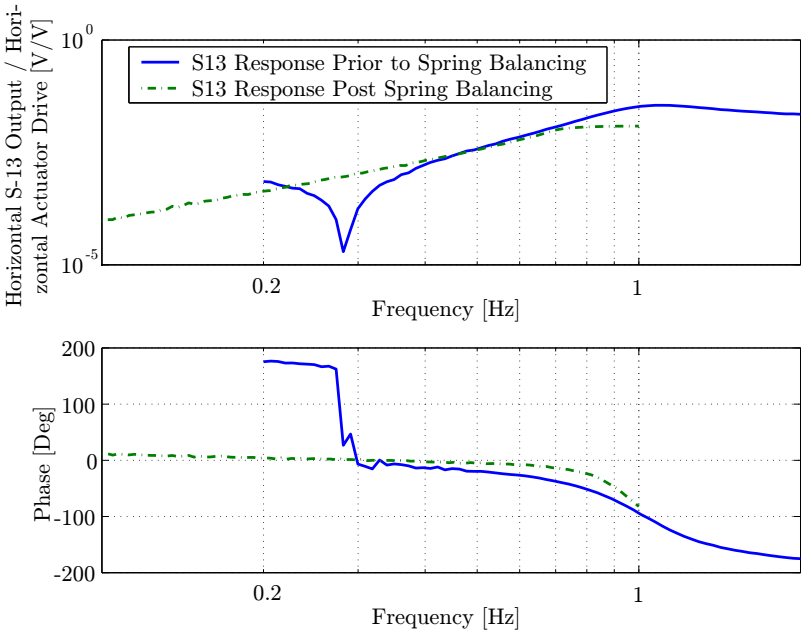

Figure 4.11: Driven transfer function of the horizontal S13 output from the horizontal actuator in the hydraulic test stand. Prior to the installation of additional horizontal springs, tilt dominated the output of the horizontal S13 at frequencies below the tilt-horizontal coupling zero at 0.3 Hz.

the horizontal motion measured by the seismometer is exactly canceled by tilt motion. This observation can be used to determine the preexisting stiffness and center of stiffness of the platform if the added stiffness and the position of the added stiffness are known.

To investigate this procedure, consider the platform in figure 4.12. The stiffness of the horizontal spring K_T at the offset distance h_{kt} represents the rotational stiffness of the platform including the horizontal stiffness of the coil springs, the transverse stiffness of the vertical actuator, and the rotational stiffness implied by the separated vertical coil springs. An additional spring, K_A, is added in parallel, but offset on the opposite side of the actuator a distance h_{ka}. The sum of moments acting at the center of mass (in line with the horizontal actuator) is:

$$\sum M = \theta(K_T h_{kt}^2 - K_A h_{ka}^2) - x(K_T + K_A) \tag{4.6}$$

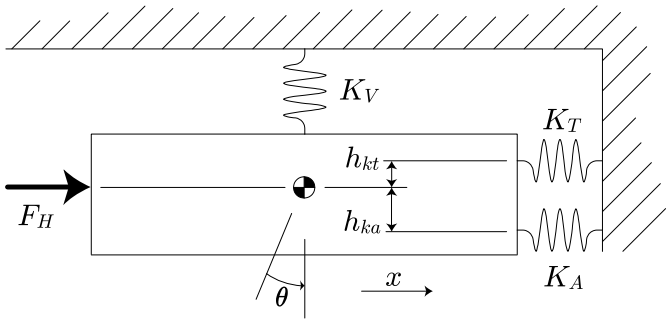

Figure 4.12: A two-dimensional schematic of the compliant elements attached to the hydraulic test platform.

The sum of moments is equal to zero at the tilt-horizontal coupling zero, and at this frequency, $x/\theta = g/\omega_{tilt}^2$ (appendix D.2). Therefore, at the tilt-horizontal coupling zero, equation 4.6 reduces to:

$$g/\omega_{tilt}^2 = \frac{K_T h_{kt}^2 - K_A h_{ka}^2}{K_T + K_A} \tag{4.7}$$

In this equation, the tilt-horizontal coupling frequency, ω_{tilt}; the added spring stiffness, K_A; and the added spring offset, h_{ka}, are known. If ω_{tilt} is known for two values of K_A, it is possible to solve for both K_T and h_{kt}. Once K_T and h_{kt} are known, K_A and h_{ka} can be selected to move the center of stiffness in line with the horizontal actuator. Figure 4.11 shows the effect of this process.

Offload Springs

Coil springs are commonly employed as an offload mechanism to support the static weight of a payload. The popularity of coil springs may be attributed to their compact size and widespread availability. However, the well known shortcoming of the coil spring is axial to rotational coupling. This coupling is caused by the *uncoiling* of the spring as the spring is stretched. This problem is overcome in this experiment by assembling springs from right and left-hand wound halves as is shown in figure 4.13.

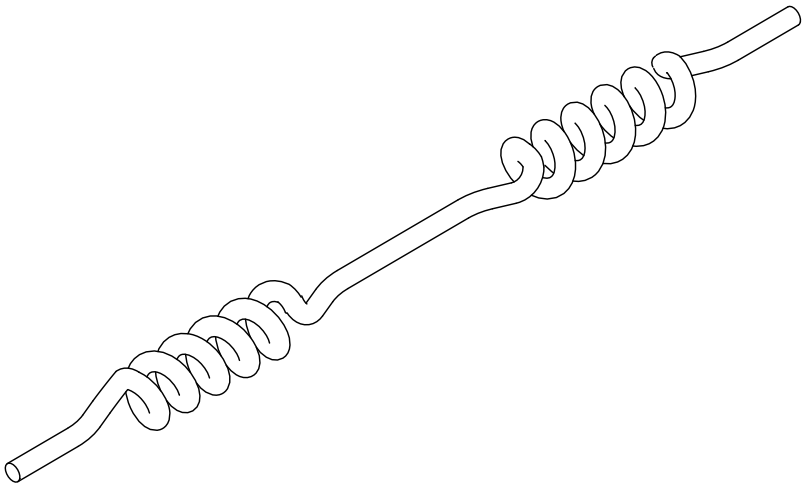

Figure 4.13: An offload spring for the test platform. The offload springs feature both clockwise and counterclockwise turns to reduce coupling of vertical motion into rotational modes.

4.3.3 HEPI

The Hydraulic External Pre-Isolator (HEPI) integrates the final design of the quiet hydraulic actuator with the existing LIGO vacuum chambers. The HEPI system is installed on two styles of vacuum chamber in LIGO, the Beam Splitter Chamber (BSC), figure 4.14, and the Horizontal Access Module (HAM), figure 4.15. Both chambers are similar in that they are surrounded by four vertical piers that support the optical payload inside the vacuum chamber. The hydraulic actuators are placed on top of these piers and beneath the crossbeams which carry the payload inside the vacuum chamber.

A vertical and horizontal actuator are integrated with the offload springs into the housing assembly. A housing assembly is then placed on top of each of the piers and oriented such that the horizontal actuators are arranged tangentially around the chamber. In total there are four vertical and four horizontal actuators per chamber. In this configuration there are eight actuators for six degrees of freedom meaning that the system is overdetermined.

An overdetermined system is not ideal because of the risk of wasting control

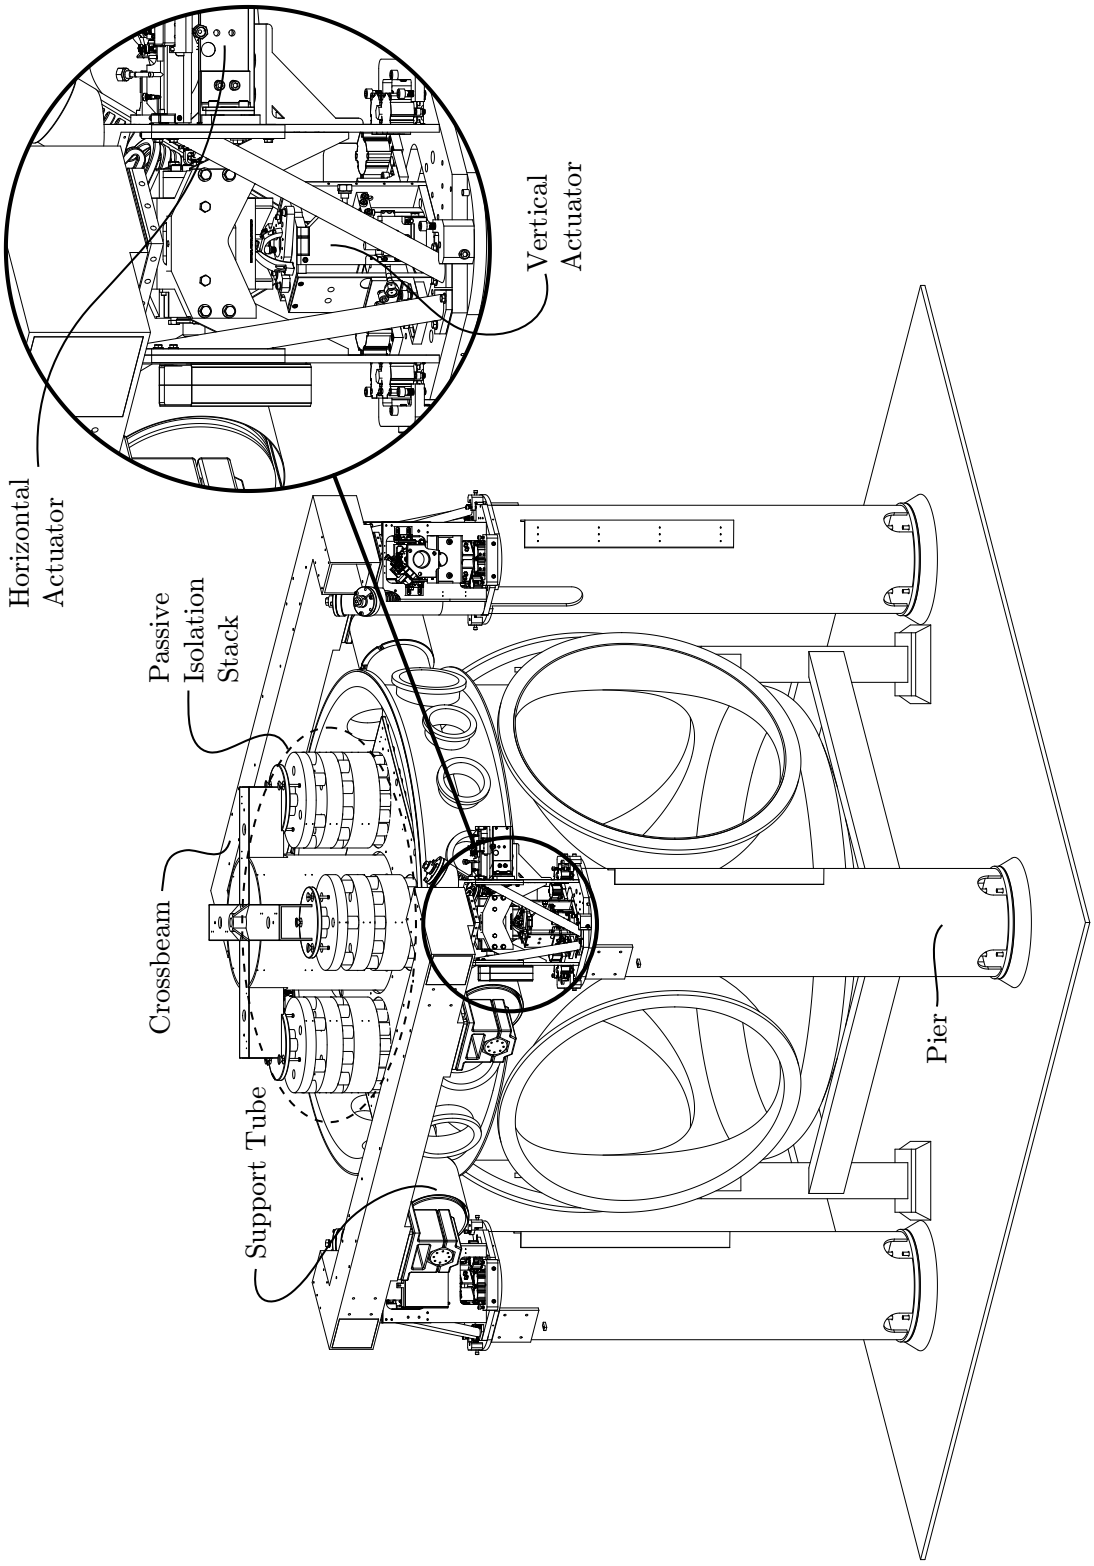

Figure 4.14: The BSC with a detail view of one corner of the HEPI system.

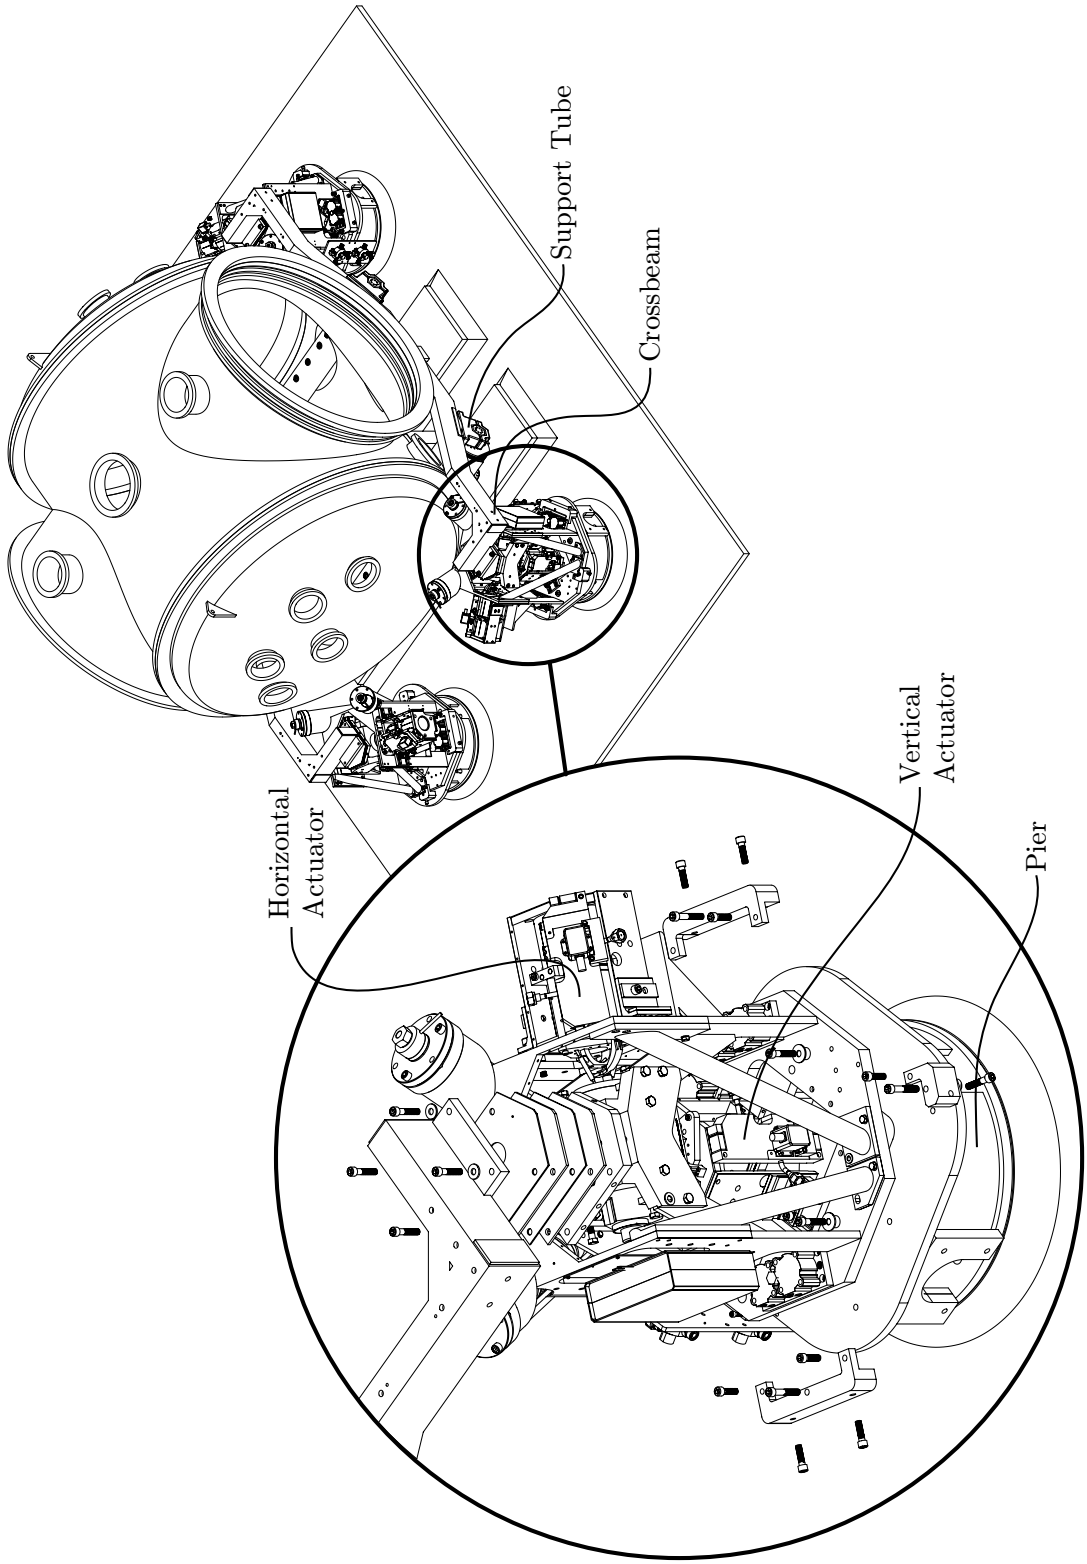

Figure 4.15: The HAM with an exploded detail of one corner of the HEPI system.

effort when two actuators attempt to apply force in opposite directions. While it is straightforward to avoid this problem in the system control algorithm, this particular application is less sensitive because of the flexibility of the payload. The crossbeam weldment and the piers are relatively soft in the overdetermined modes (discussed further in chapter 6) greatly reducing the requirements of the controller.

The HEPI Housing

The HEPI housing, figure 7.1, was designed by LIGO engineer Ken Mason with input from the author and others. The HEPI housing integrates a vertical and horizontal actuator with a pair of offload springs. The design requirements of the housing differ slightly from those discussed for the test platform, but the fundamental themes of collocation and stiffness based design continue to be of great importance.

The additional requirements for the HEPI housing stem from needs for adjustment. In order to reach the LIGO specification of manual payload positioning of ± 5 mm, the HEPI housing has pneumatically actuated ball bearings which facilitate displacement across the pier top.

The attachment surfaces between the pier top and the crossbeam cannot be expected to be parallel. Hence, the housing must incorporate means to adjust the actuator positions to account for this misalignment. Furthermore, it is important that these connections be stiff once the adjustments have been made so as to avoid undesirable flexibility in the foundation. The adjustment is done by adding shims between the sides of the actuators and the housing attachments. The bolt holes for attaching the actuators are generously slotted to allow for a substantial amount of adjustment in the actuator position. However, this adjustment comes at the cost of difficulty during installation. Experience indicates that it can be quite challenging to install and position the actuators with multiple shims. Alternatives to this design are discussed in chapter 7.

Also unlike the test platform, the HEPI system is less sensitive to tilt-horizontal coupling generated by misalignment between the horizontal actuators and the local (per corner) center of stiffness. This is a result of the large separation distance between offload springs on neighboring piers. The distance between offload springs, r_k in figure

4.8, is greater than 3 meters, and due to the substantial weight of the payload, the offload springs are comparable in stiffness to the rotational stiffness of the actuator tripod. Therefore, the alignment of the horizontal actuator with the local center of stiffness is less important in this case.

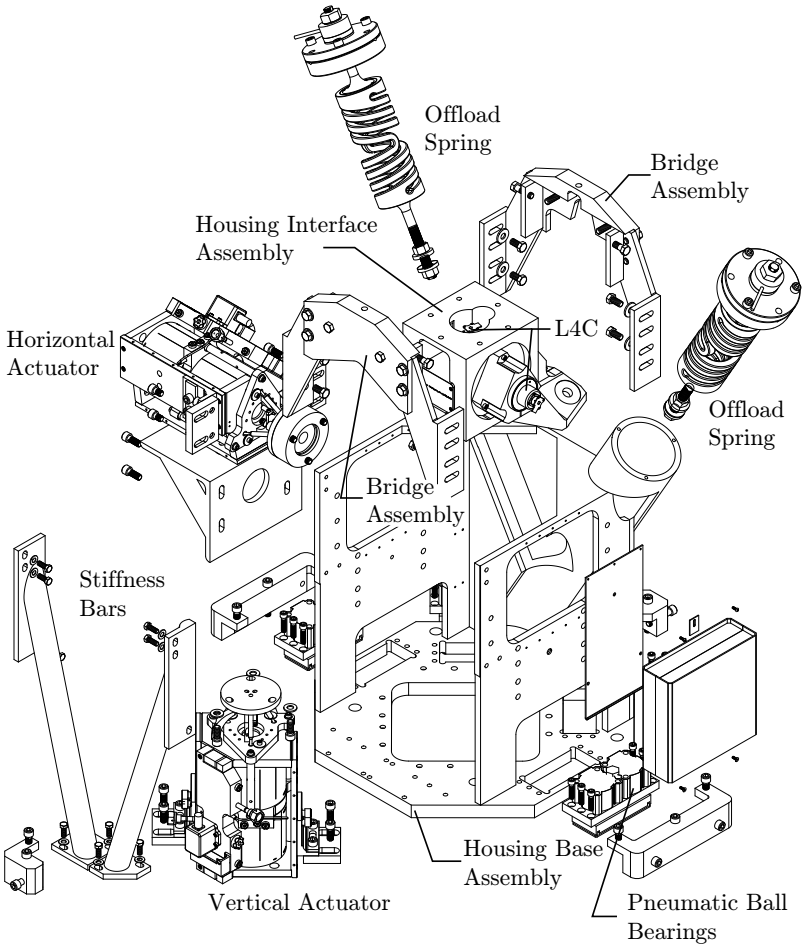

Figure 4.16: An exploded view of the HEPI housing. Courtesy of Oddvar Spjeld.

On each corner of the HEPI system there are two feedback displacement sensors (Kaman DIT-5200) and two feedback seismometers (Marks Products L4C). Both of the displacement sensors and the vertical seismometer are integrated directly into the actuator. The horizontal L4C is clamped into the housing interface assembly which is mounted directly to the horizontal actuator. This arrangement is less desirable than

attaching the L4C directly to the tripod plate because it places the tripod stiffness between the actuation plate and the sensor. Since the tripod is axially stiff, and there is no convenient way to attach the L4C closer to the actuation plate without introducing tilt, this is an acceptable compromise.

4.4 Summary

The underlying theme of this chapter is mechanical design for feedback control applied to the quiet hydraulic actuator and the associated platforms. Many examples of this are discussed in this chapter. The most important of these is collocation. The experience in this lab with the platforms discussed here and several others [12] is that feedback control systems are often limited by how well collocated the actuator is to the sensor. This is the reason why much of the instrumentation is integrated directly into the actuator.

A substantial limitation of all vibration isolation systems is tilt-horizontal coupling. If seismometers are used to measure horizontal motion, the platform must be designed to minimize tilt. Minimizing tilt is a difficult design problem that is frequency dependent, and must be addressed in the design of the platform. However, using the techniques discussed in this chapter, it is possible to obtain reliable information from horizontal seismometers down to 0.5 Hz.

Finally, the results from the first prototype actuator on the test platform (figure 4.5) show that the bellows exhibit a breathing compliance that causes the bellows breathing resonance. This sharp peak can easily be eliminated in software, but this approach is accompanied by a significant loss in robustness. In the following chapter, alternative methods are used to eliminate this peak in hardware.

Chapter 5

The Bellows Breathing Resonance

At the conclusion of Chapter 4 an undesirable resonance was identified in the measured open loop behavior of the actuator. The resonance is undesirable for control of the actuator because the peak must be inverted if bandwidths near or above the peak are desired. This peak can be inverted by a number of standard controls techniques, and it will be shown that the actuator can perform well in the presence of this resonance. However, this type of plant inversion does not lend itself to the long term robustness that is desired for LIGO, and therefore work has been devoted to suppressing this resonance passively.

In this chapter, two means for suppressing this resonance are investigated. The first, more obvious, solution is to redesign the bellows for an improved breathing stiffness. The second method is to introduce a network of hydraulic resistances and capacitances to passively damp the undesirable mode.

5.1 The Bellows

The bellows must be axially compliant to allow the actuator to move, but it is undesirable for the bellows to volumetrically expand, or breathe, when a varying internal pressure is applied. The resonance of chapter 4 was the first indication that the bellows exhibited a breathing stiffness that is significantly lower than the stiffness expected by the properties of the working fluid. This hypothesis was easily verified

with three-dimensional finite element analysis (FEA) models. The ultimate goal in redesigning the bellows is to create a bellow which has a breathing stiffness comparable to that of the stiffness inherent to the working fluid.

In order to develop such a bellow, it is necessary to understand the relationship between the various elements of the bellow geometry and the breathing stiffness. Analytical results are used to explore this design space, and in many cases these results are confirmed by FEA models. Before discussing the trends in geometry, it is useful to define breathing stiffness in a way which can be compared to the stiffness of the working fluid.

5.1.1 Breathing Stiffness

Breathing stiffness is a measure of a structure's ability to maintain a constant volume in the presence of a changing internal pressure. This is analogous to rigidly attaching the ends of the bellow to a pair of flanges that are fixed in position and then introducing a varying internal pressure. The effect of the internal pressure is to cause features of the bellow convolution geometry to distort or bulge particularly in the straight sections of the convolutions (Figure 5.1).

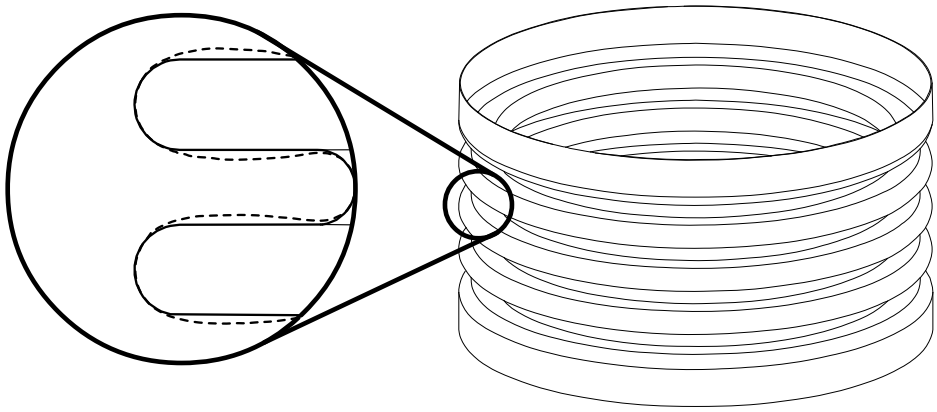

Figure 5.1: Bellow convolution distortion due to internal pressure.

Quantitatively, the breathing stiffness is the change in volume caused by a change in pressure: $\frac{\Delta V}{\Delta P}$. However, in order to compare different bellow geometries, it is

convenient to define a spring rate representative of this stiffness. This can be obtained by manipulation of Hooke's law, $dF/dx = K$. If the change in volume of the bellow is expressed as the bellow cross-sectional area multiplied by the change in displacement: $\Delta V_b = A_b \Delta x$, and the change in pressure in the bellow is expressed as the force divided by the bellow cross-sectional area: $\Delta P_b = \frac{\Delta F}{A_b}$, then the breathing stiffness can be expressed as a linear spring stiffness:

$$K_{breathing} = \frac{\Delta P_b A_b^2}{\Delta V_b} \tag{5.1}$$

The breathing stiffness of the bellows introduces a new compliance to the dynamic equations of motion for the actuator. In particular, equation 3.19 and 3.20 become:

$$\dot{P}_1 \left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}} \right) = Q_1 + A(\dot{z_a} - \dot{z_g}) \tag{5.2}$$

$$\dot{P}_2 \left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}} \right) = Q_2 - A(\dot{z_a} - \dot{z_g}) \tag{5.3}$$

where $A_b^2/K_{breathing}$ results from the breathing compliance of the bellows and is added in series with the fluid compliance term, V_b/β.

5.1.2 Bellow Design

In view of Figure 5.1, it is evident that omitting the straight sections of the convolutions will improve the breathing stiffness. However, the axial compliance of the bellow which is essential to the function of the actuator is dependent on the convolution geometry and is greatly improved by the straight sections in the convolutions. Hence, in the search for an optimum bellow convolution geometry, it is not merely the breathing stiffness which must be maximized, but rather, the ratio of breathing stiffness to axial stiffness: $\frac{K_{breathing}}{K_{bel}}$.

The design of the bellows is based on convolution geometry that has no flat sections. While it is evident that this change will improve the breathing stiffness, the challenge here is to design a bellow that is sufficiently compliant axially while demonstrating the best possible ratio of breathing to axial stiffness.

Roark's book [24] (Table 30, Case 6b and 7b) features solutions for a bellow with circular convolutions (sections of a toroid) based on a paper by Clark [25]. These solutions include both axial and internal pressure loadings, and are useful for determining the scaling effects caused by changes in the convolution geometry.

The convolution geometry is described by three variables: the outside radius, R_{bel}; the convolution radius, r_{bel}; and the thickness, t_{bel} (figure 5.2). Each of these variables is confined to a specific range by either size or manufacturing limitations. As discussed in section 4.1.2, much of the actuator size is dictated by the overall dimensions of the bellow. In order to ensure that the actuator will be compatible with the systems at LIGO, the length and diameter of the bellow are limited to 12.5 cm (5 inches). Adopting this even aspect ratio has the additional advantage of making the bellow insensitive to squirm.

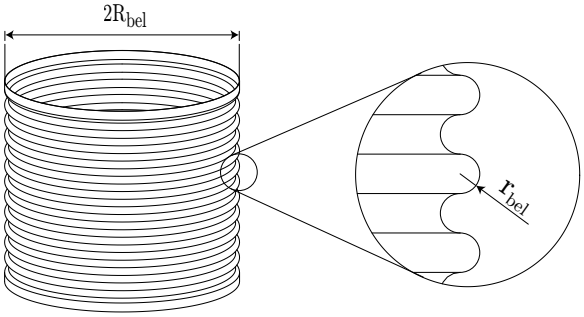

Figure 5.2: Some of the pertinent dimensions of a bellow with no flat sections in the convolutions.

The effect of changing various elements of the bellow geometry is complicated by the axial stiffness requirement. In order for the actuator to function, the bellow must be an order of magnitude more axially compliant than the offload spring stiffness (section 4.1.2). Therefore, the maximum allowable axial stiffness is set at 8.5×10^4 N/m. This encumbers the scaling because the minimum number of convolutions required to meet this axial stiffness changes for each combination of bellow geometry. The approach taken here is, for each set of bellows dimensions, to calculate the number of convolutions needed to exceed the axial compliance requirement. For that number of convolutions, the breathing stiffness is calculated, and the ratio of breathing to

axial stiffness determined.

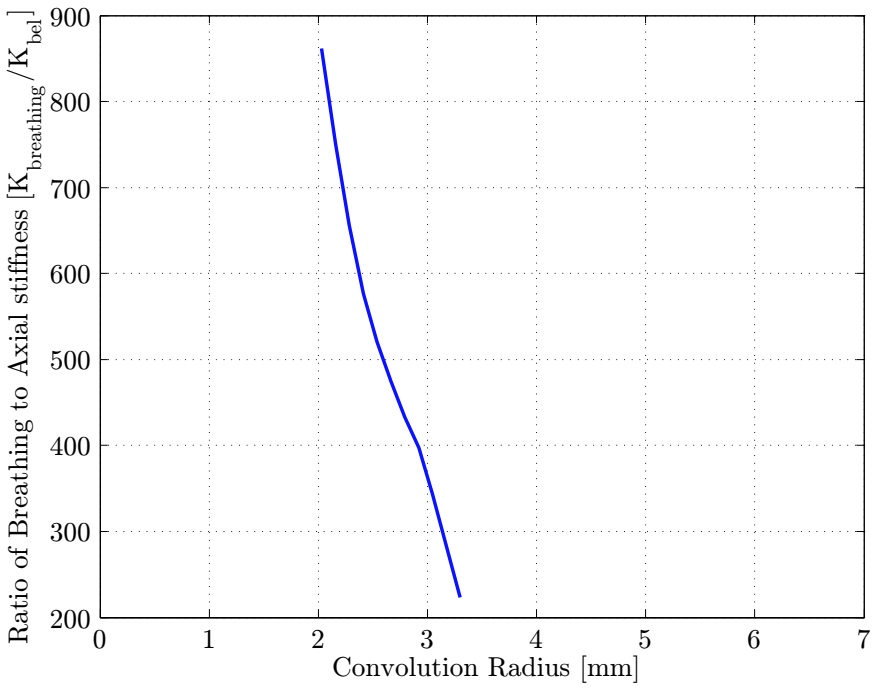

Figure 5.3: The effect of varying convolution radius on the breathing to axial stiffness ratio for a bellow thickness of 0.15 mm (0.006 in.) and outside diameter of 12.5 cm (5.0 in.). The trace for most of the convolution radii shown has been omitted because for those convolution radii, the axial stiffness requirement causes the number of convolutions to be large. In these cases, the large number of convolutions cause the bellow to be unacceptably long.

If the outside diameter of the bellow is set, for the moment, at 12.5 cm (5.0 in.), the ratio of breathing to axial stiffness increases sharply with decreasing convolution radius as shown in figure 5.3. In this figure, the thickness is set to 0.15 mm (0.006 in.). For larger thickness and the same range of convolution radii, the number of convolutions required to meet the axial compliance causes the bellow to be too long. It is fortunate that a thickness of 0.15 mm is acceptable because it is also at the lower limit of the available fabrication capabilities, and the upper limit for allowable stress under the actuator operating pressures.

The thickness of 0.15 mm (0.006 in.) provides the broadest range of possible

convolution radii for an outside diameter of 12.5 cm. In figure 5.4, the thickness is held constant at 0.15 mm, and the both the convolution radius, r, and the outside radius, R are varied. Here again, the length limit of 12.5 cm is imposed. This limits the number of allowable convolutions and causes most combinations of r and R to yield unacceptable results. The areas in the figure where the trace is shown are combinations that are compatible with the size limitations. In the figure, it is apparent that increasing the outside radius and decreasing the convolution radius results in the best ratio of breathing to axial stiffness.

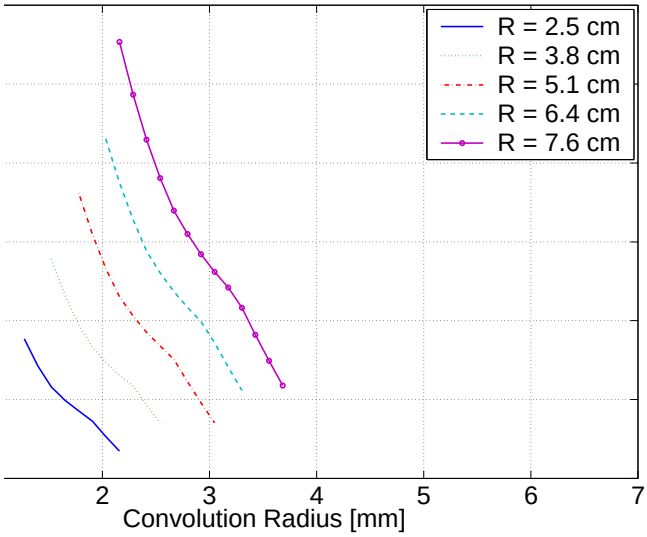

Figure 5.4: The effect of varying convolution radius and outside radius on the breathing to axial stiffness ratio for a bellow thickness of 0.15 mm (0.006 in.). The traces on the plot are not shown in the regions where the number of convolutions required to exceed the axial stiffness limit cause the bellow to be too long.

Based on the results of figures 5.3 and 5.4, the final dimensions chosen were: outside diameter of 12.5 cm (5.0 in.), length of 12.5 cm (5.0 in.) and convolution radius of 2.1 mm (0.850 in.) with 12 convolutions. For this combination of dimensions, a FEA model was created in the Algor software environment. The ratio of breathing to axial stiffness predicted by Algor was 876 which compares decently with the 750 predicted by the analytical methods.

5.1.3 Final Bellow Design

The final design of the bellow was manufactured to the dimensions listed above. The breathing stiffness of this bellow was predicted to be 8.5×10^7 N/m, a factor of nine better than the original bellow, but still below the stiffness of the working fluid, 1.5×10^8 N/m for these dimensions. Assuming that the resonant frequency is dependent on the bellows breathing stiffness, this should result in a factor of three increase in frequency (as frequency changes with the square root of stiffness); moving the bellow's breathing resonance out to nearly 100 Hz. Unfortunately, this was not observed. The resonance moved from 27 Hz to 38 Hz with the new bellows. The cause for this shortcoming stems from the compliance of the connection between the actuator and the mass. At 38 Hz, the connection stiffness becomes the dominant compliance and masks the improved stiffness of the bellows.

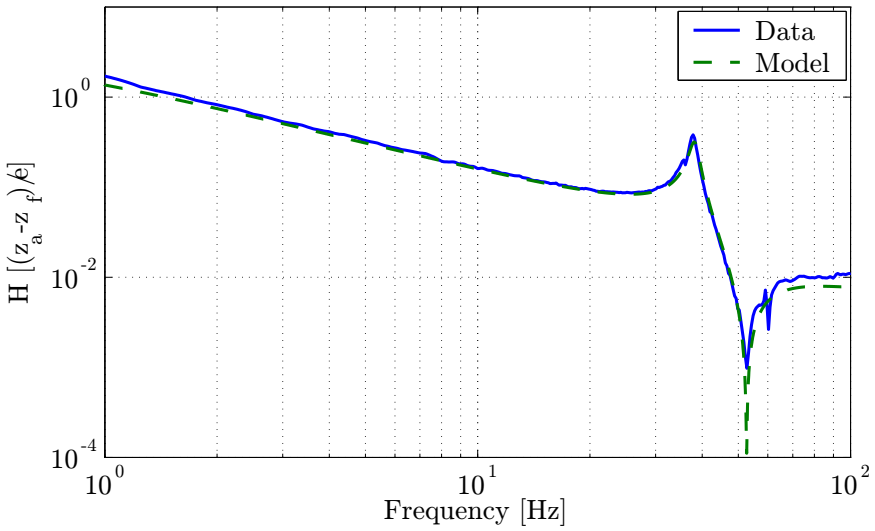

Figure 5.5: A transfer function from the valve to the displacement sensor with the new bellows geometry.

5.2 The Bypass Network

As an alternative to increasing the stiffness of the bellows, the concept of the bypass network is to dissipate undesirable energy associated with the aforementioned resonance. The bypass network is the hydraulic equivalent of a resistor and capacitor in series between the two bellows.

The bypass network is a path between the two bellows (Figure 5.6) with impedance to flow that varies with frequency. While at no time is there fluid exchange between the two bellows, at high frequencies the diaphragm becomes compliant and there is an exchange of volume. The oscillation of the diaphragm draws fluid through the resistance, R_{byp}, and in so doing, dissipates energy in the form of heat.

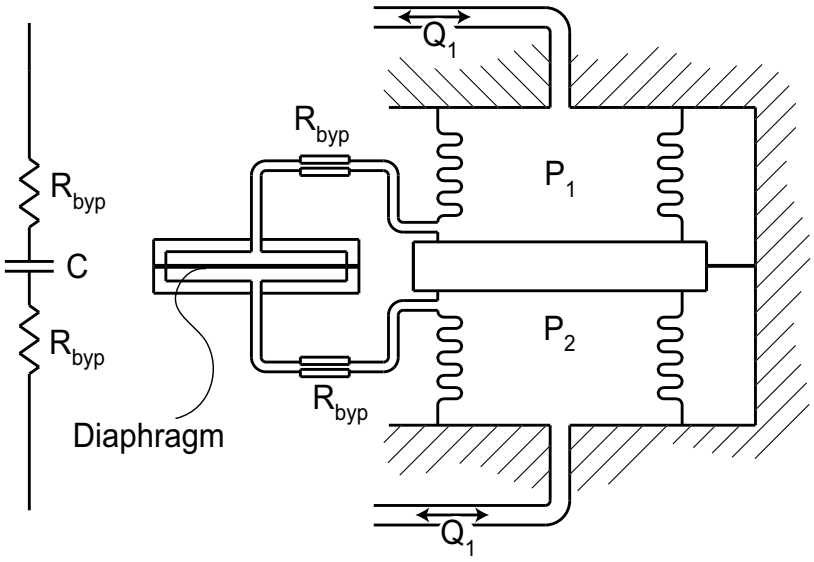

Figure 5.6: The electrical representation of the bypass network, and a schematic of the hydraulic bypass network shown between the stacked bellows.

Bypass Capacitor

The bypass network incorporates a clamped plate to serve as a diaphragm. The *volumetric compliance* of the diaphragm creates the effect which is analogous to electrical capacitance. Volumetric compliance is similar to the breathing stiffness discussed

in Section 5.1.1 except for the shape of the compliant surface. For the form factor of a flat plate, the capacitance of the diaphragm is dependent on both the stiffness and surface area of the diaphragm. The expression for capacitance, C_{byp} is derived starting with Hooke's law:

$$F = K_{byp}x \tag{5.4}$$

$$\Delta P A_{byp} = K_{byp}\frac{Q}{A_{byp}s} \tag{5.5}$$

$$\frac{\Delta P}{Q} = \frac{K_{byp}}{A_{byp}^2 s} \tag{5.6}$$

$$C_{byp} = \frac{A_{byp}^2}{K_{byp}} \tag{5.7}$$

where K_{byp} is the stiffness of the diaphragm, A_{byp} is the area of the diaphragm, ΔP is the differential pressure, and Q is a flow through the bypass resistance.

5.2.1 Bypass Network Pole Frequency

The choice of frequency for the pole associated with the bypass network, $f_{byp} = \frac{1}{2\pi R_{byp} C_{byp}}$, is a compromise between control authority and suppression of the undesired resonance. Clearly, if the pole is set too low (the diaphragm compliance large), the actuator will cease to deliver the required force, and if the pole is too high, the suppression of the resonance will be minimal.

The effect of the bypass pole frequency can be qualified using a Matlab model of the actuator installed on the LIGO BSC chamber (the system which the actuator is ultimately intended for). Using this model, the frequency of the bypass pole is moved to three different representative frequencies: 1, 8 and 50 hertz. The reduction of the resonance (Figure 5.7) worsens considerably when the pole frequency is shifted from 8 to 50 hertz, but based on the plot, a pole frequency lower than 8 hertz may appear compelling.

While a low frequency bypass pole may be advantageous for damping performance, it is difficult to obtain the required compliance in the diaphragm without making the diameter undesirably large. The minimum size of the diaphragm is limited by the

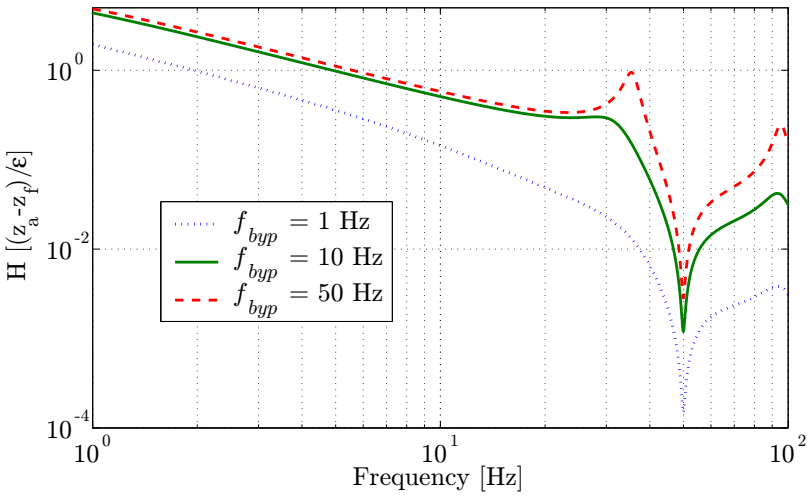

Figure 5.7: The transfer function from valve drive to platform displacement with various bypass pole frequencies. The model used to generate this figure is derived in appendix B.

stress incurred by the differential pressure across the bellows. For any desired diameter of diaphragm, the thickness can be reduced to reach the specified compliance, but the resulting stress is often unacceptably high.

The predicted stress in the diaphragm for various pole frequencies and diaphragm radii is plotted in Figure 5.8. In the plot, it is assumed that the bypass resistance is held fixed at $2{\cdot}10^9 \frac{Pa{\cdot}s}{m^3}$, and only the value of the capacitance is varied for each pole frequency. The value of capacitance is used to determine a diaphragm stiffness, K_{byp}, and given the diaphragm stiffness and radii, an appropriate thickness can be calculated. The stress in the diaphragm is then predicted for an applied pressure of $2.8{\cdot}10^5$ Pa (40 psi) (approximately twice the expected differential pressure). The stress prediction is calculated using Roark's Book [24] Table 24, case 10b for a circular plate with fixed edges.

Based on the Matlab model of the actuator installed on the LIGO BSC chamber (Figure 5.7), a relatively low-frequency bypass pole (near 1 Hz) may seem attractive for the standpoint of improved damping. However, it is argued here that the improvement in damping is not sufficient to justify compromising the mechanical design with a large diaphragm plate. In consideration of this, a compromise is made by setting

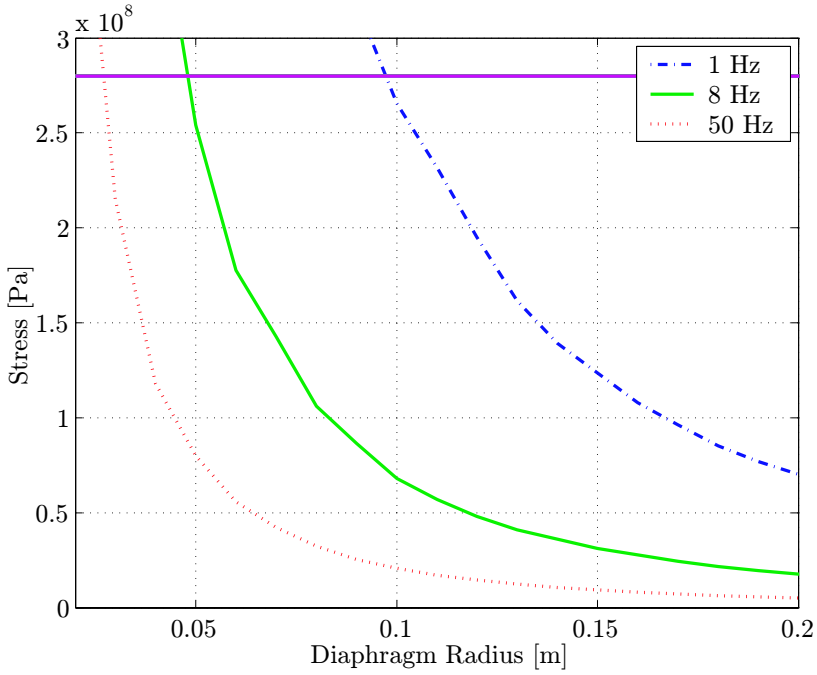

Figure 5.8: The predicted stress in the diaphragm for various pole frequencies and diaphragm radii. In this figure, the bypass resistance is held fixed while the capacitance is varied for each pole frequency. For each value of capacitance, a volumetric stiffness, K_{byp}, is calculated for a given radii using Equation 5.7. Given a stiffness and radius the required diaphragm thickness can be calculated and the associated stress is predicted assuming $2.8{\cdot}10^5$ Pa (40 psi) of pressure.

the bypass pole frequency at 8 Hz.

5.2.2 Bypass Resistance

Instead of increasing the capacitance to lower the frequency of the bypass pole, an alternative is to increase the bypass resistance. While this is a natural approach in electrical circuits, it is not easily realized in hydraulics because the onset of parasitic inductance occurs at surprisingly low frequencies.

Parasitic inductance first appeared in a prototype of the bypass network. In this first attempt, the fluid paths to the diaphragm were designed to be small enough to

create the bypass resistance. The inductance of a round pipe is:

$$L_{round} = \frac{4\rho \cdot length}{\pi \cdot diameter^2} \tag{5.8}$$

Hence, if the length is relatively large and the diameter small to provide ample resistance, the inductance can be overwhelming. In the case of the prototype, the inductive impedance dominated the bypass path in the frequencies of interest rendering the bypass completely ineffective. Based on this experience, a more realistic model of the actuator and bypass network is shown in Figure 5.9.

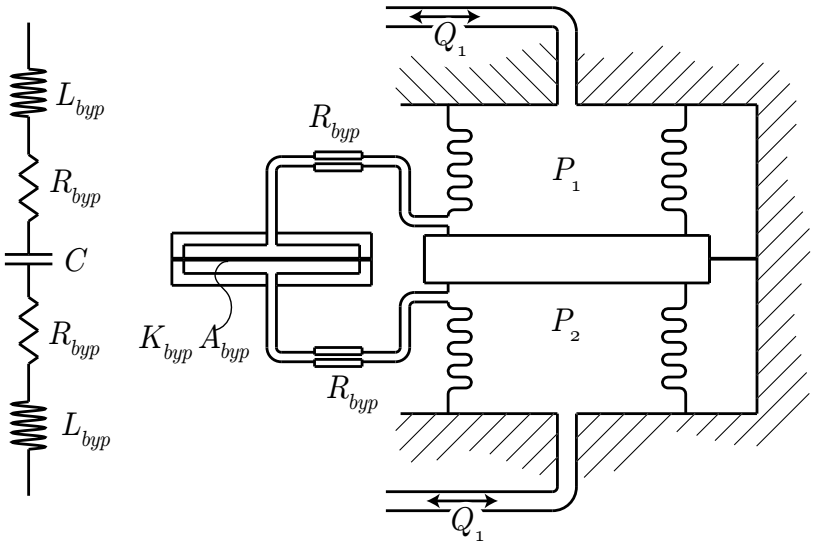

Figure 5.9: A more realistic schematic of the hydraulic actuator and bypass network.

After the failure of the prototype, the design of the bypass resistance was reevaluated with attention to the associated inductance. Both a parallel plate and a tubular configuration were considered for this evaluation. The resistance is designed by forcing the ratio of the inductance to the resistance to be small at the frequency of the resonance.

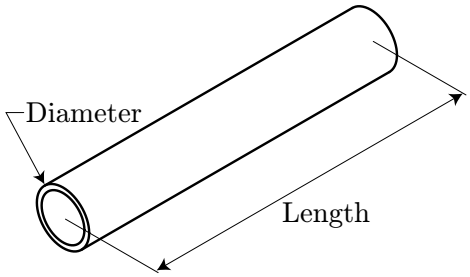

Figure 5.10: A circular tube type resistor.

Circular Tube

For a circular tube, the ratio of inductive to resistive impedance is:

$$\frac{L_{round}}{R_{round}} = \frac{\frac{4\rho \cdot length}{\pi \cdot diameter^2}s}{\frac{128\mu \cdot length}{\pi \cdot diameter^4}} = \frac{\rho \cdot diameter^2}{32\mu}s \tag{5.9}$$

In this relationship, the frequency variable, s, is set to the frequency of the breathing resonance at ∼30 Hz. The ratio of the inductance to the resistance, L_{round}/R_{round}, is set to 1/100 to ensure that the dominate impedance at the resonance is resistive. With this and equation 5.10, it is possible to solve for the diameter. If the viscosity is set to be 100 times water and the density equivalent to water (reasonable values for mineral oil), the resulting diameter is 1.64×10^{-4} m. For a desired bypass resistance of 2×10^{9} Pa·s/m^3, the associated length can be calculated to be 3.42×10^{-7} m.

There are several problems with this result. The diameter is very small which will make the resistor sensitive to clogging. More importantly, since the diameter is larger than the length, it is exceedingly unlikely that flow through the resistor will be fully developed and laminar. Without laminar flow, these equations do not hold and unsteady flow may result.

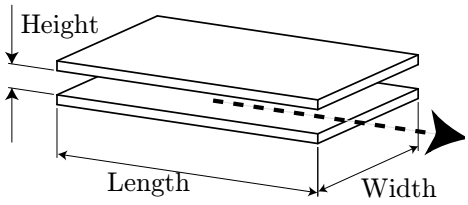

Figure 5.11: A parallel plate type resistor.

Parallel Plates

Another means for making a hydraulic resistance is a pair of parallel plates. Here again, the ratio of inductive to resistive impedance is the important figure of merit:

$$\frac{L_{parallel}}{R_{parallel}} = \frac{\frac{\rho \cdot length}{width \cdot height}^{s}}{\frac{12\mu \cdot length}{width \cdot height^3}} = \frac{\rho \cdot height^2}{12\mu} s \tag{5.10}$$

As was done for the circular tube, this equation is used to solve for the height. For the values above, the height is 2.5×10^{-4} m. However, when solving for the length, the parallel plate resistance offers another degree of freedom: the width.

This additional degree of freedom is valuable because it enables the length to be set to create steady flow conditions. To guarantee steady flow through a constricted area, the length must be 10 times the height [27]. Given a height of 2.5×10^{-4} m, the length can be set to 2.5×10^{-3} m. With the length and height defined, the equation for parallel plate resistance is used to solve for the width. This yields 0.1 m which is rather large by comparison to the dimensions of the actuator.

5.2.3 Final Bypass Design

The crucial lesson learned from the bypass prototype was that the fluid path to and from the bypass must be extremely short in order to avoid parasitic inductance. The most obvious solution to this dilemma is to place the bypass network directly between the two bellows. However, the best candidate resistance, the parallel plate

configuration, is too large to fit within the bellow diameter.

The solution to this predicament is to distribute the width of the parallel plate resistance around a circular perimeter (Figure 5.12). Using this method, the resistance width can be made almost arbitrarily large by adding additional circular plates to the *resistor stack*. This resistor stack is affixed to the actuator plate and conveniently hidden within the lower (actuator in the vertical orientation) bellow.

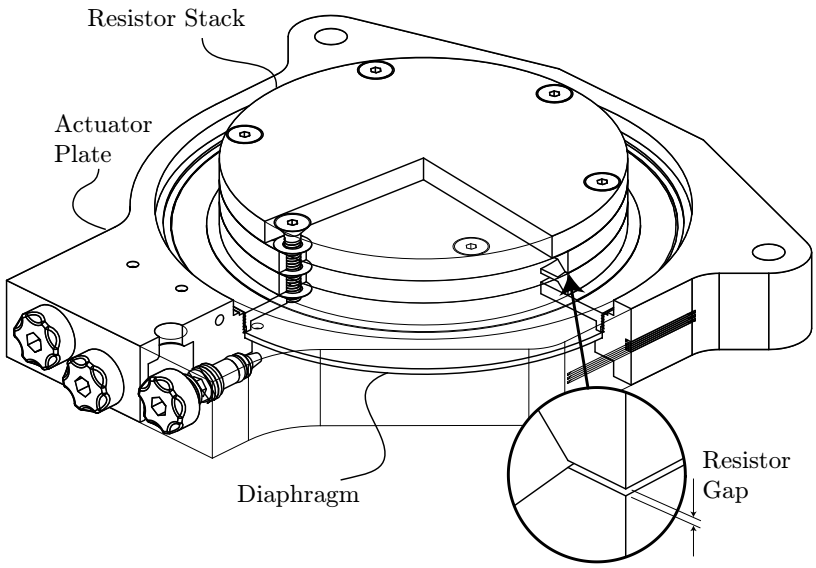

Figure 5.12: The final design of the bypass network. The parallel plate resistance is comprised of long, narrow gaps in the resistor stack (shown in the detail view).

Due to the ease of increasing the width of the parallel plate resistance, the final bypass design features an inductive to resistive impedance ratio that is 1/600 at the bellows resonance. The corresponding dimensions are: height of 1.0×10^{-4} m, length of 1.0×10^{-3} m and width of 0.6 m. This width is distributed over three gaps that are 6.8×10^{-2} m in diameter.

5.2.4 Bypass Performance

When properly designed, the bypass network provides exceptional reduction of the bellows breathing resonance (Figure 5.13). The model for the actuator with the

bypass network is discussed in appendix B.

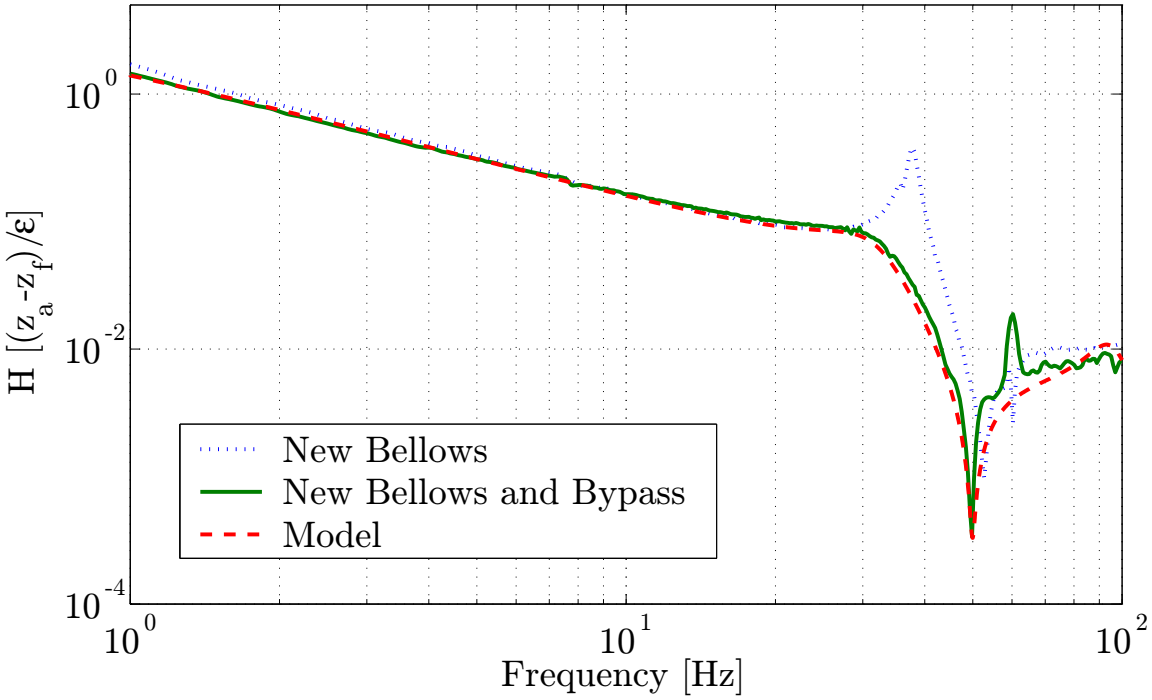

Figure 5.13: A transfer function of valve drive to displacement sensor output for a single vertical actuator in the quiet hydraulic test platform. The performance of the bypass network is evident in the reduction of the bellows breathing resonance near 40 Hz. The resonance is at a slightly higher frequency than discussed previously because of the improved bellow geometry.

Chapter 6

Control Synthesis

The performance requirements for LIGO specify parameters for both alignment and isolation. These dual, and sometimes opposing, objectives require several different sensors to be used together in control schemes that allow for position control while simultaneously providing isolation.

The control architecture for the hydraulic actuator evolved with the actuator and the platforms that support the actuator. The first platform, the Hydraulic Test Platform (section 4.3), served as a simple, two degree-of-freedom, demonstration of the hydraulic actuator, and a test bed for developing the two underlying control techniques: sensor blending and sensor correction.

Once the actuator was shown to meet specification on the Test Platform, eight actuators were installed on the BSC chamber at the LIGO Advanced System Test Interferometer (LASTI) laboratory at MIT. The LASTI BSC platform posed several challenges including overdetermined control, eight actuators to control six degrees of freedom.

While the alignment performance of the actuator is of great importance to the functionality of the LIGO instrument, the results shown in this chapter focus on isolation performance. The justification for this is that the if the displacement sensor noise characteristics exceed the LIGO position specification and there is ample gain on the displacement sensor feedback loop then alignment performance can be assumed. This assumption has since been verified by the hydraulic actuators installed at the

LIGO Livingston Observatory (LLO).

6.1 Sensors

To satisfy the alignment criteria, a displacement sensor is incorporated into the design of the actuator. This provides alignment information down to zero-frequency. Alternatively, active isolation from the ground requires an inertial reference. This is provided by seismometers stationed on both the ground and the platform. Sensitivities and noise information for all of the sensors used on this project may be found in appendix D.1.

6.1.1 Displacement Sensors

In the context of this experiment, the displacement sensor measures the difference between the position of the actuator plate and the foundation: $z_a - z_f$ (figure 3.1). If the actuator tripod and the foundation are assumed to be infinitely stiff then this is the same as the difference between the mass and ground position: $z_m - z_g$.

The displacement sensor is incorporated into the actuator to provide the best opportunity for collocation with the actuation plate. Collocation of the sensor and actuator is critical to obtaining controllability at higher frequencies (in this case, above about 10 hertz, but in general as discussed in appendix C.3).

The choice of displacement sensor adapted with the development of the actuator. The first actuator featured an optical displacement sensor designed and manufactured at Stanford. Preliminary results with this sensor indicated the noise from this sensor was unacceptably close the LIGO noise requirements. Because of this, later editions of the actuator featured the DIT-5200 manufactured by Kaman sensors. The noise floor of the DIT-5200 is impressive, approaching 10^{-10} m/$\sqrt{\text{Hz}}$ in the 1 to 10 Hz. band.

6.1.2 Seismometers

Seismometers are useful in this work as an inertial reference for either the ground or the platform. An inertial reference is something that remains fixed relative to the motion of the environment. This is best explained by describing the construction typical of these type of sensors.

The most common embodiment of a seismometer is a suspended magnet surrounding by a readout coil (figure 6.1). The readout coil is attached to the case of the seismometer and when the case moves at frequencies higher than the natural frequency of the magnet suspension, the relative motion induces a voltage in the surrounding coil proportional to the relative velocity.

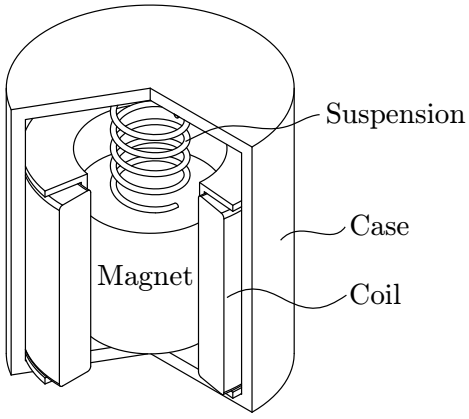

Figure 6.1: A representative section view of a seismometer. The suspension spring suspends the magnet whose motion is read out by the surrounding coil. Practical seismometers typically feature a more compact suspension system and two counter-wound coils designed to buck stray magnetic fields. This configuration is also referred to as a geophone.

The difficulty with seismometers is that below the natural frequency of the suspended mass, the mass begins to move with the case and the relative motion ceases to be detectable. At this point, the velocity signal output is reduced while the noise remains and rapidly dominates as the frequency decreases.

When measurement to a lower frequency is required, it is common to use active feedback to control the position of the suspended mass (the magnet in figure 6.1).

In this embodiment, the output of the seismometer becomes the control authority necessary to maintain the suspended mass in its centered position (proportional to the acceleration of the mass). Active feedback allows for very high natural frequencies and a wide bandwidth for acceleration measurement. It also offers improvements in linearity since the suspended mass is held fixed with respect to the case by the feedback control.

6.2 Control Techniques

Both sensor blending and sensor correction are useful techniques for either single input/single output (SISO) or multi-input/multi-output (MIMO) systems. Discussions of sensor correction techniques can be found in the literature [30], and an optimal solution for sensor blending has recently been developed by Wenshang Hua [31].

6.2.1 Sensor Blending

Sensor blending is commonly used when no individual sensor can provide all of the desired information over the frequency band of interest. In these instances, the output of two more sensors are combined to make a *supersensor*.

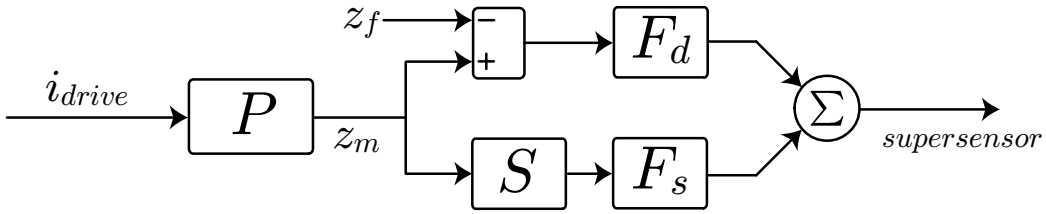

Figure 6.2: A block diagram for the open-loop sensor blending between a displacement sensor and a seismometer. The plant, P, includes the actuator and suspension dynamics and is commanded by electrical current input, i_{drive}, to the valve. The displacement sensor has neglible dynamics and measures the difference between the platform position, z_m, and the foundation position, z_f. The seismometer, S, is filtered by the blending filter, F_s, before being combined with the filtered displacement sensor output to make the supersensor.

In applications where alignment and vibration isolation are important, as in LIGO, a common embodiment of the supersensor is a displacement sensor blended with a seismometer. For feedback control, the supersensor is more desirable than either of the individual sensors for several reasons. Feedback control based solely on the displacement sensor provides good low-frequency alignment with respect to the ground, but feedback control with only the displacement sensor will cause the platform to move with the ground. This is directly in opposition to the goals of vibration isolation. Alternatively, a seismometer on the platform may be used to provide an inertial reference. Tracking the output of an inertial reference will provide isolation from the ground, but only at frequencies above the natural frequency of the seismometer. At frequencies below the seismometers' natural frequency, the signal to noise ratio becomes rapidly unacceptable with decreasing frequency (as discussed in the previous section, seismometers are inherently AC coupled devices). A controller based solely on seismometer output cannot maintain platform alignment and cannot be used at low frequencies. Hence, a supersensor that is the combination of a displacement sensor and a seismometer can provide both isolation and alignment over a broad range of frequencies. Figure 6.3 shows how a displacement sensor and seismometer may be combined in practice. The displacement sensor and seismometer filters are developed by manually shaping the output in the frequency domain. In figure 6.3, the displacement sensor filter is an elliptic filter that provides substantial suppression at frequencies above 10 hertz. The purpose of this aggressive filtering is to ensure that the displacement sensor output does not impede isolation performance at higher frequencies. Similarly, due to the suppression of the displacement sensor, the seismometer output is emphasized at higher frequencies providing inertial information where isolation performance is desired.

Filters developed in this manner have no guarantee of creating a supersensor sensor that is amenable to control, or even stable, in a feedback scenario. The most common problem is the introduction of non-minimum phase (NMP) zeros. NMP zeros often appear when two sensors outputs are combined with phase differences approaching 180 degrees in a band where their magnitudes are near equivalent. Most often, both sensors have equal magnitude only in the vicinity of the blend frequency,

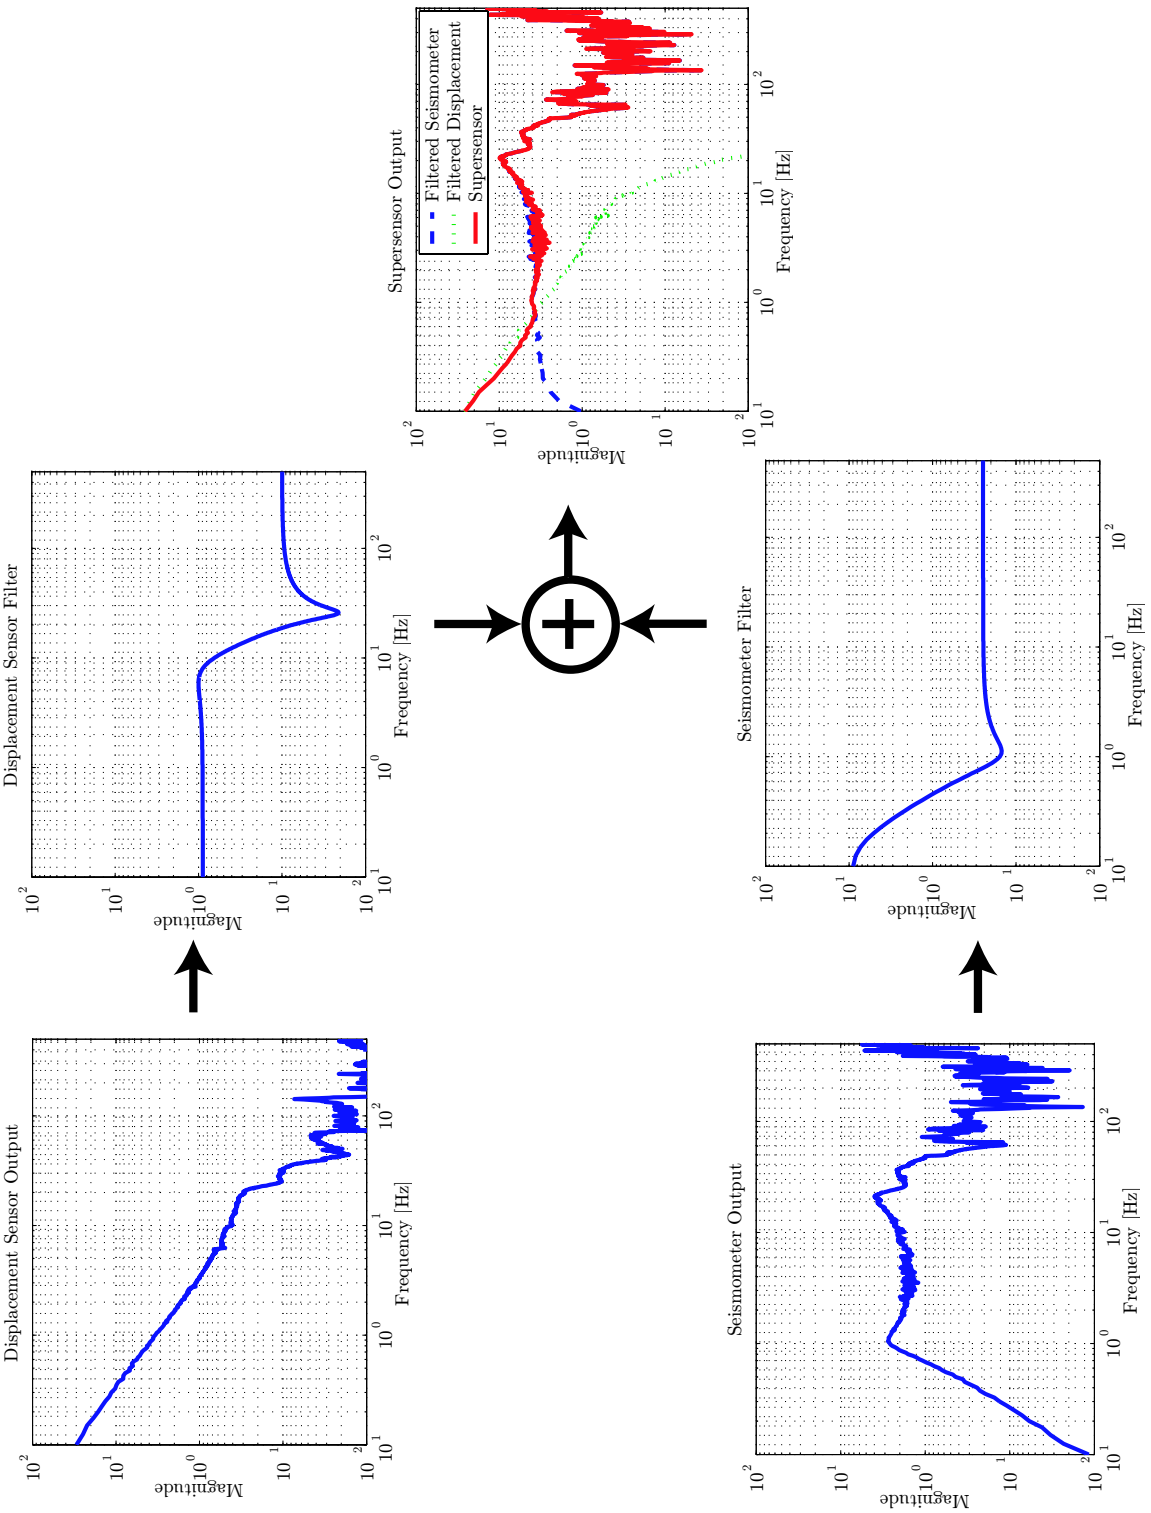

Figure 6.3: An example of sensor blending for a displacement sensor and seismometer. The left most plots are driven transfer function measurements (measured sensor output to random noise input applied to the valve).

and therefore, it is prudent to arrange the filtered sensor outputs to have similar phase at the blend frequency. This is difficult because, at the blend frequency, one sensor must decrease in magnitude while the other must increase. As is obvious in the figure, the usual compromise is that the blend does not occur at a particular frequency, but instead, over a range of frequencies.

Complimentary Filters

Complimentary filters are two filters that sum to unity over all frequency. Complimentary filters can be used for sensor blending and offer some advantages over blending filters which are not complimentary.

A simple example of a complimentary filter set is a pair of first-order high and low-pass filters. If these two filters share the same pole frequency, they can be summed to unity (figure 6.4):

$$F_{low} = \frac{a}{s+a} \tag{6.1}$$

$$F_{high} = \frac{s}{s+a} \tag{6.2}$$

$$F_{low} + F_{high} = \frac{a}{s+a} + \frac{s}{s+a} = 1 \tag{6.3}$$

The advantage of using complimentary filters is that they do not contribute to the overall loop dynamics. Once two sensors are blended, the dynamics of the blending filters are not visible in the supersensor output. This becomes significant when higher order filters are used for the blending.

Normalized Complimentary Filters

It is often useful to normalize a sensor blending filter set with respect to each other. The advantages of this are that the resulting normalized filters are complimentary and the stability (in particular, the existence of non-minimum phase zeros) is made obvious during the division of the filter sum. This is especially useful as the blending filters become more sophisticated. Given two blending filters, F_1 and F_2, the

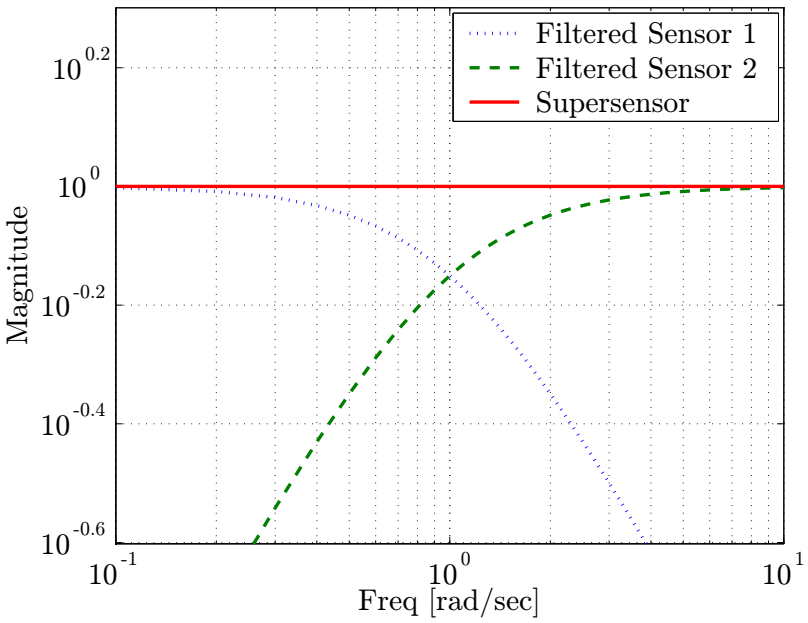

Figure 6.4: A simple example of complimentary filtering. Two sensors, sensor 1 and 2, have no dynamics and are filtered with F_{low} and F_{high}. The two filtered sensor outputs are combined to make the supersensor.

normalized filters are:

$$F_{1n} = \frac{F_1}{F_1 + F_2} \quad and \quad F_{2n} = \frac{F_2}{F_2 + F_1} \tag{6.4}$$

6.2.2 Sensor Correction

Sensor correction is a method of correcting the output a sensor by using information from another sensor. In this experiment, sensor correction is used to eliminate spurious ground motion from the displacement sensor.

The displacement sensor measures the difference between the position of the ground and the payload: $z_m - z_f$ (where z_f is the same as the ground position if the foundation stiffness is assumed infinite, figure 3.1). A controller attempting to null the displacement sensor output will minimize the difference between the payload and the ground position. As mentioned before, this causes the payload to follow the

ground and prevents isolation from the ground. However, position control is necessary to align the payload for locking the LIGO interferometer. A compromise to these conflicting requirements is to remove the ground position term, z_f, from the displacement sensor output in specific range of frequencies.

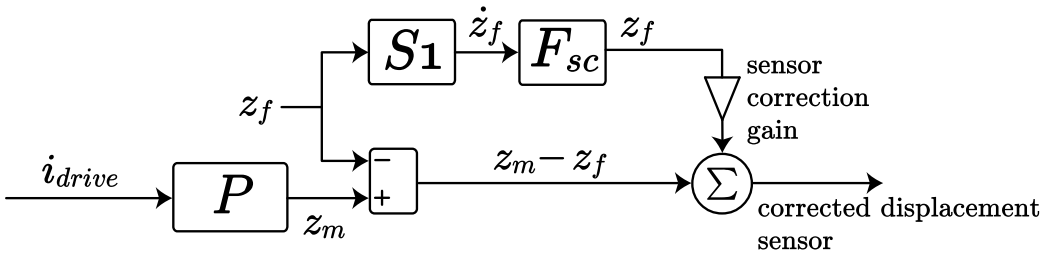

Figure 6.5: A block diagram of a displacement sensor corrected by a ground based seismometer. The seismometer, $S1$, measures the velocity of the ground, $\dot{z_f}$, which is filtered by F_{sc} to provide the ground position, z_f. The position output from F_{sc} is used to correct the displacement sensor output, $z_m - z_f$.

In order to accomplish this, it is necessary to obtain a measurement of the ground position. This can be done by using the output of a low-frequency seismometer stationed on the ground near the platform. Because of the nature of the instrument, the output of seismometer will either be the velocity or acceleration of the ground, but with some filtering, this can be integrated to provide the ground position over a band of frequencies. The position of the ground, z_f, is then added to the displacement sensor output, $z_m - z_f$, leaving only the payload position, z_m. Finally, the feedback loop acts to null the remaining output of the displacement sensor (z_m), and thereby, reduces the motion of the payload.

Since ground position information is applied directly to the plant, sensor correction is commonly mistaken for feedforward. The critical difference between sensor correction and feedforward is that no knowledge of the feedforward plant dynamics is necessary. This dramatically simplifies the implementation and improves the robustness (since plant dynamics often change with time) of sensor correction.

Sensor Correction Implementation

In order to implement sensor correction, the output of the ground stationed seismometer must be filtered (integrated) to provide the ground position. This signal must then be amplified by the sensor correction gain to drive the closed loop actuator with a magnitude that mimics the effect of the ground.

The task of the ground stationed seismometer filter, F_{sc} (figure 6.5), is to merely integrate the ground velocity, $\dot{z_f}$, to provide position, z_f. However, this simple objective is made difficult by the characteristics of the seismometer. Since seismometers are AC coupled devices, measurement of ground velocity is only available down to a specific frequency. For conventional seismometers, this frequency corresponds to the natural frequency of the instrument. However, the ground seismometer used in this experiment is the Streckheisen STS2 which is actively compensated, and hence, the low-frequency limit is set by the cutoff frequency built in to electronics of the instrument ($\sim$8 mHz). Therefore, the sensor correction filter must change sharply from an increasing slope below the seismometer cutoff frequency to a -1 slope (integrating) past the natural frequency.

There are sharp filters which would make this possible, but the affiliated disruption in phase is unacceptable for canceling the displacement sensor output of the ground position. Therefor, the sensor correction filter must compromise between rejecting noise below the seismometer natural frequency and integrating to provide position information above this frequency. An example of a sensor correction filter used in this experiment is shown in figure 6.6.

The sensor correction gain can be predicted if the sensitivities of the seismometer and the displacement sensor are well known. The sensor correction gain is equal to the ratio of the displacement sensitivity of the filtered seismometer output to the sensitivity of displacement sensor.

More commonly, the sensor correction gain is determined explicitly by taking the ratio of two transfer functions. For a simple platform (figure 6.7) the two transfer functions are:

$$H_1 = \frac{S_2}{S_1 * F_{sc}} \ \ and \ \ H_2 = \frac{S_2}{F} \tag{6.5}$$

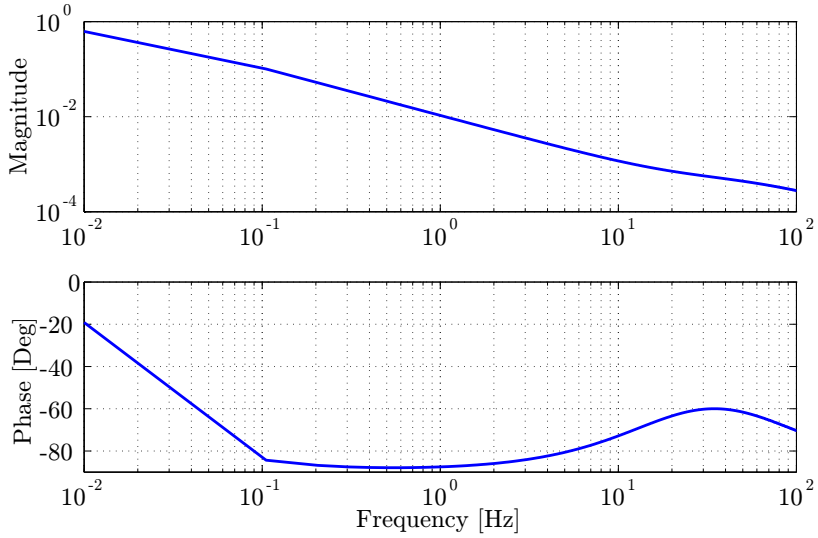

Figure 6.6: A filter typical of those used for sensor correction (F_{sc}). Note that between 0.1 and 5 hertz, the phase deviates less than 10 degrees from a prefect integrator (-90 degrees).

where $S_1 * F_{sc}$ is the filtered output of the ground stationed seismometer, S_2 is the output of the payload based seismometer, and F corresponds to the control input to the closed loop hydraulic actuator. Both of these transfer functions are taken with the actuator operating in closed loop (displacement sensor feedback). The ratio of these two transfer functions is:

$$\frac{H_2}{H_1} = \frac{S_1 * F_{sc}}{F} \tag{6.6}$$

This ratio should be flat over the range of frequencies where the filtered ground based seismometer, $S_1 * F_{sc}$, outputs ground position. The magnitude of this flat section is the desired sensor correction gain.

6.3 Control of the Quiet Hydraulic Test Platform

The control strategies discussed above were first implemented on the quiet hydraulic test platform (figure 4.10). The goal of the platform is to demonstrate the success of

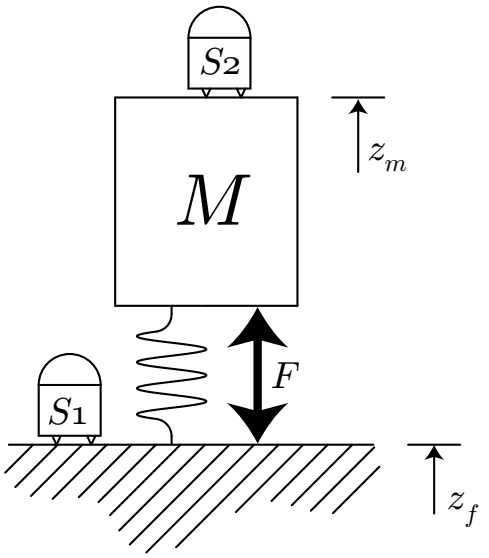

Figure 6.7: A simplified platform shown here to discuss the measurement of sensor correction gain. The force F is supplied by the hydraulic actuator, S_1 is a seismometer located on the ground and S_2 is a witness seismometer placed on the payload.

these strategies on a single actuator, and then, two actuators operating in concert.

6.3.1 Controller Design and Implementation

The design of the controllers evolved with the actuator, and notably with the introduction of the bypass network. All of the controllers include both sensor blending and sensor correction. These controllers are implemented on a real time controller in the commercially available Dspace environment.

Due to the orthogonal arrangement of the actuators in the test platform, there is little coupling between actuators, and the control can be implemented in a SISO fashion. The control diagram for each actuator (figure 6.8) includes both sensor correction and sensor blending. Within this diagram, there are four filters that must be developed: F_d, F_s, F_{sc} and F_c.

The blending filters, F_d and F_s, are developed along with the overall filter, F_c. The design of the blending filters is a compromise between emphasizing the displacement

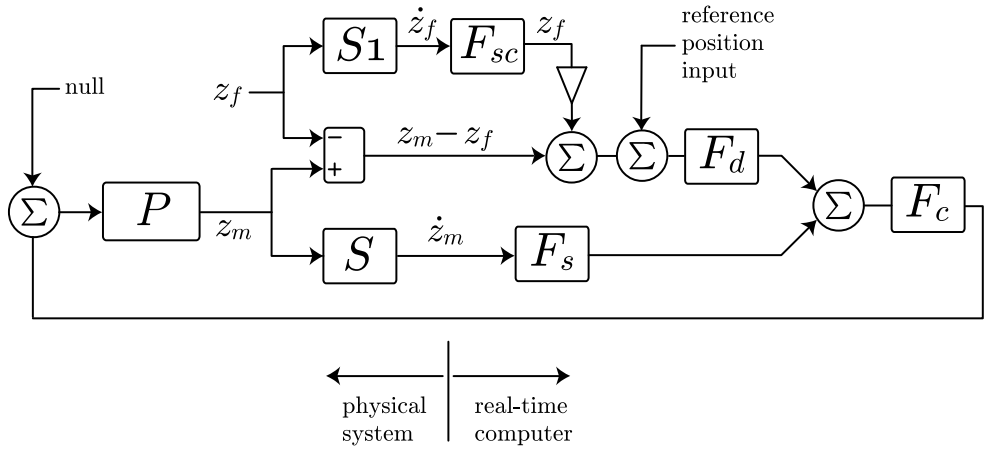

Figure 6.8: The closed loop system block diagram for a single actuator in the test platform. This diagram combines figures 6.2 and 6.5 with the addition of the overall controller, F_c.

sensor at low-frequencies and the seismometer at high-frequencies. For the first actuator prototype, the sensors involved are the Stanford designed optical displacement sensor and a Geospace manufactured HS-1 seismometer. Since the natural frequency of the HS-1 seismometer is 4.5 hertz, the blend frequency is set to 8 hertz (figure 6.9). Note that at the blend frequency, the filtered displacement sensor output is less than 180 degrees out of phase with the filtered seismometer output. The resulting supersensor is relatively continuous across the blend frequency.

The overall filter, F_c (figure 6.8), improves the phase margin of the supersensor near the unity gain frequency and increases the gain around 10 hertz. For the first actuator prototype, the overall filter also inverts the bellows' breathing resonance at 28 hertz. On the subsequent prototypes, the absence of the bellows' breathing resonance (figure 6.10) enables the overall filter to be less aggressive (more robust),

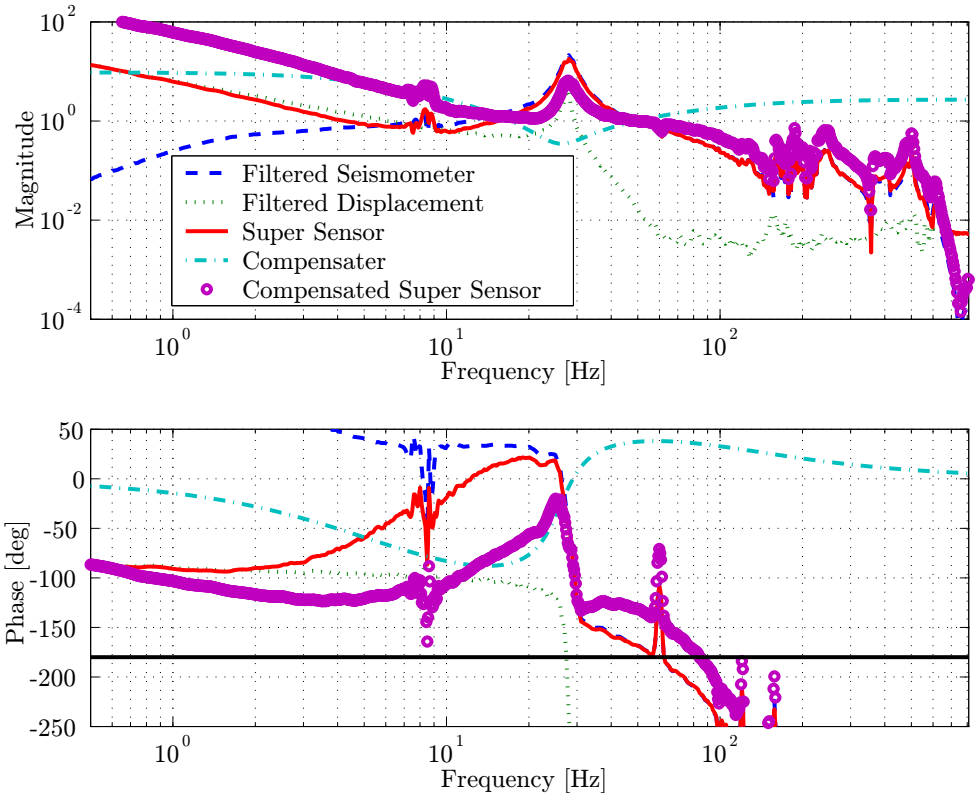

Figure 6.9: Blending filters and overall filter design for the first actuator prototype on the test platform. The compensator here is the same as overall filter, F_c, in figure 6.8.

and the loop gain is much greater across the bandwidth of the controller.

In both examples of sensor blending (figures 6.9 and 6.10), the displacement sensor blend filter is a simple gain. The dynamics of the displacement sensor output are unaltered because the sensor correction provides ample isolation over the frequencies where the displacement sensor output is prominent.

The sensor correction filter, F_{sc}, is based on inverting the dynamics of the seismometer used to measure ground motion as discussed in section 6.2.2. In all of the work presented here, the ground based seismometer is the Streckheisen manufactured STS2. The STS2 is a 3-axis, active feedback seismometer that is capable of providing inertial information down to 8 mHz. The filter used for sensor correction on the test

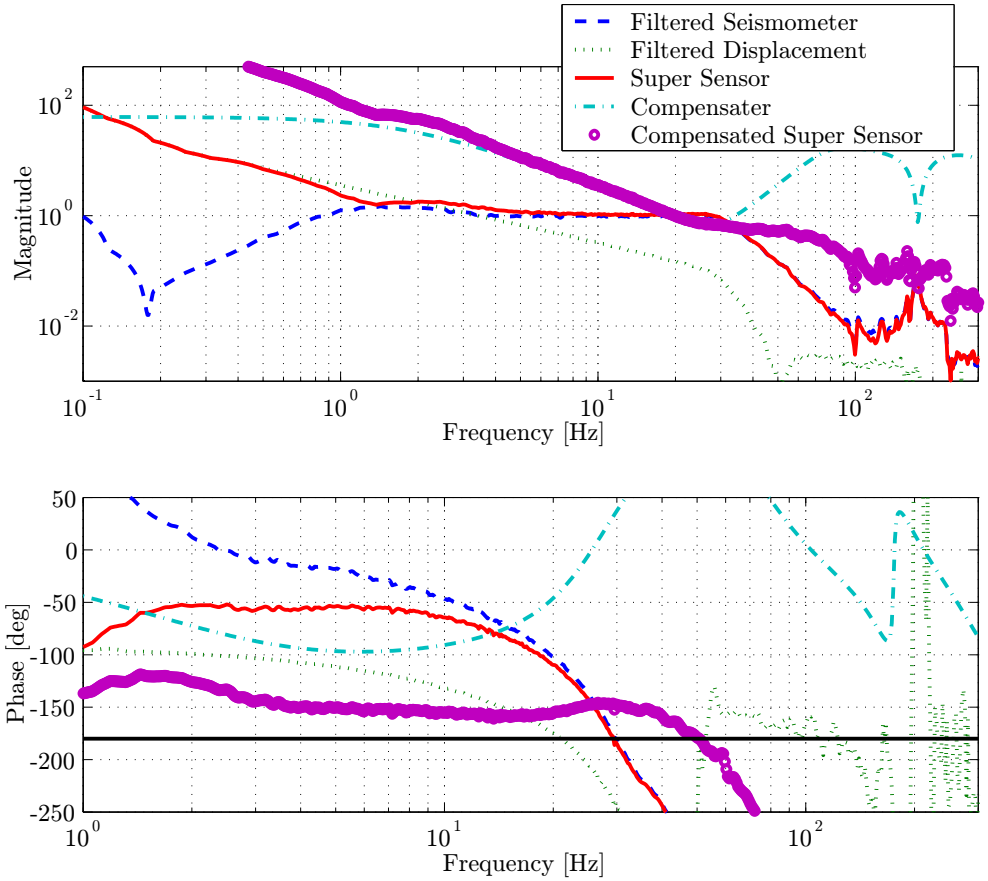

Figure 6.10: Blending filters and overall filter design for the final actuator prototype on the test platform. Note the absence of the bellows' breathing resonance and the increased low-frequency gain.

platform is similar to that shown in figure 6.6, and the output is useful for correcting the displacement sensor from 0.1 to 5 hertz. The sensor correction gain is determined by measuring the transfer functions H_1 and H_2 discussed in equation 6.5, and carrying out the division of equation 6.6.

For the purpose of measuring performance, there are two witness sensors installed directly inside of the payload. These witness sensors, Geotech manufactured S-13 seismometers, are external to the feedback loop. This is critical in order to obtain an unbiased measurement of performance since the loop suppresses the motion measured by the feedback seismometers. Since the total motion of the payload may be correlated

to the feedback seismometer output, the feedback seismometers can give an artificially promising measurement of performance.

6.3.2 Results from the Test Platform

The results shown in this section were generated by the final actuator prototype installed in the test platform. The modeled results are based on measured data of the plant and noise sources combined with the known controller transfer functions. All of the measured performance data is from the S-13 witness sensors embedded in the payload.

The feedback sensors and associated read-out electronics contribute noise which limits the performance of the actuator. In the test platform installation, the feedback sensors are the DIT-5200 displacement sensor and a 2 hertz natural frequency HS-1 seismometer.

Another source of spurious noise is the analog to digital interface (A/D and D/A) that separates the realtime digital controllers from the physical system. The noise from these processes is based on the number of bits in the converter and the amount that the signal is oversampled. The smallest level that can be resolved by the converter is called the quantizer step size, Δ [28]:

$$\Delta = \frac{X_m}{2^B} \tag{6.7}$$

where X_m is the input (or output) range of the converter and $B+1$ is the number of active bits in the converter. The number of active bits refers to the number of useful bits that the converter can practically resolve. This is typically less than than the manufacturer's specification. The amplitude spectral density of the noise introduced by quantizing the continuous signal is [28]:

$$N_{cd} = \Delta\sqrt{\frac{2T}{12}} \tag{6.8}$$

where T is the sample period. In the case of the Dspace manufactured realtime computer attached to the test platform, measurements undertaken in the lab indicate

that the A/D converters are 16-bit with 12 active bits, and the D/A converters are 14-bit wherein 12 are active. The input and output range (X_m) is ±10 volts, and the sampling rate used for the results shown here is 5000 samples per second.

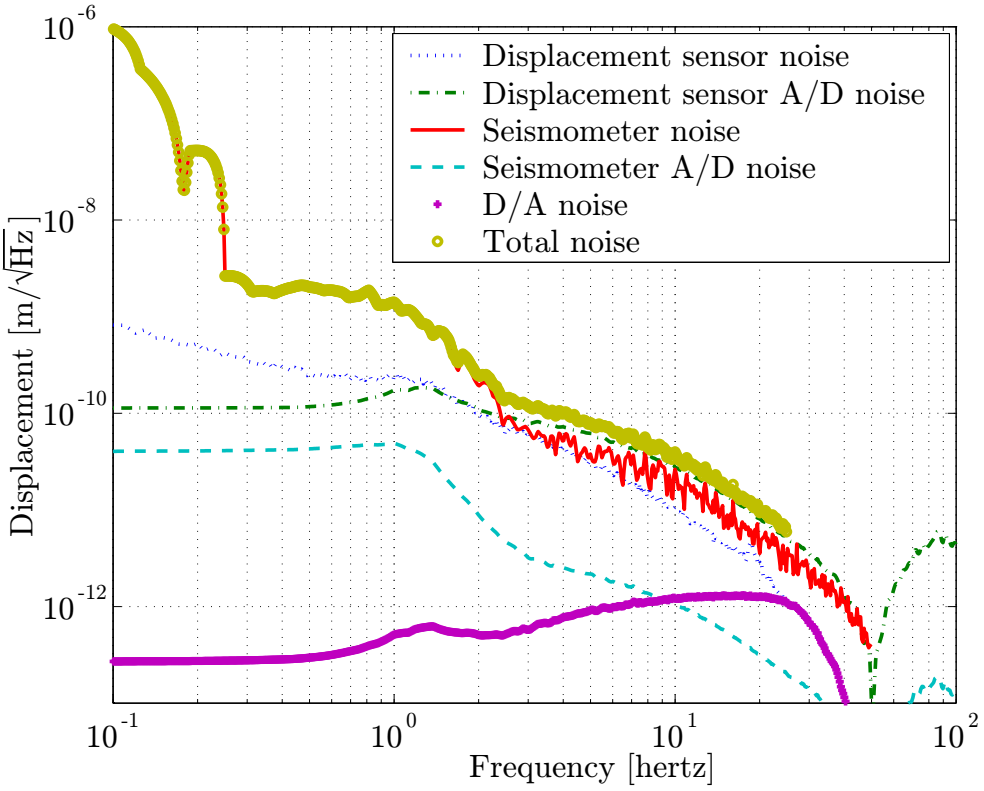

Figure 6.11: Amplitude spectral densities for the known noise sources of the test platform.

The closed loop amplitude spectral densities (figure 6.11) of the noise sources exhibit some interesting features. The A/D noise for the seismometer and the displacement sensor are not flat because of the loop shaping in the controller and are not the same because the blending filters are unique to each sensor. The noise from the seismometer increases dramatically at low-frequency both because of the electronics noise from the sensor readout board and the dynamics of the sensor. The total noise in the system is dominated by the seismometer noise. This suggests that the performance of the system would be diminished if the seismometer was further emphasized

in the sensor blending.

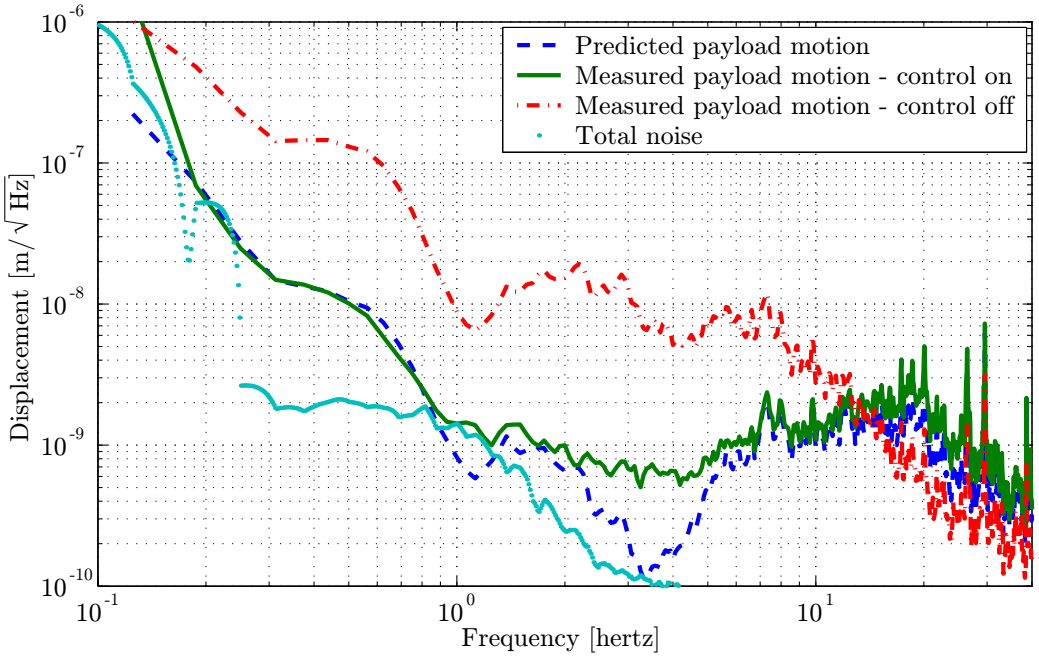

Figure 6.12: The measured and predicted displacement performance of the final actuator prototype mounted on the test platform.

The isolation performance of the system with sensor correction (figures 6.12 and 6.13) show that the actuator is capable of meeting the isolation performance specification for LIGO. Also evident in the figures are the limitations to performance which vary with frequency.

At frequencies below 1.0 hertz, the isolation performance is derived entirely from sensor correction. The success of the sensor correction is determined by the ability of the ground based seismometer, the STS2, and the sensor correction filter to provide accurate information about ground position. The STS2 can provide velocity information down to frequencies approaching 8 mHz., but at these low frequencies, it is difficult to develop a sensor correction filter to purely integrate the velocity output without distortion. The filter used for the test platform (figure 6.6) only resembles an integrator (phase of -90 degrees) in the band between 0.1 and 2.0 Hz, and even within this band, there is measurable ($\sim$ 3-7 degrees) phase distortion. This phase

distortion causes matching errors on the order of 10% thus limiting the isolation performance to a factor of 10 reduction in the 0.1 and 2.0 hertz band. Outside of this band, the isolation performance from sensor correction decreases with the increasing phase distortion.

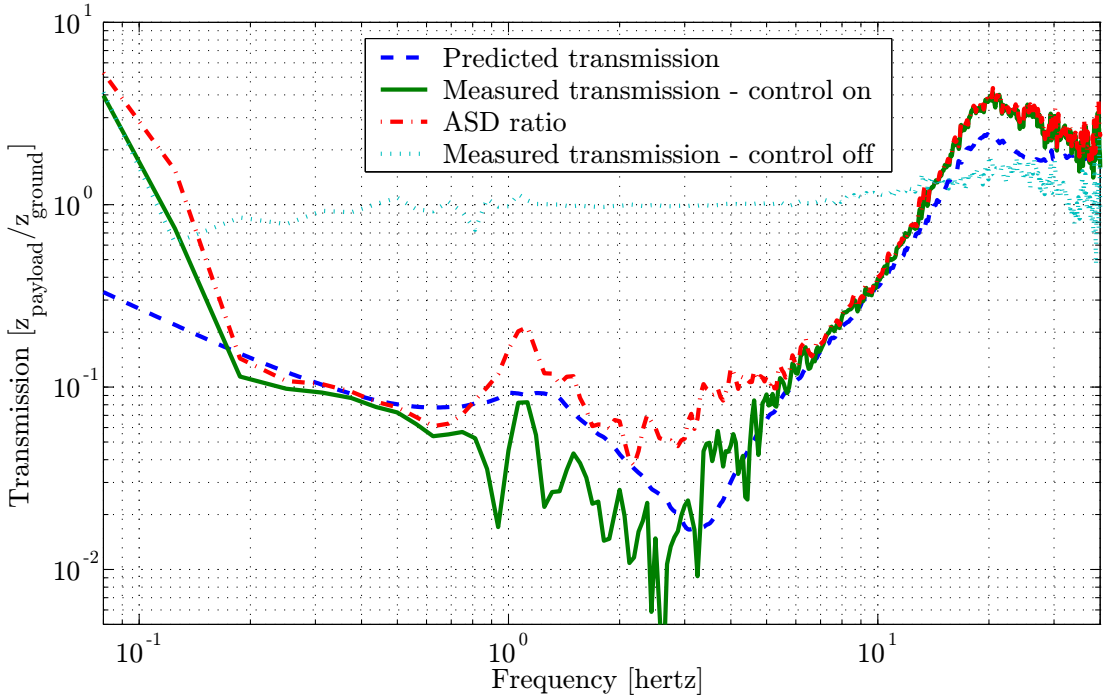

Figure 6.13: The measured and predicted transmission performance of the final actuator prototype mounted on the test platform. The measured transmission is the transfer function between the payload mounted S13 and the ground base STS2 seismometer. The ASD ratio is the ratio of the amplitude spectral densities for the S13 and the STS2. The discrepancy between the measured transmission and the ASD ratio is the uncorrelated noise between the S13 and the STS2.

The performance at frequencies around 1.0 Hz is limited by sensor noise. In this narrow region, the total noise (largely due to the HS1) passes above the predicted performance estimate and meets with the measured performance. This should improve with a more sensitive (lower natural frequency and larger generator coefficient) seismometer.

Above 2.0 hertz, the isolation performance comes from feedback to the HS1 seismometer. In this region, the performance is dependent on the feedback loop gain, and therefor, the prediction of performance matches well with the measured performance. However, in the 2 to 5 hertz band, the measured performance deviates considerably from the prediction. This is in contrast to figure 6.13 where the prediction follows the measured transmission more closely. The discrepancy between these two measurements of performance, the amplitude spectral density (figure 6.12) and the transmission (figure 6.13), is the measurement method. The amplitude spectral density measurements are based on the individual outputs of the ground and witness sensors whereas the transmission measurement compares these two outputs simultaneously. During this simultaneous comparison, the signal analyzer, a Stanford Research Systems SR785, removes noise that is uncorrelated between the ground and the payload. Hence, the transmission measurement is falsely optimistic since the goal of the platform is reduce the total motion; not merely motion correlated to the ground. This excess of noise is not surprising since the payload is not in a vacuum, and is therefor exposed to acoustic coupling and air currents.

6.3.3 Test Platform Results in the Horizontal Direction

In the horizontal direction, the isolation performance is not as impressive as in the vertical direction. This difference stems from tilt-horizontal coupling, and the lack of feedback gain caused by the absence of a bypass network in the horizontal actuator (first actuator prototype). However, the results figure 6.14 demonstrate that a vertical and horizontal actuator can operate in concert and significant isolation performance can be realized with the original actuator design. The improvement of the spring stiffness tilt correction discussed in section 4.3.2 is evident in figure 6.14.

6.4 The LASTI Installation

Once the performance of the actuator had been verified on the test platform, eight actuators were manufactured and installed on the BSC tank (figure 4.14) at the LIGO

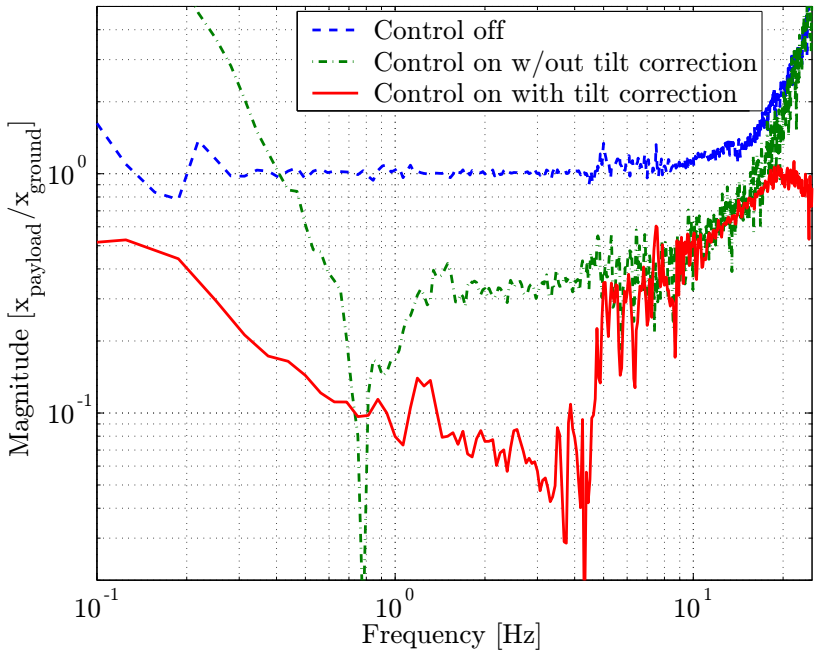

Figure 6.14: Measured performance of the first actuator prototype in the horizontal direction from supersensor feedback and sensor correction. The vertical performance was unaffected by the horizontal loops.

Advanced Science Testing Installation (LASTI) located at MIT. The BSC chamber and payload at LASTI are identical to those in use at the LIGO observatories, and the purpose of the LASTI installation was to demonstrate the isolation performance of the hydraulic external pre-isolator (HEPI) system in a full scale installation. The HEPI installation at LASTI was the final demonstration of the actuator prior to installation in the LIGO observatories.

The BSC installation offers several interesting control challenges. Surprisingly, the least troublesome of these is the overdetermined arrangement of the platform involving eight actuators for six degrees of freedom. The overdetermined control is facilitated by the compliance in the platform and payload. However, this compliance combined with the considerable size and mass of the system presents the greatest challenges for the control implementation.

Based on the experience with the test platform, the Sercel manufactured L4C

is used for the feedback seismometers in the BSC installation. The L4C exhibits substantially better noise performance by comparison to the HS1 thereby alleviating any concerns of sensor noise limitations for the HEPI system.

6.4.1 Controller Design and Implementation for the LASTI BSC

The control architecture is similar to that used for the test platform, but generalized to a multi-degree of freedom system. Sensor blending and sensor correction are applied to each actuator individually, but once combined, the resulting supersensor outputs are applied to modal directions involving groups of actuators. The result is a multi-input, multi-output (MIMO) controller that is diagonalized in the chosen modal directions.

Within each of the modal directions, groups of four actuators are operated in concert. Each pier supports both a vertical and a tangentially oriented (with respect to the BSC chamber) actuator. Thus, for each BSC chamber, there are two groups of four actuators: four oriented vertically, and four oriented tangentially (shown in detail in chapter 4).

The controlled modes of the system (figure 6.16) are chosen to coincide with the directions stipulated in the LIGO performance specifications. This is important in order to guarantee performance in the orientations that are critical to the LIGO instrument. The controller is diagonalized along these modal directions because there is little cross-coupling between modes. This can be attributed to the flexibility of the crossbeams and optics table (inside of the vacuum tank) which connect the four corners of the system together. The highest cross-coupling exists between horizontal actuators in the Rz mode, but even within this mode, the magnitude of the collocated sensor output is almost a factor of ten larger than the next uncollocated sensor (figure 6.15). This observation led to a diagonalized controller that ignores cross-coupling between modes.

Each of the eight actuators are equipped with a collocated displacement sensor and seismometer making 16 feedback sensor outputs. There are three outputs from the ground stationed STS2 seismometer for sensor correction of x, y and z, and

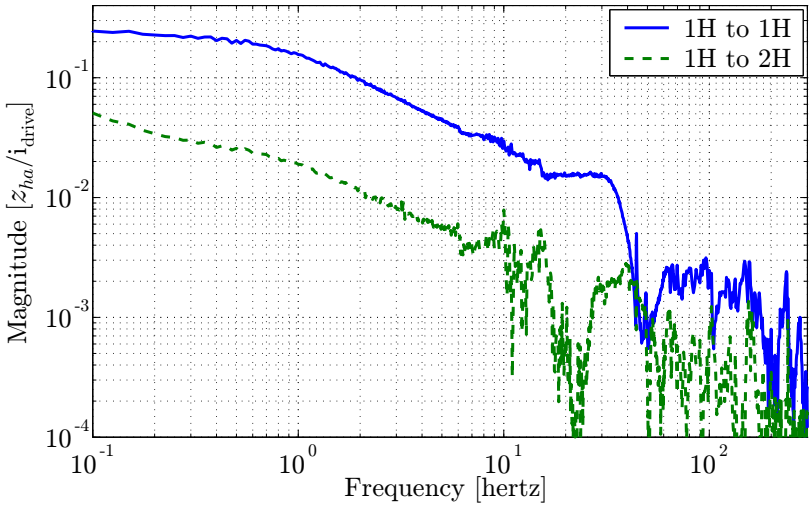

Figure 6.15: Transfer function measurements of cross coupling between horizontal displacement sensors. 1H to 1H is the horizontal displacement sensor output from the horizontal actuator on pier 1, and 1H to 2H is the horizontal displacement sensor output on pier 2 from the horizontal actuator on pier 1.

witness sensors on the crossbeams and optical table (inside the vacuum chamber) for out-of-loop isolation performance measurements.

After these outputs are sampled by the A/D converters, the first step of the control is to apply sensor correction to the displacement sensor outputs. This begins with filtering the outputs of the ground stationed STS2 seismometer in order to obtain measurements of ground position. Since all three axes of the STS2 are used, the sensor correction filter F_{sc} (figure 6.18) becomes a 3x3 matrix of transfer functions where the sensor correction filters populate the diagonal and all other elements of the matrix are zero.

Following this, the sensor outputs are separated into vertical and horizontal sets of four. In the vertical direction, the four displacement sensor outputs are filtered by the vertical displacement sensor filters, F_{dV}, and combined with the vertical, filtered seismometer outputs to make four vertical supersensor outputs. The blending filters applied here are similar to those shown in figure 6.3, and generally more aggressive than those used on the test platform. The use of more aggressive blending filters

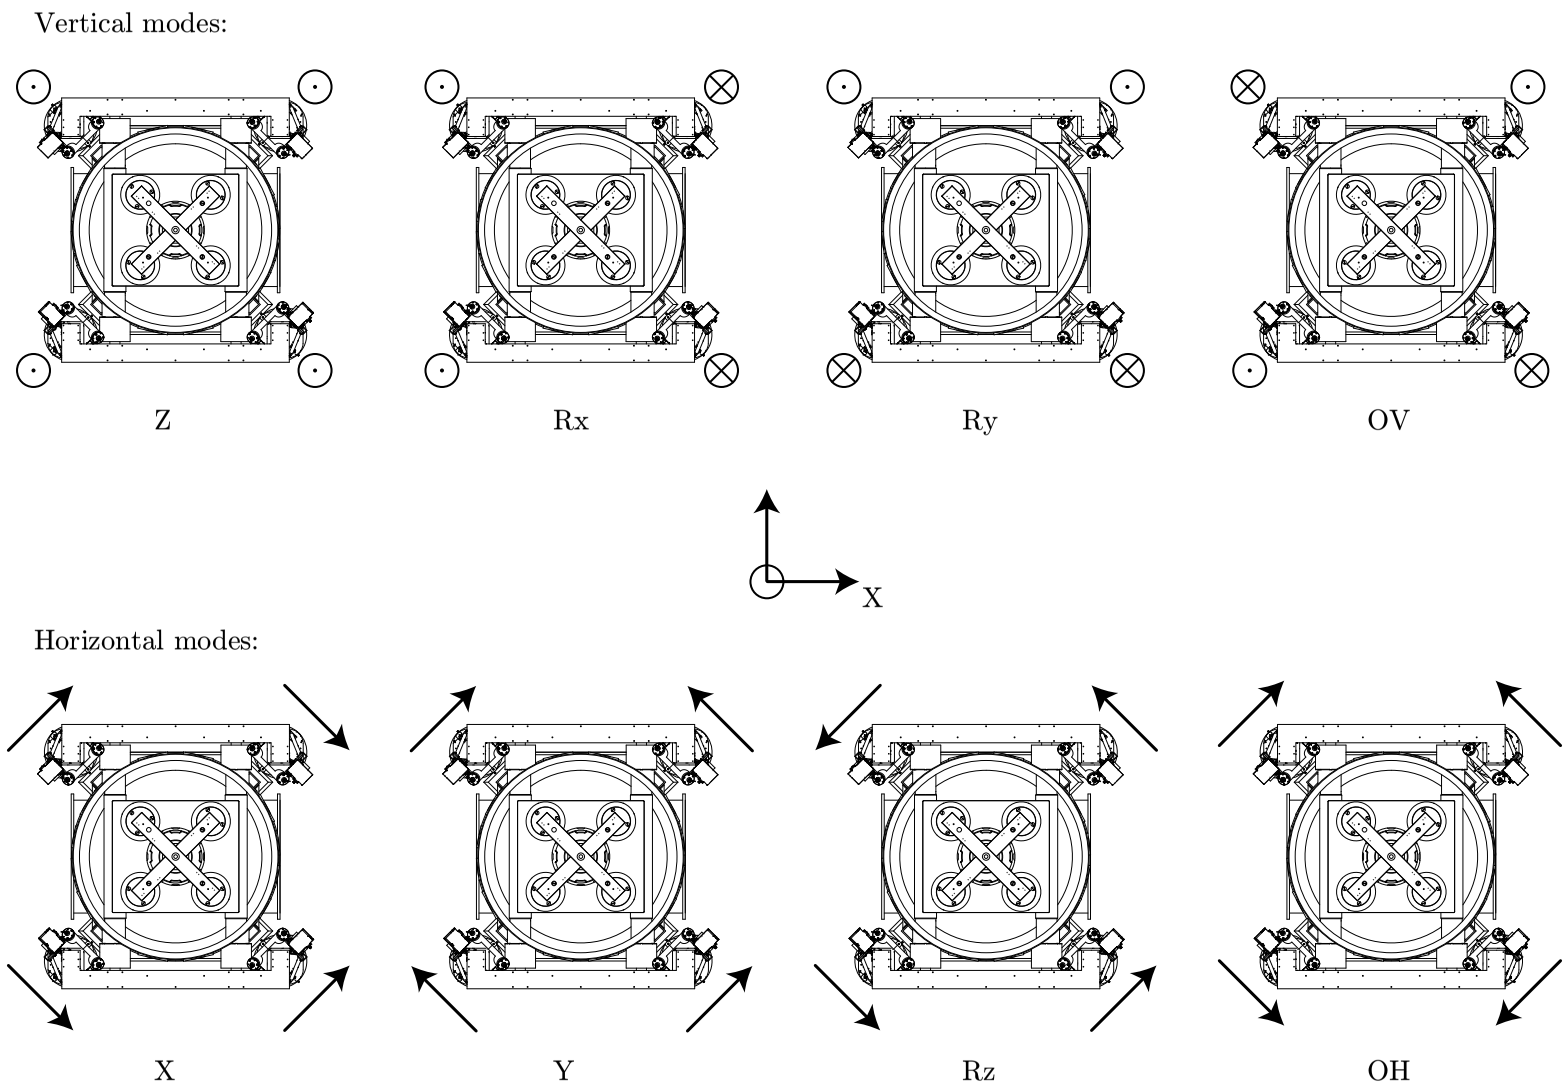

Figure 6.16: The controlled modes of the BSC platform. There are two overconstrained modes, OV in the vertical direction and OH in the horizontal direction.

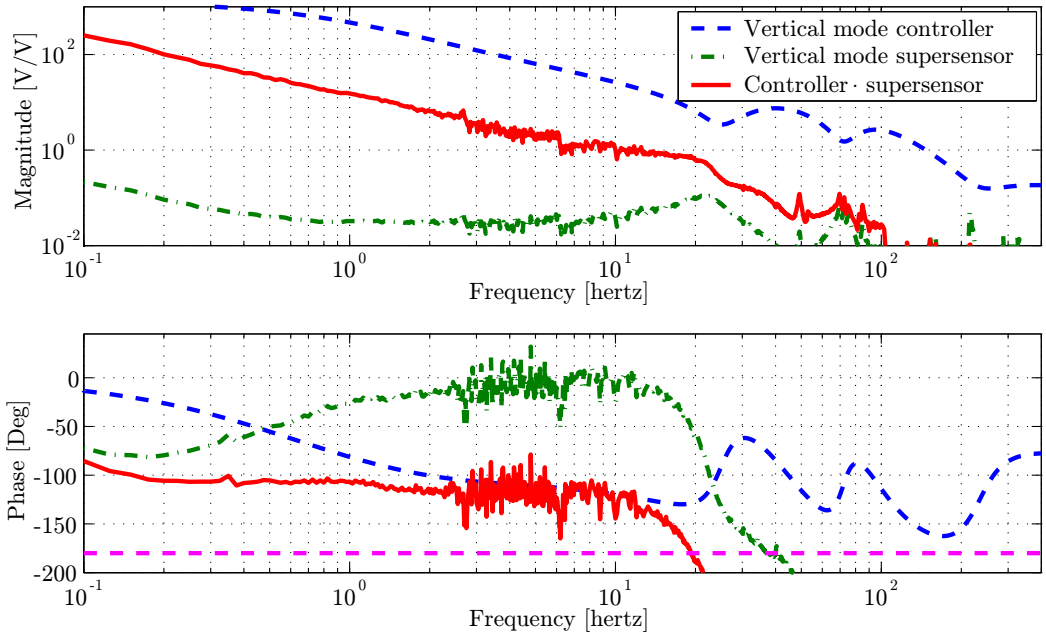

Figure 6.17: Transfer functions of the Z mode controller and supersensor.

is motivated by the higher quality feedback seismometers used on the HEPI system. With a lower natural frequency and an improved signal to noise ratio, the objective of the blending filters shifts to further emphasize the seismometer at lower frequencies (even below 1 hertz), and to dramatically suppress the displacement sensor above 10 hertz (where the sensor correction is no longer effective).

The four vertical supersensor outputs are then projected onto the Z mode direction (pure vertical) by the convert to modal coordinates matrix gain block (figure 6.18). In this orientation, a SISO controller is applied to this modal supersensor (figure 6.17). The resulting control signal is rotated back into the actuator basis and added to the Rx, Ry and OV control signals before being sent to the vertical actuators.

The pitch, yaw and vertical overdetermined modes (Rx, Ry and OV) are controlled separately by a set of simple displacement feedback controllers contained in the F_{cV} transfer function matrix block. Again, this block is populated only along the diagonal since cross-coupling between modes is ignored. Displacement feedback control alone is sufficient for the rotational modes since the ground motion exhibits very little local

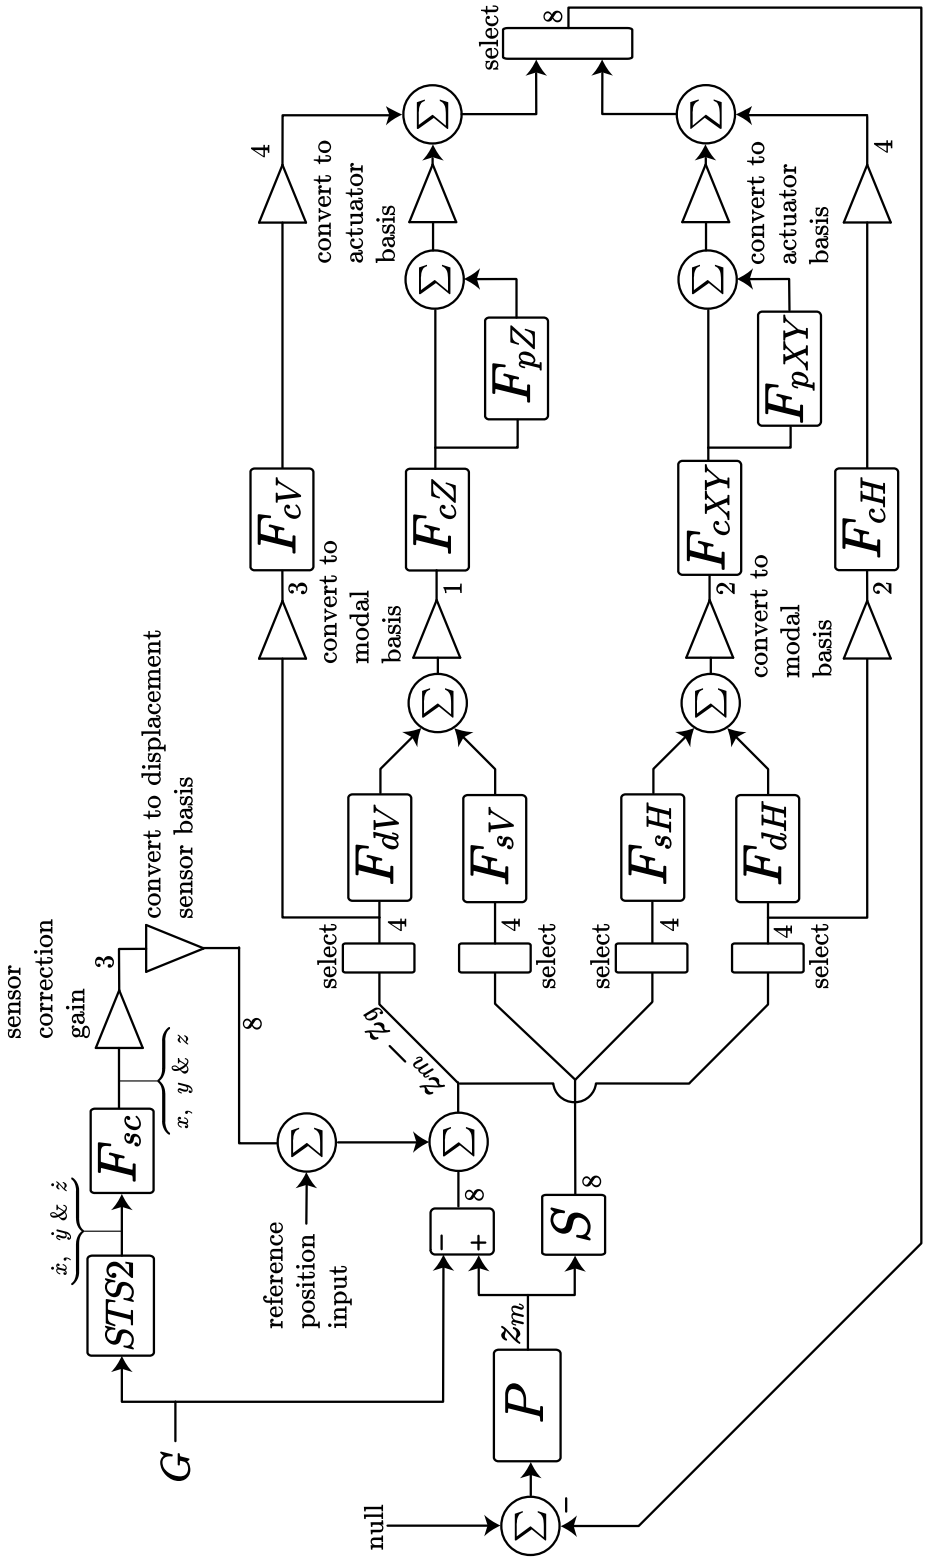

Figure 6.18: The block diagram for the eight actuator HEPI installation at LASTI.

rotational motion.

Similarly, the vertical overdetermined mode is controlled with displacement feedback. In circumstances where the payload and platform are stiff, applying control effort to an overdetermined mode could be detrimental, but the HEPI payload and support piers are sufficiently compliant to alleviate these concerns.

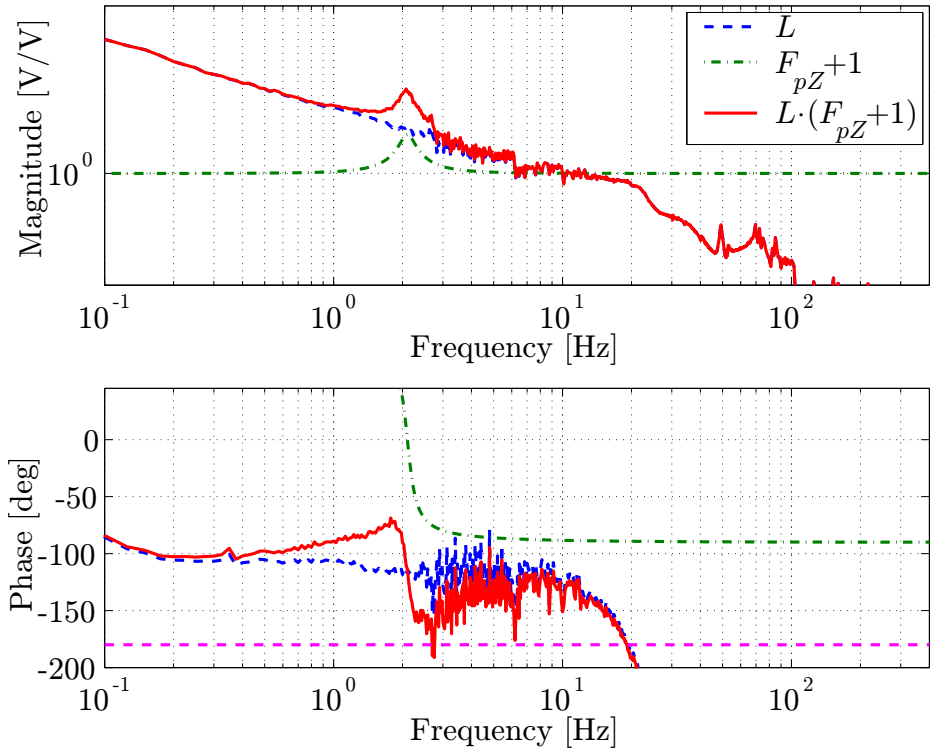

Figure 6.19: Transfer functions of the Z mode controller with the conditionally stable gain peak loop. L is the open loop compensated plant (from figure 6.18, $P(SF_{sV} + F_{dV})F_{cZ}$), and F_{pZ} is the additive filter. The additive path containing F_{pZ} is only engaged after the main loop is closed. Once engaged, the effect on the open loop compensated plant is represented by multiplying by $(1 + F_{pZ})$.

In addition to providing isolation from the ground at low-frequencies, another goal of the HEPI system is to damp modes of the passive isolation stack. The passive isolation stack for LIGO I consists of several masses separated by internally damped compression springs. Despite the internal damping intended to suppress the modes of the stack, the are several mass/spring resonances that are easily excited and disruptive

to the interferometer optic. In order to ensure that these stack modes are sufficiently suppressed, gain is added to the feedback loop at the resonances of the stack (figure 6.19). These gain peak loops, F_{pZ} and F_{pXY} in figure 6.18, are conditionally stable meaning that they can only be closed after the main feedback loop has already been closed.

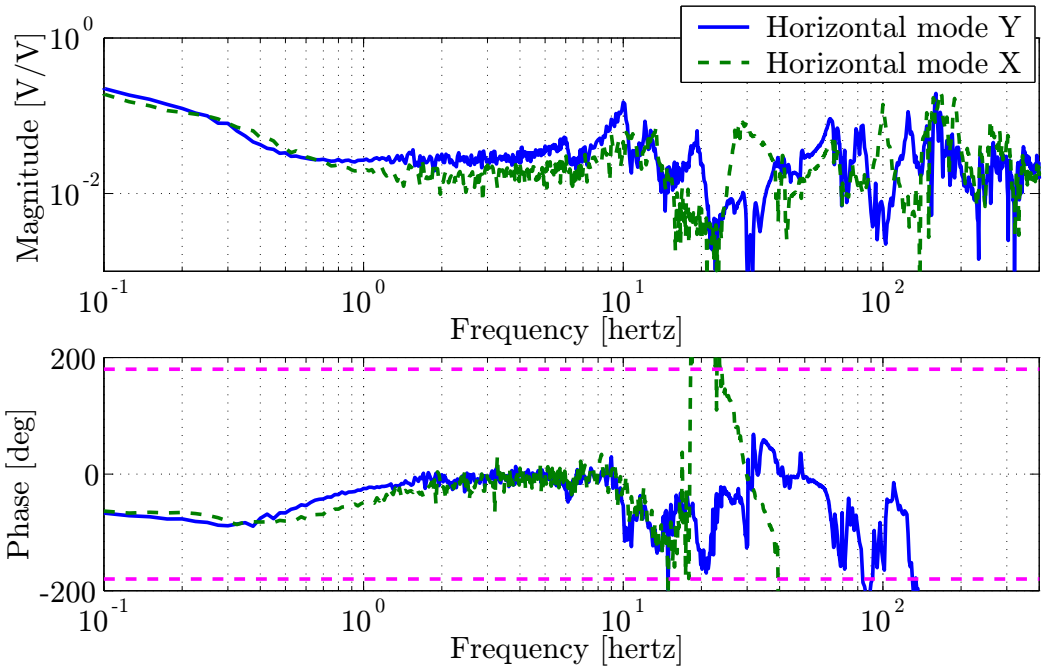

Figure 6.20: Transfer functions of the horizontal blended supersensor outputs projected onto the X and Y modes.

In the horizontal direction, the performance requirements apply to both the X and Y modes, and to meet these requirements, these two modes are configured the same as the Z direction. However, for the LASTI installation, it is prohibitively difficult to incorporate the feedback seismometers into the control loop. The complication results from substantial compliance separating the horizontal actuators from the horizontal seismometers coupled with platform/pier compliance that is pronounced in the horizontal direction. This compliance is responsible for the plethora of modes populating the X and Y direction transfer functions above 10 hertz (figure 6.20). While it is conceivable to control each of these modes individually, the resulting controller would

be very sensitive to any perturbations in the plant, and therefore, utterly lacking the requisite robustness.[1]

In the absence of useful seismometer output, the horizontal feedback control is based entirely on the horizontal displacement sensors. In the X and Y directions, isolation performance is obtained by sensor correction, while the yaw and overdetermined horizontal modes (Rz and OH) are controlled with uncorrected displacement sensor loops.

6.4.2 Results from the HEPI Installation at LASTI

The purpose of the LASTI installation is to demonstrate the functionality and isolation performance of the hydraulic actuator in the HEPI system. While the functionality of the actuator was nearly flawless, the isolation performance was challenging to obtain. Some of the obstacles to performance have already been discussed in the preceding section, but the LASTI installation, and the environment surrounding it, present some interesting hurdles that made the isolation performance difficult to measure.

A significant problem with the HEPI installation at LASTI stems from the compliance of the floor in the laboratory. On a compliant floor it is possible to unintentionally create a positive feedback loop between the payload and the ground based seismometer, the STS2. This occurs when the 2000 kg HEPI payload undergoes a large magnitude ($\sim$ 0.5 mm) horizontal translation with reasonable speed (even as slow as 1 mHz. peak to peak). This movement of mass causes the ground beneath the piers to deform and the floor beneath the STS2 to tilt. The STS2 is extremely sensitive to tilt. Owing to tilt-horizontal coupling, the STS2 perceives this tilt as horizontal motion. The horizontal motion measured by the STS2 is applied to the payload through sensor correction worsening the situation and causing the system to enter into an oscillation limited by the maximum velocity of the actuator (80 μm/s).

This problem is exacerbated by LASTI's location. The LASTI laboratory is located on the MIT campus in the center of Cambridge Massachusetts, and due to

[1]This problem is specific to the LASTI installation. Subsequent installations of HEPI in LIGO feature better collocation between the horizontal actuators and feedback seismometers.

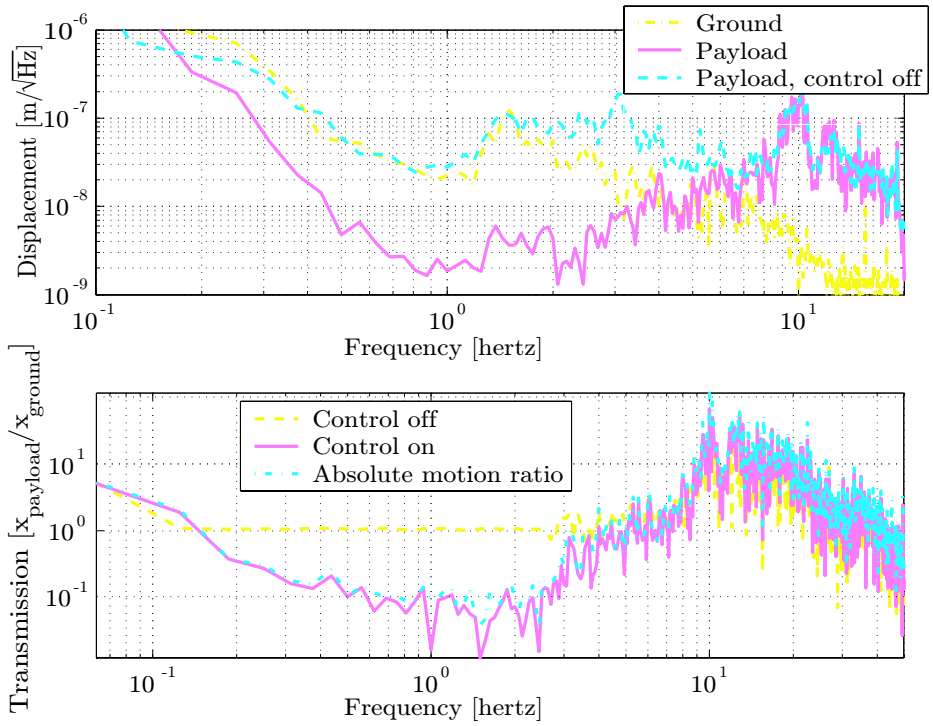

Figure 6.21: The isolation performance of the HEPI installation at LASTI in the X direction. The upper plot is an amplitude spectral density measurement of the ground and the payload motion while the lower plot shows the transmission to the payload from the ground. Note that the performance at frequencies below 0.15 Hz is limited by a modified sensor correction filter that suppresses tilt motion, and above ∼7 Hz, the lack of performance can be attributed to the absence of useful horizontal feedback seismometers.

the urban location, the ground motion at the lab is considerably higher than that experienced at the LIGO sites. Furthermore, the ground motion at the lab cannot be assumed to be stationary because of frequent bursts of large magnitude motion caused by nearby trucks, subway trains and pile driving. These large magnitude bursts are often accompanied by low-frequency tilt that initiates the horizontal displacement required to couple the payload motion to the sensor correction. Once this cycle is initiated, the sensor correction must be turned off before system can operate properly.

Measurements of performance were encumbered by the combined effects of payload motion coupling into the sensor correction and the lack of stationarity in the ground

motion. In order to maintain stability against tilt motion, the poles of the X and Y mode sensor correction filters were moved from 0.01 Hz to 0.05 Hz. This reduces the system's sensitivity to floor tilt at the cost of abandoning low-frequency sensor correction performance. The performance in the horizontal direction (figure 6.21) reflects both the poor sensor correction performance at low frequencies and the lack of useful feedback seismometer output at higher frequencies.

While the vertical performance for the HEPI installation at LASTI approached the performance observed on the test platform, the isolation performance in the horizontal directions is unimpressive. Unfortunately, due to the extent of the modifications necessary, the verification of performance in the horizontal directions had to wait for the HEPI installation in a LIGO site.

6.5 HEPI at the LIGO Livingston Observatory

At this time several HEPI systems have been commissioned at the LIGO observatory in Livingston Louisiana (LLO). The commissioning and control of these systems has been carried out by a team led by Professor Joe Giaime. The group led by Giaime have accomplished some impressive results which are included here to show the performance of the HEPI system with improvements motivated by the lessons learned at LASTI.

Most notable among the improvements at LLO is the stability of the floor which supports the HEPI installations. This floor is almost 8 time as thick as the floor at LASTI, and thus far, no coupling between payload motion and the STS2 has been observed. However, the absence of coupling may also be attributed to the use of an improved sensor correction filter. The sensor correction filter in use at LLO is an FIR filter developed by Wensheng Hua that integrates nearly perfectly down to 0.1 Hz. before abruptly switching to a positive slope (derivative) to reject low-frequency tilt motion.

Figure 6.22 shows the performance in the X direction at LLO. Note that the microseismic peak at $\sim$0.15 Hz. is reduced by a factor of 10 and this reduction continues up to 2 Hz.

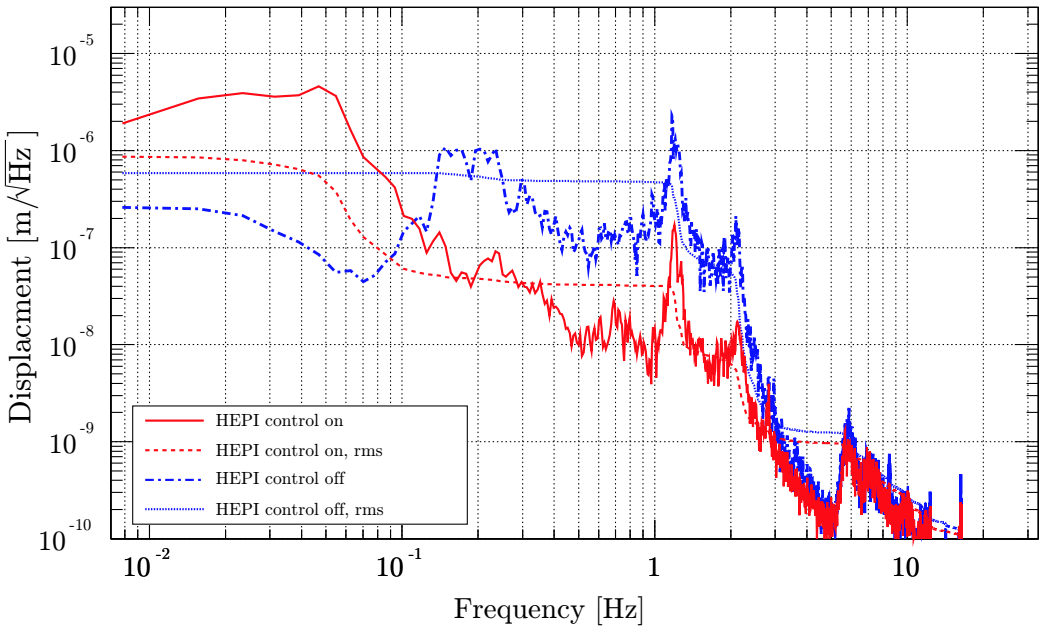

Figure 6.22: Transfer functions of the Z mode controller and supersensor. Courtesy of Professor Joe Giaime.

6.6 Summary

The control of the hydraulic actuator evolved with the design of the actuator, and in the case of the HEPI system, with the development of the platform. In both situations, the performance improved with experience of working with the actuator and the surrounding systems.

The performance of the test platform exceeds that observed on any of the HEPI BSC installations. This can be attributed to the relatively high foundation stiffness of the test platform made possible by its small size. In addition, the test platform payload is dramatically less complex, and the stiffness between the payload and the actuator is much higher. While these are the foremost causes for the performance discrepancy between the two systems, it may be possible to improve the performance of HEPI by incorporating the ideas of section 4.3 into the HEPI housing. This and other improvements will be explored in the following chapter.

The recent introduction of HEPI at LLO has already been recognized as a considerable improvement to the functionality of the observatory. Prior to the commissioning of HEPI, it was impossible for the interferometer to maintain lock during the daytime due to nearby logging activities. This meant that the observatory was only capable of producing useful data during the nighttime (in the absence of trains). After the installation of HEPI, it has been possible to maintain lock during the day even through the occasional passing freight train. As a result, LLO has moved from operating 23 percent of the time (during science run 1) to over 90 percent (recently, in science run 3).

Chapter 7

Future Work

This chapter presents design alternatives to the core components of the Hydraulic External Pre-Isolator (HEPI) system. A new actuator housing is proposed that offers improvements over the existing housing design both in facilitating adjustments/installation and improving overall stiffness. And, an alternative actuator concept is put forth for improving the range and stiffness of the existing actuator. While the proposed housing design could readily be built and installed, the new actuator will require development and experimental verification before it would be qualified for use in LIGO.

7.1 Actuator Housing

The EPI system is comprised of vertical and horizontal actuator pairs mounted on top of the piers that surround the BSC (figure 4.14) and the HAM vacuum chambers (figure 4.15). The purpose of the actuator housing is to provide mounting and alignment of these actuator pairs. The housing must also support the offload springs and some sensors. The principal design objectives for the housing are: stiff connections between sensors and actuators; adjustment of the EPI system position to center the optic ($\sim$5 mm); adjustment between the actuator pair and the vacuum chamber crossbeam in 6 DOF that results in a stiff connection (since, experience indicates, the crossbeam is never aligned in any DOF with respect to the pier top); and ease of

installation of both the housing and individual actuators.

7.1.1 The As-Built HEPI Actuator Housing

The actuator housing currently installed in LIGO was discussed in chapter 4, and is shown again here in figure 7.1. The housing is made up of three constituents: the base, bridge and interface assemblies. The base assembly is a four sided steel box that supports the foundation end of the offload springs and incorporates both actuators. The stiffness bars add to the stiffness of the base assembly in the horizontal direction. The bridge assembly also adds stiffness to the base assembly, and provides locking bolts to locate the suspended portion of the platform. The interface assembly attaches to the crossbeam of the vacuum tank and supports the suspended end of the offload springs.

The actuator housing is installed by attaching part of the interface assembly to the vacuum chamber crossbeam. Since the crossbeam is typically misaligned in multiple DOF with respect to the pier top, it is necessary to add shims between the interface assembly and the crossbeam. Once the upper portion of the interface assembly has been attached to the crossbeam and shimmed into alignment with the pier, the remainder of the housing, including the actuators, is put into place.

Once the housing is in place and the payload balanced on the offload springs, it remains to align and attach the actuators to the interface assembly. In order to compensate for the misalignment between the actuators and the interface assembly, the base assembly is designed with oversized fastener holes and the actuators have several adjustable shim plates (shown attached to the actuators in figure 7.1).

This adjustment, both at the interface assembly and the actuators, is more than sufficient to align the system and the actuators. However, the process for installation is labor intensive and the shim plates do not always result in a stiff connection even when they are clamped. Furthermore, it is especially difficult to install the vertical actuator because it is designed to fit below the base plate of the housing base assembly. Therefore, the vertical actuator must be suspended in place while it is attached to the underside of the interface assembly (difficult for even a particularly strong individual

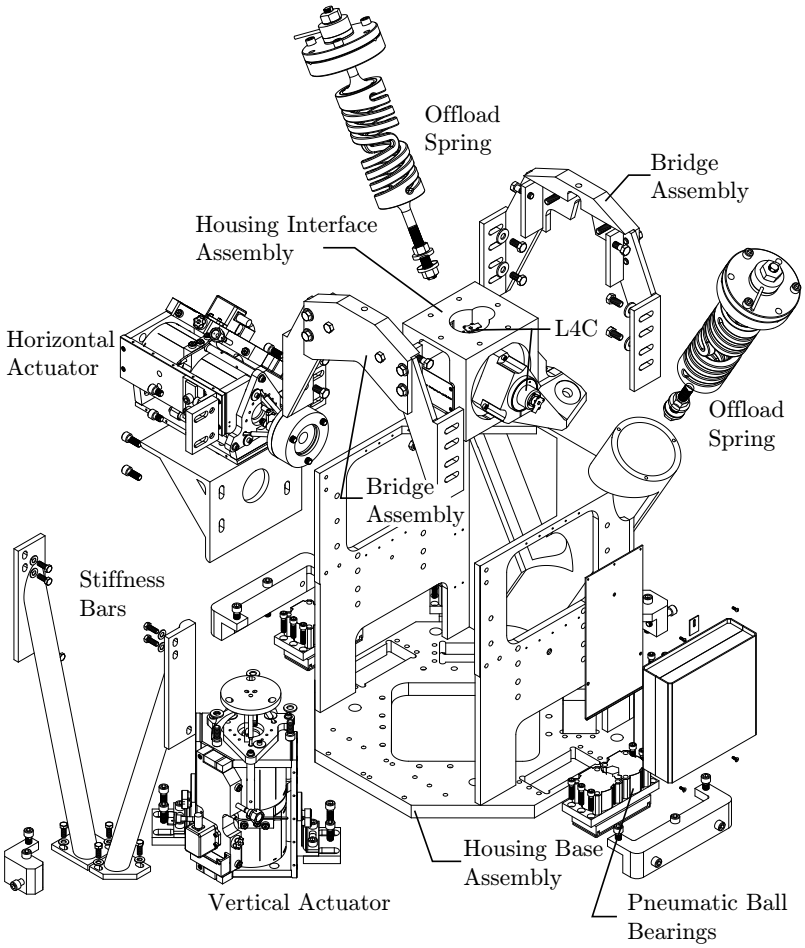

Figure 7.1: An exploded view of the as-built actuator housing design. Courtesy of Oddvar Spjeld.

since the actuator has a mass of nearly 40 kg).

With regard to stiffness, the housing is intended to provide a stiff foundation for the actuators and stiff connections between sensors and actuators. The base housing assembly has many features which are intended to create a stiff foundation. In the horizontal direction, both the stiffness bars and the bridge assembly help to close the structure loop across the open faces of the housing assembly. However, much of the material in the stiffness bars and the bridge assembly is in bending with respect to horizontal forces, and hence, does little to improve the overall stiffness of the housing.

Additionally, the vertical actuator is suspended from its sides over an open hole in the base assembly. The shim plates on the sides are loaded in bending and not attached to the housing walls. The combination of these two issues lead to a vertical stiffness that could be improved.

The link between the actuators and the L4C seismometers is the critical sensor/actuator connection in the housing assembly. Both the vertical and horizontal L4C seismometers are located inside of the interface assembly. This is convenient for the vertical actuator since the sensor rests directly above the actuator and the link between them is stiff. In the horizontal direction, the sensor is mounted on the side of the interface assembly, and in as much as the interface assembly is stiff in compression, the horizontal sensor is well collocated with the horizontal actuator.

7.1.2 A Proposed Actuator Housing

The new design proposed here preserves many of the aspects of the previous, as-built, housing and benefits from experience accrued while working with the original design. However, many aspects of the proposed design (figures 7.2 and 7.3) differs substantially from the original assembly. The proposed design shown here is not detailed and should be regarded as a conceptual design.

The goal of the proposed housing design is to reduce the number of parts while greatly facilitating the installation of both the housing and the individual actuators. This objective is primarily enabled by a 5-DOF mount which connects the actuators to the the interface assembly.

The 5-DOF mount is comprised of a ball connected with two orthogonal bars to a cylinder. The 5-DOF mount in combination with a set of slotted holes on the interface assemble allows for alignment compensation in all 6-DOF. Once the mount has been placed as desired, it may be locked in place by securing the clamping fasteners on the ball and cylinder clamps.

The result of this is the that the housing may be installed on the pier top with disregard to misalignment between the two installation surfaces and without shims. The primary benefit is ease of installation, but other aspects of the design are enabled

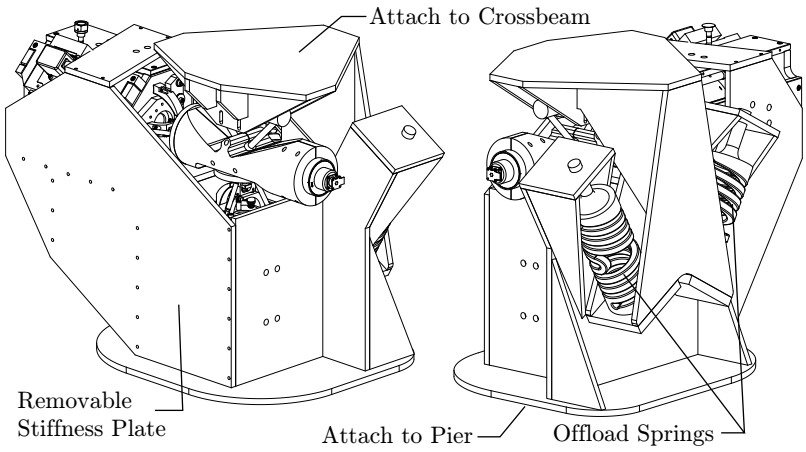

Figure 7.2: Isometric views of the proposed actuator housing design. Not shown in this design are the required payload locating hardware and rolling balls to translate the housing and payload on top of the pier.

by the use of this mount. Specifically, since the actuator positions no longer need to be adjustable, the actuators may be placed to maximize stiffness and accessibility. Moreover, the interface assembly does not need to be shimmed, and hence, the housing may be installed independently of the actuators.

Freedom of actuator placement allows for the for the structure to be manufactured from plate elements. This results in a shell structure that is inexpensive to manufacture and exhibits very high stiffness. The proposed housing is comprised entirely of 12 mm thick aluminum plates which are intended to be water-jet cut and dip-brazed together. Two removable stiffness plates are added to the base assembly to improve the horizontal stiffness.

There are some other simplifications to the installation procedure which are particular to the proposed design. The overall height of the proposed design does not exceed the gap between the crossbeam and pier top which allows for the assembly to be put in place directly from any number of angles. The vertical actuator rests on top of the base plate of the base assembly as opposed to down inside the pier. This simplifies both the installation and removal procedures.

With regard to sensor and actuator collocation, the proposed design assumes that

Alignment	
DOF	Method
X	Cyl/Ball/Clamp
Y	Cylinder
Z	Clamp
θ_x	Ball
θ_y	Ball
θ_z	Ball

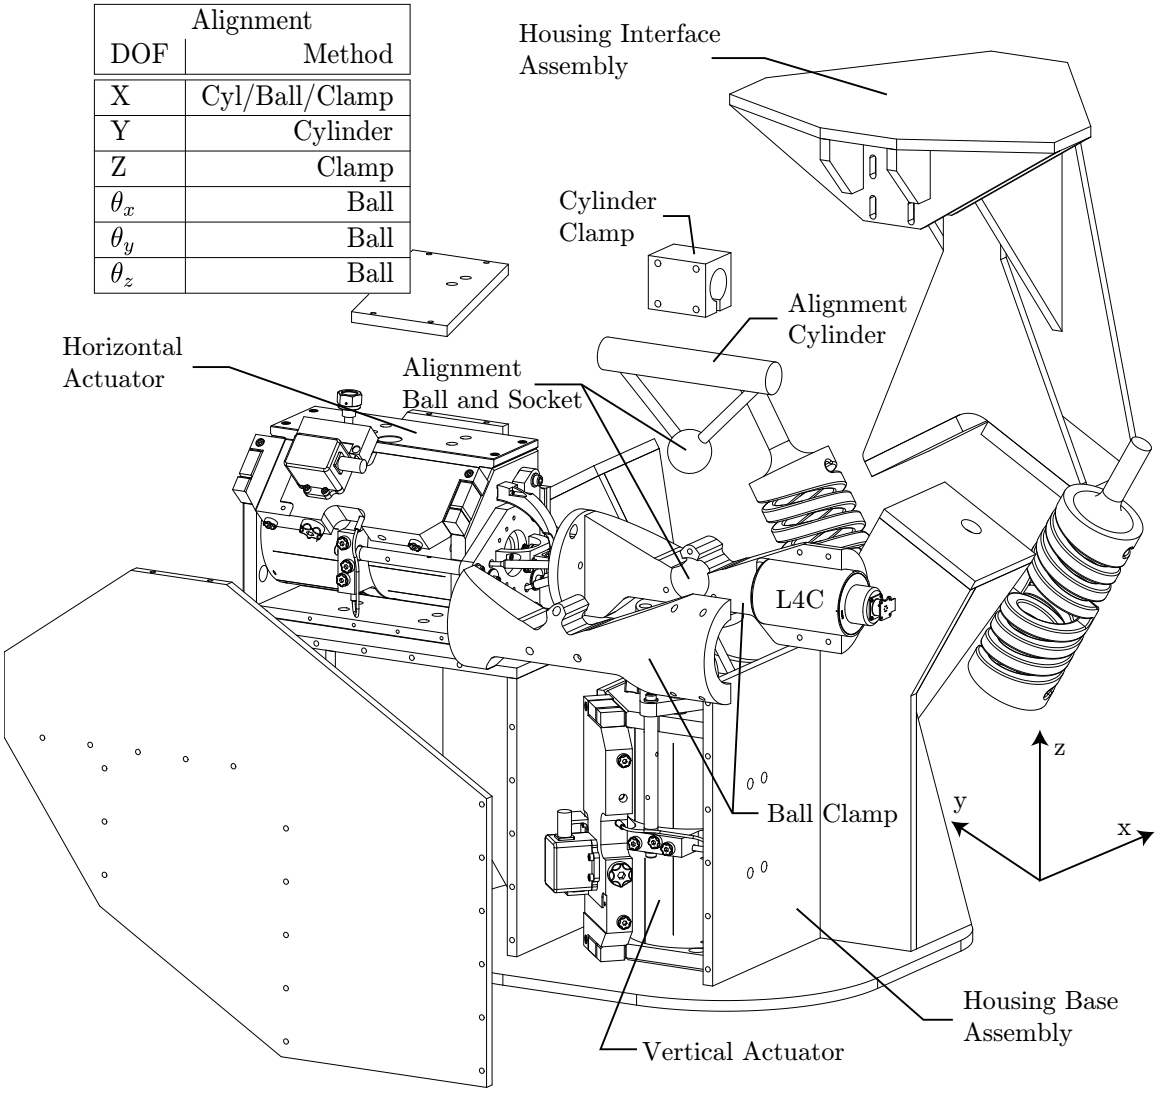

Figure 7.3: An exploded view of the proposed actuator housing design. The adjustment in Z is given by moving the cylinder clamp in the slotted holes on the Housing Interface Assembly.

the vertical L4C is placed inside of the actuator. The horizontal L4C is clamped by the ball clamp and shares a direct connection with the top of the horizontal actuator.

Further Enhancements to the Proposed Housing

The proposed housing design offers many improvements for interfacing the existing actuator design to the LIGO vacuum chambers. However, many more improvements are possible if modifications to the actuator design are allowed. In particular, the design of the actuator can be simplified if the actuator assembly is more tightly integrated with the housing.

For example, the side foundation plates of the actuator interface directly with vertical plates of the base housing assembly and are therefor redundant. An actuator could be assembled without side plates and held together with only the manifold. When the actuator is inserted into the housing assembly, the top and bottom foundation plates are bolted to the housing base assembly. In this configuration, the three side plates of the actuator may be excluded. The resulting actuator is less costly due to the reduction of parts and significantly lighter weight.

7.2 The Floating Seal Actuator

The component which most critically limits the potential of the quiet hydraulic actuator discussed in the previous chapters is the bellows. The bellows limit the maximum displacement and bandwidth of the actuator, and contribute significantly to the manufacturing cost (in part due to the added complexity of the bypass network). It would be a great improvement if the impressive noise performance of the quiet hydraulic actuator could be duplicated with a more conventional cylinder geometry that would likely offer a considerably higher hydraulic stiffness.

In this section, a new configuration of quiet hydraulic actuator is described, called the floating seal actuator (FSA). The force produced by the FSA does not depend on the compliance of bellows. Instead, the cylinder and piston of the FSA are comprised of a pair of hydraulic bearings (figure 7.4), and the bellows in this design serve only as a compliant seal (unpressurized).

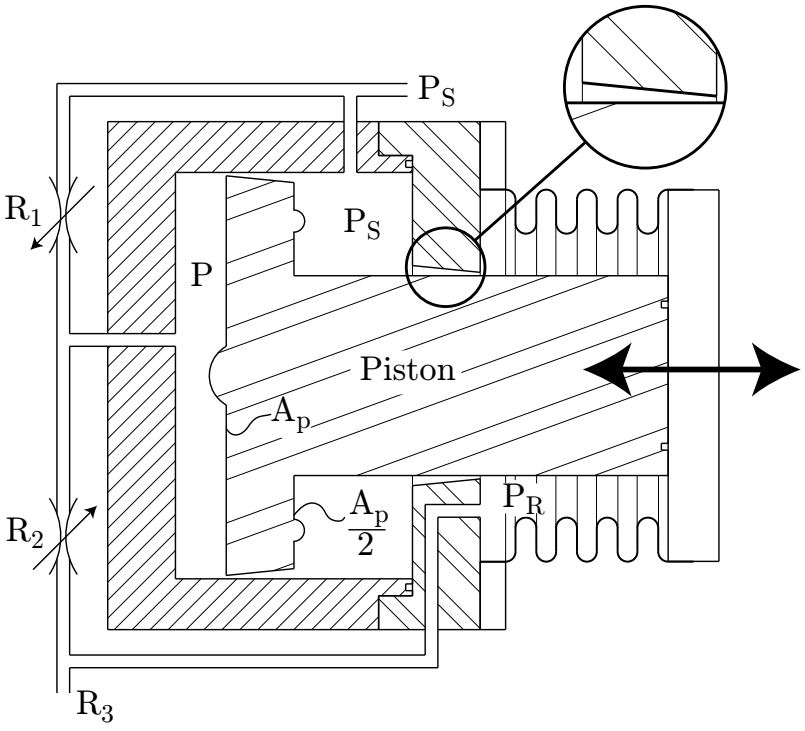

Figure 7.4: The floating seal actuator schematic.

7.2.1 Hydraulic Bearings

The functionality of the FSA is based upon hydraulic bearings. Of particular concern is the transverse stiffness of the hydraulic bearings. Clearly if the piston touches the cylinder wall, the actuator will not meet the noise performance specification. This subsection is devoted to the design of the hydraulic bearings. In general, there are two types of hydraulic bearings: externally and internally compensated.

Externally Compensated Hydraulic Bearings

The externally compensated bearing [33] may be the most common and easily understood. A common configuration for this type of bearing is a circular puck used as a 2-axis translational bearing to support a load (figure 7.5).

The puck supports the load by capturing a pocket of pressurized fluid in the cavity beneath the puck. This fluid flows out around the perimeter of the puck at a rate

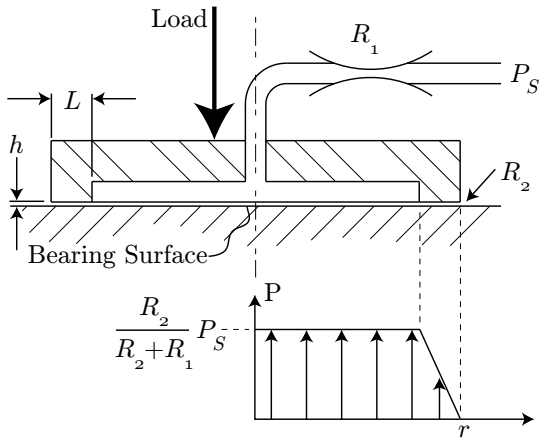

Figure 7.5: A schematic of an externally compensated hydraulic bearing puck. The puck is intended to support a vertical load by injecting pressurized fluid, at P_S, through the external resistance, R_1, into the cavity beneath the puck and out the perimeter of the puck. The parallel surfaces created by the perimeter of the puck (separated by a height, h, for a distance L) create the internal resistance, R_2.

which is modulated by the parallel plate resistance created by the lip at the perimeter of the puck. The lip has a height from the bearing surface of, h, and a length, L. If the height is reduced by an increase in load, the resistance of the parallel plates, R_2, increases. An increase in R_2 causes there to be a larger fraction of the supply pressure applied to the bottom of the puck. In this way the bearing exhibits stiffness and compensates against variations in applied load.

This configuration is labeled externally compensated because the first resister, R_1, is external to the puck. Since R_1 is external, size and manufacturing constraints often dictate that the restriction be a small orifice or similar geometry that is prone to clogging and turbulence. Since for this application, these two detriments are to be avoided, it is desirable to consider an internally compensated bearing.

Internally Compensated Hydraulic Bearings

An internally compensated bearing behaves similarly to an externally compensated bearing, but differs in that both resistances are built into the interface between the bearing surface and the underside of the puck. The underside of the puck has the

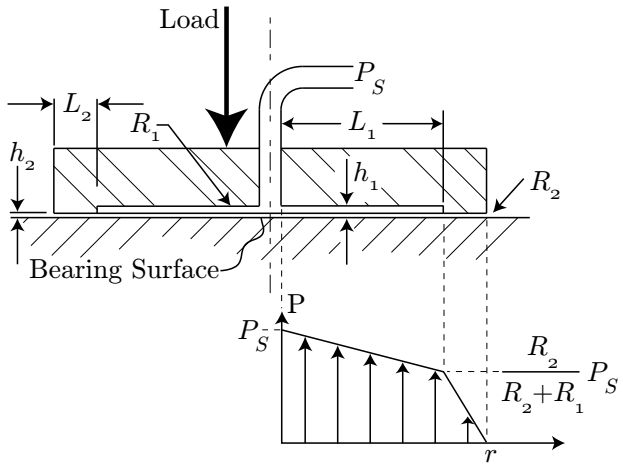

Figure 7.6: A schematic of an internally compensated hydraulic bearing puck. The underside of the puck has a stepped profile that acts as two parallel plate impedances.

same shape as the puck shown for the externally compensated case except that h_1 is much smaller in this case. The effect of this is to produce two parallel plate resistances in series. In this case, if the load increases causing the heights, h_1 and h_2, to decrease, the pressure profile across L_1 approaches a horizontal line while the profile across L_2 becomes more steep. Since the force exerted by the puck is proportional to the area under the pressure profile, the force increases as the puck approaches the bearing surface.

The internally compensated hydraulic bearing shares the load compensation capabilities featured by the externally compensated bearing but does not have the detriments typical of the external resistance. However, the internal bearing configuration of figure 7.6 has an abrupt edge between the two parallel plate resistances which may not be desirable for avoiding turbulence. Hence, the approach taken here is to adopt a tapered profile [32].

The Tapered Hydraulic Bearing

The tapered bearing (figure 7.7) is an internally compensated bearing that is functionally similar to the internally compensated bearing described above, but blends

the step between to the two resistances into one tapered profile. This section focuses on the derivation of the governing equations for this type of bearing since it is the proposed bearing for the FSA.

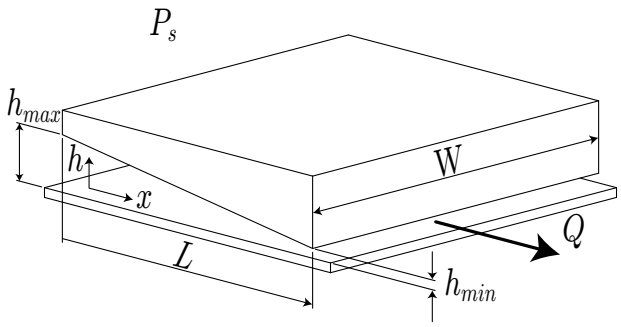

Figure 7.7: The basic geometry for a tapered hydraulic bearing. A source provides fluid at pressure, P_S, which enters the tapered parallel plate gap at h_{max} and produces a flow, Q. The width, W, is assumed to be quite large by comparison to h_{max} such that flow out the sides can be ignored.

The tapered hydraulic bearing is a parallel plate impedance with a height that varies with the length. The impedance of a parallel plate resistor with a constant height is:

$$R_{pp} = \frac{12\mu L}{Wh^3} \tag{7.1}$$

where μ is the viscosity of the fluid.[1] In terms of the variables defined in figure 7.7, the change in height along the length of the tapered interface can be defined as:

$$\frac{h}{h_{min}} = a + b\frac{x}{L} \tag{7.2}$$

$$a = \frac{h_{max}}{h_{min}} \tag{7.3}$$

$$b = 1 - a \tag{7.4}$$

To describe the resistance with a varying height, it is useful to define resistance change

[1]Note that the derivations shown here are for straight sections of a tapered bearing. This is slightly different from the curved perimeter of the puck that was discussed (as a familiar example) in the previous sections. The parallel plate resistance for a puck is $R_{pp} = 12\mu L/(2\pi h^3)\ln(r_o/r_i)$, but when the diameter of the puck is large by comparison to the length of the parallel plates, the ratio r_o/r_i approaches unity. As this occurs, $ln(r_o/r_i)$ becomes $(r_o - r_i)/r_i$ or L/r_i.

in of differential increments of length:

$$dR = \frac{dP}{Q} = \frac{-12\mu dx}{Wh^3} = \frac{-12\mu L}{Wh_{min}^3}\left(\frac{h_{min}}{h}\right)^3 \frac{dx}{L} \tag{7.5}$$

The overall resistance for the tapered bearing, R_{tpp}, is the integral of equation 7.5 over the length:

$$R_{tpp} = \frac{1}{Q}\int_0^L dP = R_{pp}\int_0^L \left(\frac{h_{min}}{h}\right)^3 \frac{dx}{L} \tag{7.6}$$

$$R_{tpp} = \frac{P_S - P_R}{Q} = R_{pp}\frac{2a+b}{2(a+b)^2a^2} \tag{7.7}$$

where P_S and P_R are the supply and return pressures. The impedance can be used to determine the pressure profile by first defining the flow as:

$$Q = \frac{P_S - P_R}{R_{tpp}} = (P_S - P_R)\frac{2(a+b)^2a^2}{2a+b} \tag{7.8}$$

The pressure profile is then:

$$\int_P^{P_S} dP = (P_S - P_R)\frac{2(a+b)^2a^2}{2a+b}\int_0^x \left(\frac{h_{min}}{h}\right)^3 \frac{dx}{L} \tag{7.9}$$

$$P(x) = P_S - (P_S - P_R)\frac{2(a+b)^2a^2}{2a+b}\frac{x(2aL+bx)}{2(aL+bx)^2a^2} \tag{7.10}$$

The pressure profile can be integrated to find the force created by the bearing:

$$F_{tpp} = \int_0^L P(x)dA = W\int_0^L P(x)dx = \frac{WL(P_Sa + P_r)}{a+1} \tag{7.11}$$

If the substitutions are made for a and b, the force is:

$$F_{tpp} = WL\frac{P_Sh_{max} + P_Rh_{min}}{h_{max} + h_{min}} \tag{7.12}$$

Two tapered bearings may be arranged around a common bearing surface, as in figure 7.8. In this case, the bearing force acts differentially. To discuss this behavior,

the variables describing the taper are defined in the context of figure 7.8:

$$a_{top} = \frac{h_{max}}{h_{min}} = \frac{H_t + h_t - z}{h_t - z} \tag{7.13}$$

$$a_{bot} = \frac{h_{max}}{h_{min}} = \frac{H_t + h_t + z}{h_t + z} \tag{7.14}$$

Similar to equation 7.12, the force created by the top and bottom bearings are:

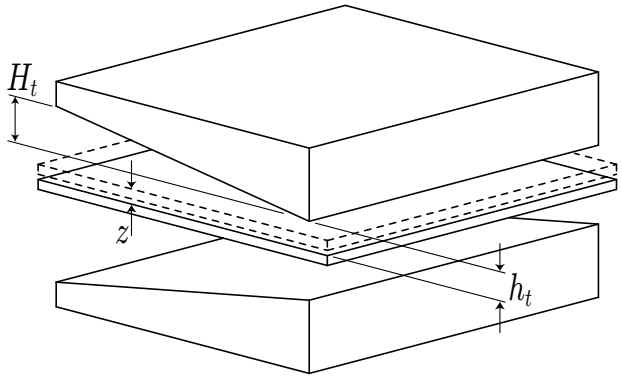

Figure 7.8: Two tapered hydraulic bearings sharing a common bearing surface.

$$F_{top} = WL\frac{P_S a_{top} + P_R}{a_{top} + 1} \tag{7.15}$$

$$F_{bot} = WL\frac{P_S a_{bot} + P_R}{a_{bot} + 1} \tag{7.16}$$

And the total force will be:

$$F_{dtpp} = WL\left(\frac{P_S a_{top} + P_R}{a_{top} + 1} - \frac{P_S a_{bot} + P_R}{a_{bot} + 1}\right) \tag{7.17}$$

Given this force, it is possible to determine a spring constant for motions of the bearing surface. The spring constant for the differential tapered bearing is the total force divided by the displacement of the bearing surface, $K_{dtpp} = F_{dtpp}/z$. If equations

7.13 and 7.14 are substituted into equation 7.17, the spring stiffness is:

$$K_{dtpp} = \frac{2H_t(P_S - P_R)}{(H_t + 2h_t)^2 - 4z^2}WL \tag{7.18}$$

$$K_{dtpp} \cong \frac{2H_t(P_S - P_R)}{(H_t + 2h_t)^2}WL \quad \text{for small z} \tag{7.19}$$

It is useful to determine a specific spring stiffness that is not dependent on the length and width:

$$K_0 = \frac{2H_t(P_S - P_R)}{(H_t + 2h_t)^2} \tag{7.20}$$

The values of H_t and h_t can be chosen to maximize the specific spring stiffness by finding the derivative of the specific spring stiffness with respect to H_t and setting it to zero.

$$\frac{dK_0}{dH_t} = \frac{2H_t}{(H_t + 2h_t)^2} - \frac{4H_t}{(H_t + 2h_t)^3} = 0 \tag{7.21}$$

$$H_t = 2h_t \tag{7.22}$$

Hence, the tapered bearing displays the highest stiffness when the entry height is three times the exit height ($h_{max} = 3h_{min}$ from figure 7.7 is the same as equation 7.22). Equation 7.22 can be substituted into equation 7.17 to solve for the maximum force when $z = h_t$:

$$F_{max\ dtpp} = (P_S - P_R)\frac{WL}{3} \tag{7.23}$$

The bearing stiffness of the tapered bearings in the FSA is less than the stiffness given in equation 7.19 because the FSA bearings are cylindrical instead of flat. The force created for a given transverse displacement is less because the bearing stiffness acts radially as opposed to along the direction of the displacement (vertical, in figure 7.9). In figure 7.9, a cylindrical bearing is displaced vertically a distance z. The vertical displacement z translates into a radial displacement of $cos(\theta)z$ and the vertical force created by this displacement is $F_p = K_r cos(\theta)z$ where K_r is the bearing stiffness in the direction of the displacement

The differential of the radial spring stiffness, K_r, is $dK_r = K_0 cos(\theta)LdW =$

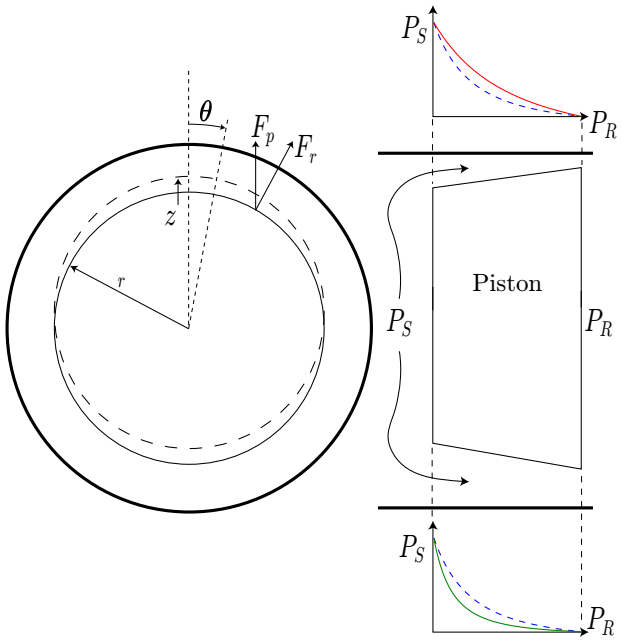

Figure 7.9: The force resultants and pressure profiles for a transverse displacement of a cylindrical tapered bearing. As the bearing (Piston) undergoes a transverse displacement (vertically shown here), the pressure profile changes from the nominal, piston centered profile (dashed) to a displaced profile (solid). At the smaller gap on top, the area under the pressure profile is increased from the centered position while the bottom has decreased.

$K_0 cos(\theta) L r d\theta$. Then the differential force in the direction of the displacement is $dF_p = K_0 cos(\theta)^2 L r d\theta$. The total force is the integrated force in the direction of the displacement:

$$F_p = L K_0 z \int_0^{\pi} cos(\theta)^2 r d\theta = \frac{1}{2} \pi K_0 L r z \tag{7.24}$$

Therefore, if a flat bearing has a width of $W = 2r$, then for a flat bearing the stiffness is $K_{dtpp} = K_0 2rL$ and the stiffness for a cylindrical bearing of the same width is:

$$K_{rdtpp} = \pi/2 K_0 rL \tag{7.25}$$

The stiffness ratio for the cylindrical to the flat case is surprisingly high at $\pi/4$.

The implicit assumption in the calculation for the cylindrical tapered bearing is

that the circumferential component of the flow is very small. This should be plausible so long as $L \ll 2\pi r$. When the length along the bearing, L, is much smaller than the circumference, $2\pi r$, the resistance to tangential flow should be high enough by comparison to that in the direction along the bearing to make flow in the tangential direction negligible.

7.2.2 Design of the Floating Seal Actuator

The component unique to the FSA, the hydraulic bearings, have been discussed in the previous sections. What remains is to size the FSA and design an interface such that the FSA can be used in multi-actuator platforms.

FSA Attachment Flexure

Similar to the hydraulic actuator of previous chapters, the FSA must operate simultaneously with other actuators in orthogonal degrees of freedom. The FSA must be able to apply force axially while being compliant in transverse degrees of freedom. Hence, the FSA is designed with an attachment flexure that is stiff axially, but soft in the other five degrees of freedom.

Unlike the bellows of the earlier design, the hydraulic bearings of the FSA resist tilt. This complicates the design of the attachment flexure because the flexure must now provide both tilt and two degrees of rotational freedom. This is accomplished by employing two tripod structures, one on either end of the actuator, separated by a length, L (figure 7.10). As the attachment plate is translated in the z direction, the attachment flexure tilts. The attachment plate and the piston are decoupled from this tilt by the two tripod flexures which are soft in rotation.

Sizing the FSA

One of the most promising advantages of the FSA is that the bellow is only used for sealing and does not compromise the stiffness of the actuator. While this has implications of higher bandwidth, it also offers the opportunity for a much larger maximum displacement which is limited on the current design solely by the bellows.

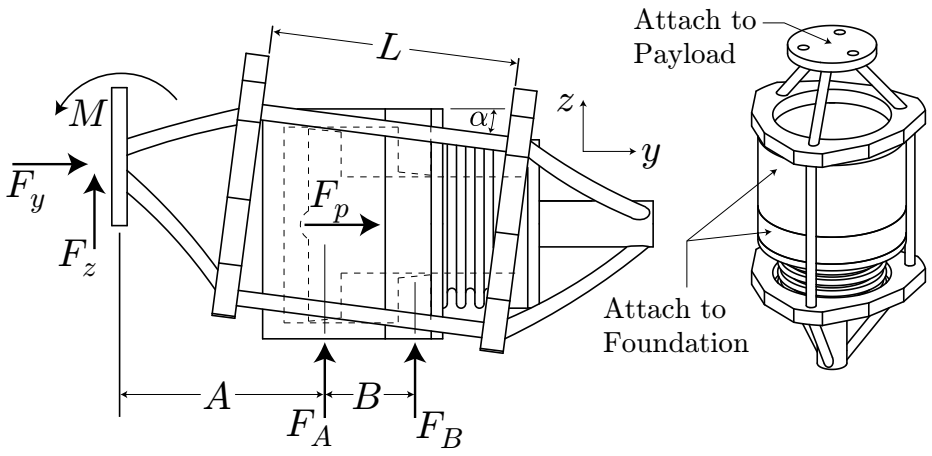

Figure 7.10: The floating seal actuator with attachment flexure.

In this section, the FSA will be sized for a large displacement: $\delta_{max} = 5$ mm. This large displacement places the FSA at the edge of the available design space (figure 7.11).

The configuration shown in figures 7.4 and 7.10 may be improved upon in several ways. In the proposed design (figure 7.13) the bellows are placed internal to the piston. This arrangement clearly offers better protection for the bellow. More importantly, the length of the bellow can be large without adding unnecessary length to the actuator. With the provision for more length, it is easier to find a bellow that will have a low axial stiffness while still capable of withstanding the full supply pressure (as a safety precaution). The bellow shown in figure 7.13 is sized to have an axial stiffness much less than $1/10^{th}$ the offload spring stiffness while still capable of surviving at 2 MPa. Another advantage of this configuration is that the bearings are separated by a large distance. This reduces the required bearing stiffness needed to accommodate transverse forces and torques.

The dimensions of the bearings in the proposed design are derived from the equations discussed in section 7.2.1. The following is a brief discussion of how these equations were applied to create the proposed design.

Given a maximum displacement of 5.0 mm and a natural frequency of 7 Hz, the spring stiffness of the offload springs can be calculated based on the known mass.

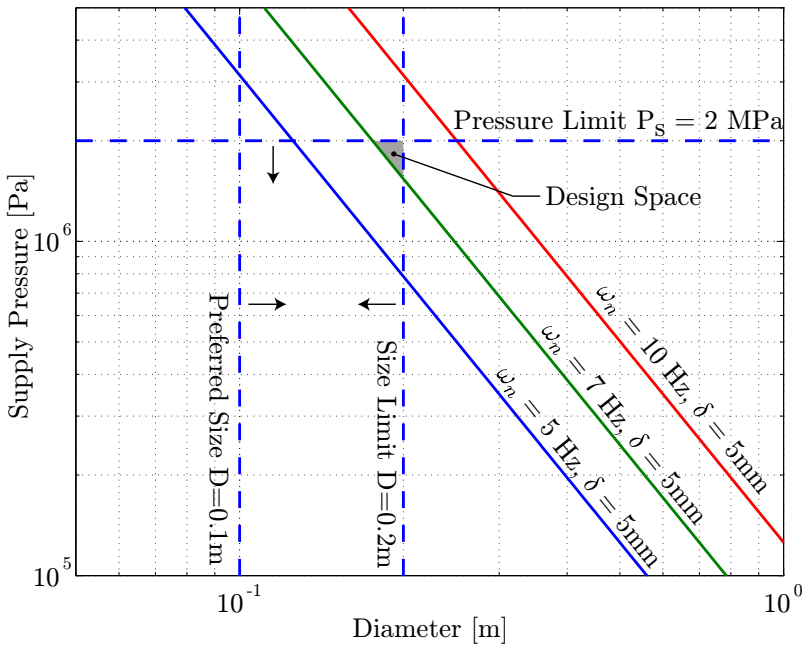

Figure 7.11: The design space available for the FSA given the natural frequency and maximum throw (displacement).

For the BSC tank in LIGO, the mass per corner is 500 kg and the offload spring stiffness is $9.7{\cdot}10^5$ N/m. As is indicated in figure 7.11, it is difficult to attain 5 mm of displacement without exceeding the maximum allowable size or supply pressure. If a cylinder diameter of 20 cm is employed, the supply pressure can be calculated based on the equation:

$$F_{max} = \frac{P_s A_p}{2} \tag{7.26}$$

where A_p is the cross-sectional area of the cylinder. The factor of 1/2 stems from the supply side of the piston having 1/2 the total area (figure 7.4). Unfortunately, solving for the supply pressure with equation 7.26 will yield a supply pressure that is too low to support the flow through the hydraulic bearings. At the maximum force, equation 7.26 dictates that the pressure on either side of bearing A (figure 7.10) will be equal. This is not acceptable for hydraulic bearings, and it is therefore necessary to add a factor of safety. A factor of safety of 5 should be sufficient to guarantee flow across the bearing in all situations. With this factor of safety, the required supply

pressure becomes $P_s = 1.5$ MPa.

The next step is to determine the required stiffness for the hydraulic bearings. The required stiffness can be determined by calculating the maximum transverse force implied by the maximum transverse displacement of a neighboring actuator. For bearing B (figure 7.10), this can be calculated by summing moments at the center of bearing A:

$$\sum M_x = 0 = M - L\sin(\alpha)F_y - F_z A + F_B B \tag{7.27}$$

$$F_B = \frac{1}{B}(F_z A + L\sin(\alpha)F_y - M) \tag{7.28}$$

$$F_B \cong \frac{1}{B}(L\sin(\alpha)F_y) = \frac{\delta_{max}}{B}F_y \tag{7.29}$$

where F_y is the maximum force of the actuator ($F_y = F_{req} = \delta_{max}K_{off}$) and the other variable are described in figure 7.10. The dimension B is 9.1 cm in the proposed design which by equation 7.29 implies that the force supported by bearing B is at most 265 N. Under this load, the bearing is allowed to displace by half the minimum bearing gap, h_{min}. This can be used to solve for the bearing length by using equations 7.25 and 7.20:

$$F_B = K_b\frac{h_{min}}{2} = \frac{3\pi r_b L_b(P_S - P_R)}{25h_{min}}\frac{h_{min}}{2} \tag{7.30}$$

$$L_b = \frac{50F_B}{3\pi r_b(P_S - P_R)} \tag{7.31}$$

where r_b is the bearing radius and L_b is the bearing length. For the proposed design, the length L_b is 9.2 mm. This is a minimum length, and since there is little penalty for increasing the length of the bearing, a bearing length of 2 cm is used in the proposed design.

Finally, the bearing height, h_{min}, is determined by establishing a maximum power dissipation across the hydraulic bearings. The hydraulic power dissipation is the square of the pressure drop divided by the resistance. The resistance for the tapered bearings is given by equations 7.1 and 7.7. If the maximum power is set at 3.0 watts,

the minimum bearing height can by solved for by:

$$h_{min} = \sqrt[3]{\frac{4\mu L_b(3.0W)}{3\pi(P_S - P_R)^2 r_b}} \tag{7.32}$$

where μ is the viscosity of the hydraulic fluid. For the proposed design the minimum gap height is 22 μm and a taper of 44 μm over the 20 mm length.

Machining a taper of 44 μm over a length of 20 mm is a step beyond common machine shop practice. However, with some fixtures and modifications, it is possible to manufacture tapers of this precision. The process [34] involves first adding a transverse degree-of-freedom to the tailstock of a lathe. The taper surface to be machined is attached to a fixture of substantially greater length (figure 7.12). The increased length given by the round stock amplifies the small angle such that the transverse motion of the tailstock is large enough to be measured with conventional tools.

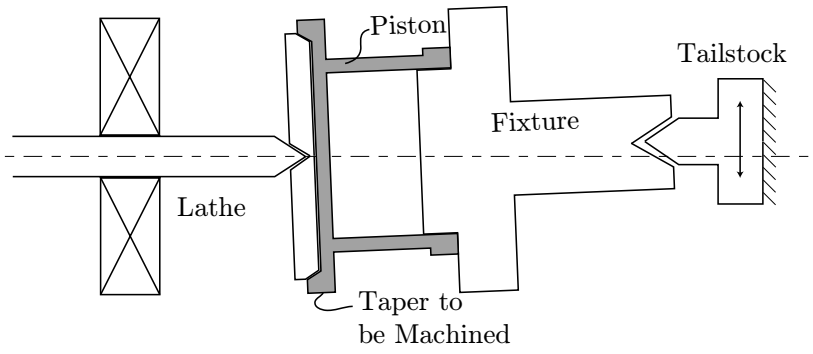

Figure 7.12: A process for machining high precision tapered surfaces.

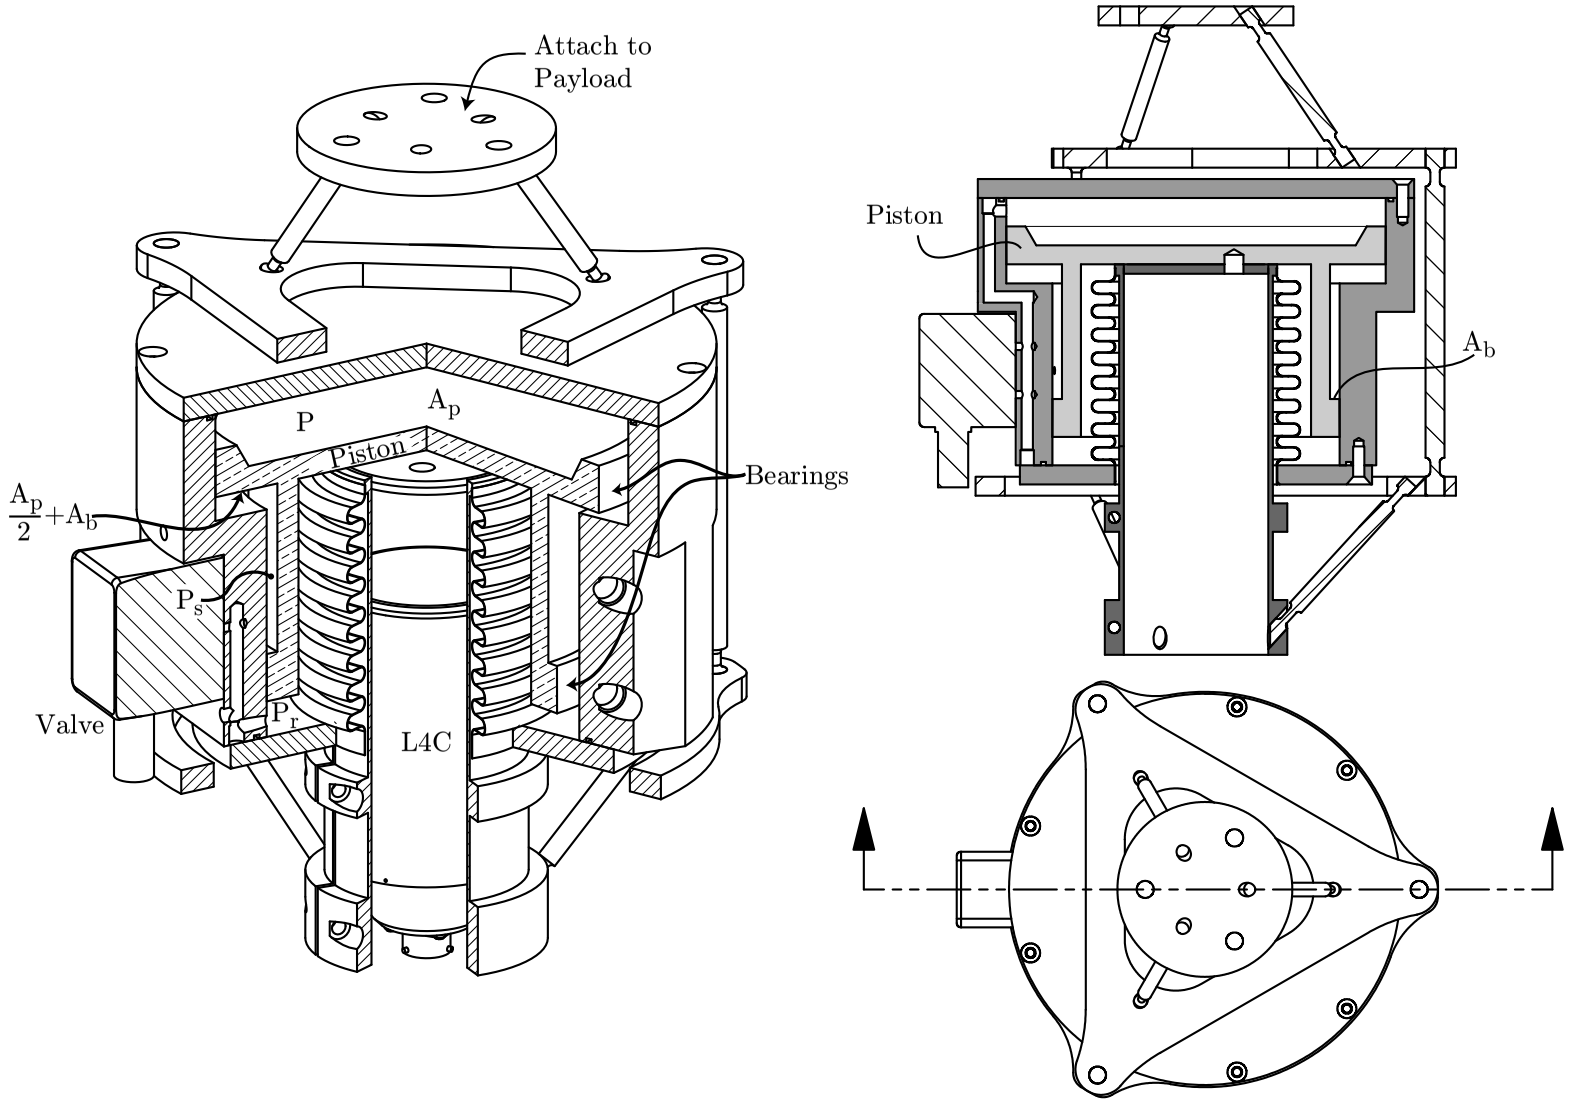

Figure 7.13: The proposed design for the floating seal actuator. Note that the bellows are hidden within the piston. This enables the bellows to be long for axial compliance without increasing the overall length of the actuator. The piston and the cylinder are made from single pieces of material which should promote the necessary parallelism of the bearing surfaces. The primary inertial sensor, the L4C, is clamped in the center of the actuator for maximum collocation. The area behind the piston is $A_p/2 + A_b$ to compensate for the area relieved around the lower (in this view) bearing.

Appendix A

The DYP2S Valve

Several Parker manufactured DYP2S valves have been in service on experiments at Stanford and MIT. This section is included to discuss some of the operational details particular to this make and model of valve.

A.1 Valve Asymmetries

The DYP2S valve has two flappers and four nozzles. As such, there are four nozzle – flapper interfaces are labelled in Figure A.1 as R1, R2, R3 and R4.

For each of the nozzle – flapper resistances, the position of the nozzle with respect to the flapper in its nominal position is tuned during calibration to generate equivalent resistances (by moving the nozzle closer or further from the flapper). Despite this, the measurement of resistance on either side of a flapper – nozzle pair during operation is not equal. This is a result of the flow coming from the upstream nozzle (R1 or R4) displacing the flapper slightly towards the downstream nozzle (R2 or R3 respectively). This causes the resistance between the upstream nozzle and flapper to be lower, and consequentially, the resistance between the downstream nozzle and flapper is slightly higher. This effect can be seen in that the calibration curves for one side of the DYP2S valve do not cross at zero flapper displacement as shown in Figure A.2.

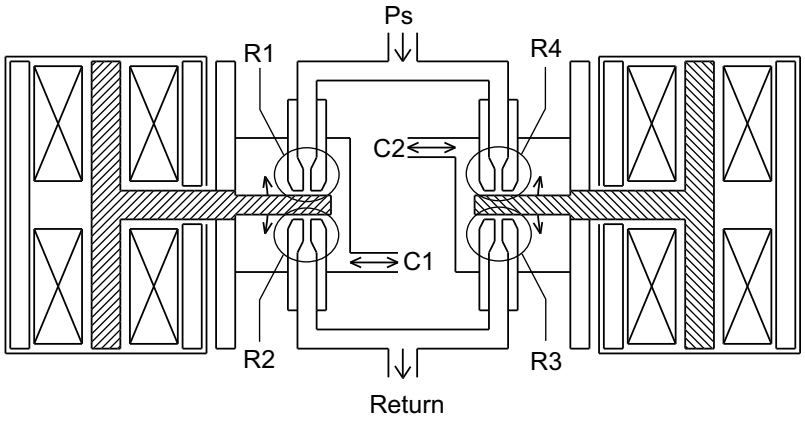

Figure A.1: The DYP2S valve.

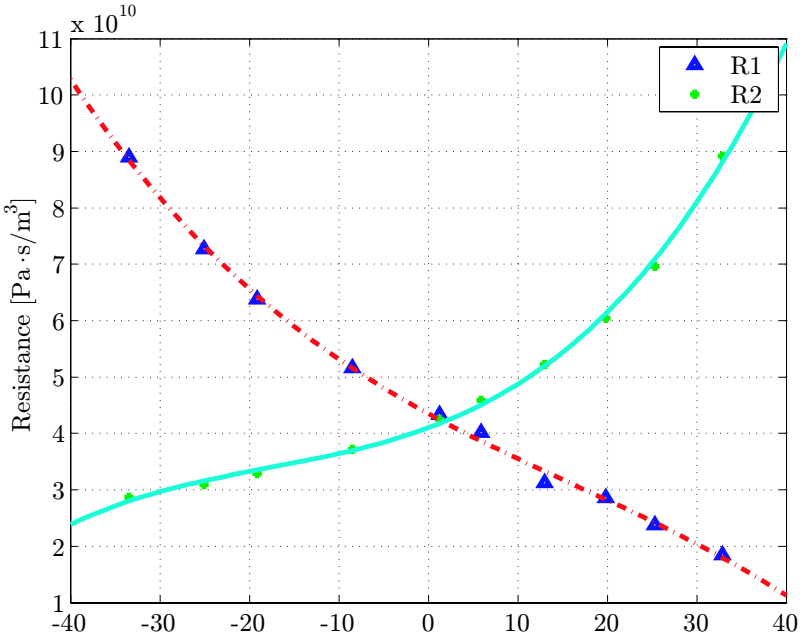

Figure A.2: Measured resistance of one side of the DYP2S valve.

A.2 Flapper Issues

There are some known issues of the DYP2S valve involving the flapper. These difficulties can be avoided by only operating the valve through a fraction of its total range. In light of this, the HEPI experiment at LASTI is drive current limited to one quarter of the maximum rated current. Driving the valve to the full allowed current may result in *flapper snap down* or flapper oscillations.

A.2.1 Flapper Snap Down

Flapper snap down occurs when the nozzle is driven close to the downstream nozzle. As the flapper moves into close proximity with the downstream nozzle, the flapper may create a seal with the downstream nozzle face perhaps with the aid of the force generated by the flow striking the upstream side of the flapper. Once a seal is created, the suction of the flow moving towards the return ensures that the flapper will remain sealed to the nozzle face.

It is possible to recover from this state by driving the valve motor in the opposite direction with a drive current approaching the maximum allowed, but the recovery is far from linear. This should be avoided in any control system application.

A.2.2 Flapper Oscillation

Flapper oscillation has been observed in the DYP2S valve when the flapper is forced up against the upstream nozzle. The flapper attempts to seal on the upstream nozzle face, but in so doing, causes the pressure at the upstream nozzle exit to build to the full supply pressure. Once this occurs, the flapper is forced back until the pressure is reduced and the process repeats. This limit cycle effect has been observed to cause some damage to the face of the flapper where it impacts the nozzle face, but otherwise, the valve can recover immediately when the drive current is reduced. This nozzle oscillation could be avoided by modifying the upstream nozzle as shown in Figure A.3. In this proposed solution the DYP2S behaves more similarly to the HSC valve with the external resistors. There are two disadvantages to this. First, the

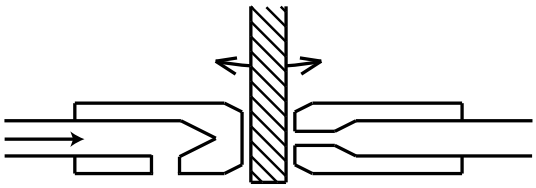

Figure A.3: A possible alternative to the DYP2S nozzle layout.

maximum pressure recovery (the maximum differential pressure) is halved because the bridge is no longer fully differential. The second issue is that of the upstream resistance. While it is tempting to reduce the diameter of the new hole in the nozzle to meet the resistance requirement, the resulting diameter will be very small and sensitive to debris. In this case it is better to leave the port diameter large and use exterior laminar flow resistors such as those available from the LEE corporation.

Appendix B

A Model of the Actuator and Foundation

A model of the actuator was developed in chapter 3. However, this model did not include the discoveries of later chapters. Here a complete, one-dimensional model of the actuator is derived that includes both the breathing stiffness of the bellows and the bypass network.

The derivation begins with force balances for the suspended mass (M_m), the actuator plate (M_a) and the foundation (M_f):

$$M_m \ddot{z_m} = K_c(z_a - z_m) + K_{off}(z_g - z_m) + D \tag{B.1}$$

$$M_a \ddot{z_a} = \Delta P A_b + K_c(z_m - z_a) \tag{B.2}$$

$$M_f \ddot{z_f} = -\Delta P A_b + K_f(z_g - z_f) \tag{B.3}$$

Equations of volume (same as equations 3.19 and 3.19):

$$\dot{P_1}\left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}}\right) = Q_1 + A(\dot{z_a} - \dot{z_g}) \tag{B.4}$$

$$\dot{P_2}\left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}}\right) = Q_2 - A(\dot{z_a} - \dot{z_g}) \tag{B.5}$$

Nodal equations (P_R is the return pressure):

$$\frac{P_S - P_1}{R_1} - Q_1 - \frac{P_1 - P_R}{R_2} - \frac{P_1 - P_R}{Z_m} = 0 \quad \text{(B.6)}$$

$$\frac{P_S - P_2}{R_4} - Q_2 - \frac{P_2 - P_R}{R_3} - \frac{P_1 - P_R}{Z_m} = 0 \quad \text{(B.7)}$$

where $Z_m = R_{byp} + K_{byp}/A_{byp}^2 s$ is the impedance of the bypass. The bypass impedance is comprised of a bypass resistance, R_{byp}, combined with a diaphragm of stiffness K_{byp} and area A_{byp}. The overall equations for the bridge (resident in the servo valve assembly):

$$P_S = P_1 + P_2 \quad \text{(B.8)}$$

$$\Delta P = P_1 - P_2 \quad \text{(B.9)}$$

these may be combined to yield:

$$2P_1 = P_S - \Delta P \text{ and } 2P_2 = P_S + \Delta P \quad \text{(B.10)}$$

Equations B.4 and B.6 can be used to eliminate Q_1:

$$\dot{P}_1 \left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}} \right) = \frac{P_S - P_1}{R_1} - \frac{P_S - P_1}{Z_m} - \frac{P_1}{R_2} + A(\dot{z_a} - \dot{z_g}) \quad \text{(B.11)}$$

For the next step, equation B.10 can be used to replace P_1 with the more convenient ΔP. In addition, the resistors R_1 and R_2 are replaced with their respective linearized models from equation 3.3.

$$\frac{\dot{P}_S - \Delta\dot{P}V_b}{2} \left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}} \right) = \frac{P_S + \Delta P}{2R_o(1 - 3\epsilon)} - \frac{P_S - \Delta P}{2R_o(1 + 3\epsilon)} + \frac{\Delta P}{Z_m} + A(\dot{z_a} - \dot{z_g}) \quad \text{(B.12)}$$

Two of the fractions on the right side of equation B.12 are nonlinear with respect to ϵ, but can be approximated with a Taylor series:

$$\frac{P_S + \Delta P}{2R_o(1 - 3\epsilon)} - \frac{P_S - \Delta P}{2R_o(1 + 3\epsilon)} \rightarrow \frac{\Delta P}{R_o} + \frac{P_S}{R_o}\epsilon + \left(\frac{P_S + \Delta P}{R_o} - \frac{P_S}{R_o} \right) \epsilon^2 \ldots \quad \text{(B.13)}$$

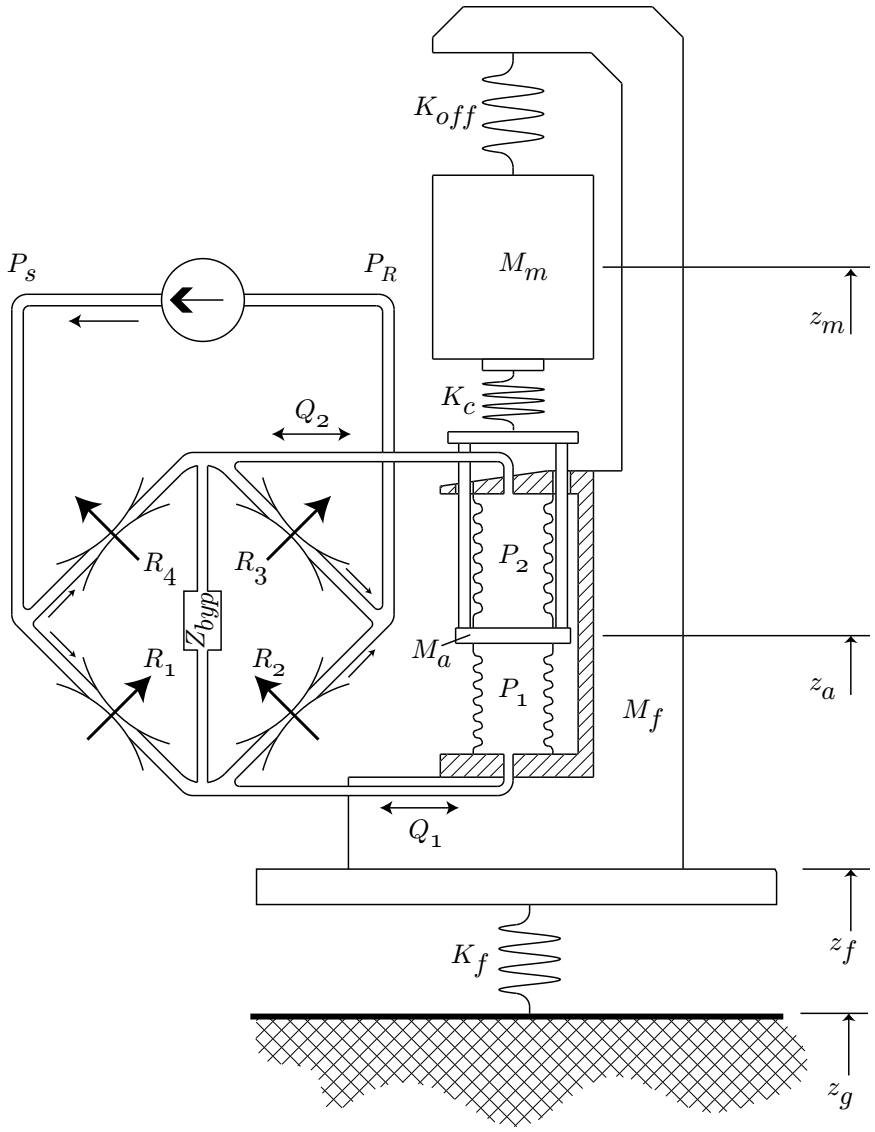

Figure B.1: A schematic of the actuator and foundation. Note the bypass impedance, Z_{byp}, in the middle of the bridge (physically located between the bellows – see chapter 5).

In the interests of linearity, only the first two terms of the expansion are maintained. The derivative of the supply pressure, $\dot{P}_S$ in equation B.12 may also be omitted because if the supply pressure is properly conditioned, variations in P_S should be small. The final equations of motion take the following form:

$$M_m \ddot{z_m} = K_c(z_a - z_m) + K_{off}(z_g - z_m) + D \quad \text{(B.14)}$$

$$M_a \ddot{z_a} = \Delta P A_b + K_c(z_m - z_a) \quad \text{(B.15)}$$

$$M_f \ddot{z_f} = -\Delta P A_b + K_f(z_g - z_f) \quad \text{(B.16)}$$

$$\frac{\dot{\Delta P}}{2}\left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}}\right) = -\Delta P\left(\frac{1}{R} + \frac{1}{Z_m}\right) + A(\dot{z_f} - \dot{z}) - \frac{P_s \epsilon}{2R} \quad \text{(B.17)}$$

Equation B.18 combines the above equations in state-space form with the bypass parameters shown explicitly. In equation B.18, Ψ is the compliance of the bellows in series with the fluid, or $\Psi = \left(\frac{V_b}{\beta} + \frac{A_b^2}{K_{breathing}}\right)$.

$$
\begin{pmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & -A_b R_{byp} A_{byp}^2 & 0 & -A_b R_{byp} A_{byp}^2 & 0 & -\frac{\Psi R_{byp} A_{byp}^2}{2}
\end{pmatrix}
\begin{pmatrix}
\dot{z_m} \\ \ddot{z_m} \\ \dot{z_a} \\ \ddot{z_a} \\ \dot{z_f} \\ \ddot{z_f} \\ \dot{\Delta P} \\ \ddot{\Delta P}
\end{pmatrix} =
$$

$$
\begin{pmatrix}
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
-\frac{K_c+K_{off}}{M_m} & 0 & \frac{K_c}{M_m} & 0 & \frac{K_{off}}{M_m} & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
\frac{K_c}{M_a} & 0 & -\frac{K_c}{M_m} & 0 & 0 & 0 & \frac{A_b}{M_a} & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
\frac{K_{off}}{M_f} & 0 & 0 & 0 & -\frac{K_f+K_{off}}{M_f} & 0 & -\frac{A}{M_f} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & K_{byp} A_b & 0 & -K_{byp} A_b & -\frac{K_{byp}}{R} & -\left(A_{byp}^2 + \frac{\Psi K_{byp}}{2} + \frac{R_{byp} A_{byp}^2}{R}\right)
\end{pmatrix}
\begin{pmatrix}
z_m \\ \dot{z_m} \\ z_a \\ \dot{z_a} \\ z_f \\ \dot{z_f} \\ \Delta P \\ \dot{\Delta P}
\end{pmatrix} +
$$

$$
\begin{pmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{1}{M_m} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
\frac{K_f}{M_f} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{P_s K_{byp}}{2R} & \frac{P_s R_{byp} A_{byp}^2}{2R}
\end{pmatrix}
\begin{pmatrix}
z_g \\ \dot{z_g} \\ \ddot{z_g} \\ D \\ \epsilon \\ \dot{\epsilon}
\end{pmatrix} \quad \text{(B.18)}
$$

Appendix C

Design Synthesis Appendix

C.1 The Electromagnetically Actuated EPI

The Livingston LIGO observatory (LLO) became operational while the development of the quiet hydraulic actuator was still in its infancy. Soon there after, the need for an updated External Pre-Isolator (EPI) became apparent as the excessive ground motion at LLO proved to be too much for the existing PEPI system. By virtue of this need, the development of the hydraulic actuator was greatly accelerated, but also in recognition of the risk associated with the new technology, a back up plan was adopted. Hence, in parallel with the development of a Hydraulic External Pre-Isolator (HEPI), the electoMagnetic External Pre-Isolator (MEPI) was constructed in case of some unforseen shortcoming of the quiet hydraulic actuator.

The results from the MEPI system offer a unique opportunity to compare two types of actuator in an otherwise identical system (figure C.1). The difference between the two systems is the lack of stiffness in the electromagnetic, voice coil actuators on the MEPI system. This lack of stiffness manifests itself as a plethora of peaks in the open loop transfer function which confounds the control design for the MEPI system.

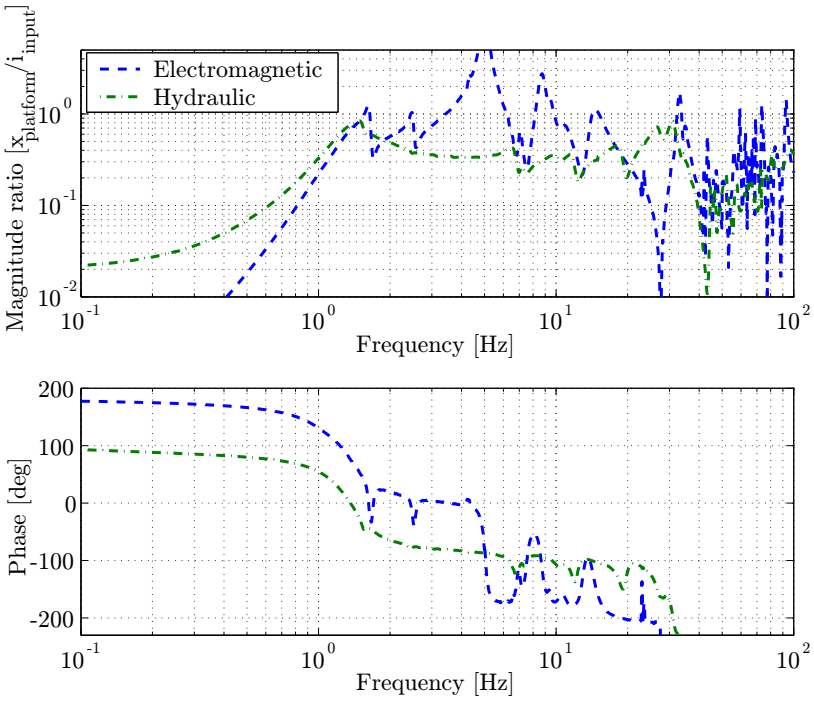

Figure C.1: The feedback seismometer output of a driven transfer function taken on both HEPI and MEPI.

C.2 Natural Frequency Temperature Dependence

In the discussion of suspensions that are defined as either soft or stiff, it is useful to define a specific frequency to discriminate between the two categories. In this work, this is defined as the suspension frequency where temperature variations cause equivalent positional drift from either: variations in suspension spring stiffness with temperature, or dimensional changes of the support structure caused by the temperature fluctuation.

While this effect is consistent for all suspension spring designs, consider a coil spring whose stiffness is defined by:

$$K = \frac{d^4 E}{\frac{2(1+\nu)}{8D^3 N_t}} \tag{C.1}$$

where d is the wire diameter, E is the modulus of elasticity, ν is poisson's ratio, D is the mean coil diameter and N_t is the number of coils. The change in the modulus of elasticity with temperature is:

$$dE = E\beta dT \tag{C.2}$$

where β represents the change in the modulus of elasticity with temperature (for steel $\beta = 2{\cdot}10^{-4}$ °C^{-1} [29]). Hence, the change in spring stiffness, K, with temperature is determined by inserting equation C.2 into equation C.1:

$$dK = K\beta dT \tag{C.3}$$

The static sag of a suspension is $Sag = mg/K$, and the change of sag with respect to the change in spring stiffness is:

$$\frac{dSag}{dK} = \frac{mg}{K^2}\frac{dK}{dT} = Sag \cdot \beta = \frac{g}{\omega_n^2}\beta \tag{C.4}$$

where ω_n is the natural frequency of the suspension and $\omega_n^2 = K/m$. Similarly, the length variation of the structure with temperature is:

$$\frac{dL}{dT} = \alpha L \tag{C.5}$$

where L is the length of the structure and α is the coefficient of thermal expansion for the structure. Then the frequency for which thermal variations cause equivalent changes in position is when $dL/dT = dSag/dK$ or:

$$\frac{g}{\omega_n^2}\beta = \alpha L \tag{C.6}$$

which can be rearranged to:

$$\omega_n = \sqrt{\frac{g}{\alpha L}\beta} \tag{C.7}$$

For a system that has a physical size of 1 m and is made of steel ($\alpha = 14{\cdot}10^{-6}$ °C^{-1} [29]), the resulting natural frequency is $\omega_n = 2$ hertz.

C.3 Collocation

The importance of collocation is easily demonstrated with a simple dynamic model. Consider two carts of mass M connected by a spring of stiffness K moving across a frictionless surface. (figure C.2). An actuator applies a force F to the left cart and the positions of the carts can be readout at two locations x and y.

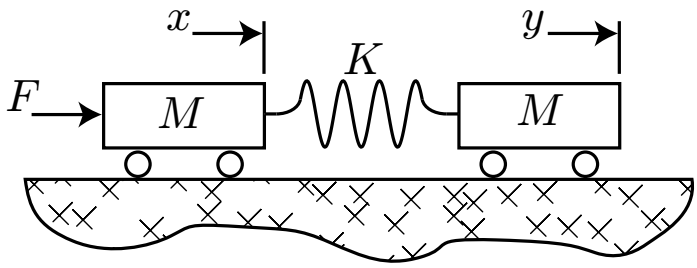

Figure C.2: A model of sensor/actuator collocation.

In the context of this work, the force F represents the hydraulic actuator and the spring stiffness K represents a compliance that may, depending on which output (x or y) is used, separates the sensor and actuator. The purpose of this example is to show that using the collocated readout at x results in a easier to control plant than using the uncollocated readout at y.

The two equations that govern the motion of the system are:

$$M\ddot{x} = F + K(y - x) \tag{C.8}$$

$$M\ddot{y} = K(x - y) \tag{C.9}$$

where the variables are referenced to figure C.2. The transfer functions to each of the sensor outputs are:

$$\frac{x}{F} = \frac{Ms^2 + K}{Ms^2(Ms^2 + 2K)} \tag{C.10}$$

$$\frac{y}{F} = \frac{K}{Ms^2(Ms^2 + 2K)} \tag{C.11}$$

where s is the frequency variable.

In the case where the uncollocated sensor output is used, equation C.11, there are no zeros in the transfer function, and hence, the phase loss does not recover after the pole. The system where the collocated sensor output is used, equation C.10, is therefor much more amenable to feedback control because of the additional phase.

Appendix D

Controls Appendix

D.1 Sensors

This experiment features both displacement sensors and seismometers. The first actuator prototype featured an optical displacement sensor designed and built in the lab. Subsequent versions of the actuator employed the Kaman manufactured DIT5200. Several different seismometers are used in this experiment for feedback, sensor correction and witnessing purposes.

D.1.1 Displacement Sensors

The optical displacement sensor built for this experiment consists primarily of a split photo diode mounted directly across from an LED. The light emitted from the LED is obstructed by a flag with a 1 mm slit. The displacement signal is calculated by taking the difference of the current generated by each side of the photo diode.

The ancillary structure surrounding these components is intended to prohibit ambient light from reaching the photo diode. A micrometer is fastened to the side of the sensor so that the sensor can manually be set to its null position with high precision. Another advantage of the micrometer is that it enables calibration of the sensor in situ.

The Kaman manufactured DIT5200 is a differential (two heads) eddy current

sensor. The advantage of the DIT5200 over the optical sensor is its impressively low noise. However, the DIT5200 is both delicate and expensive. The sensor heads of the DIT5200 are easily damaged on impact, and to reach the advertised performance, both heads must be replaced when one is damaged.

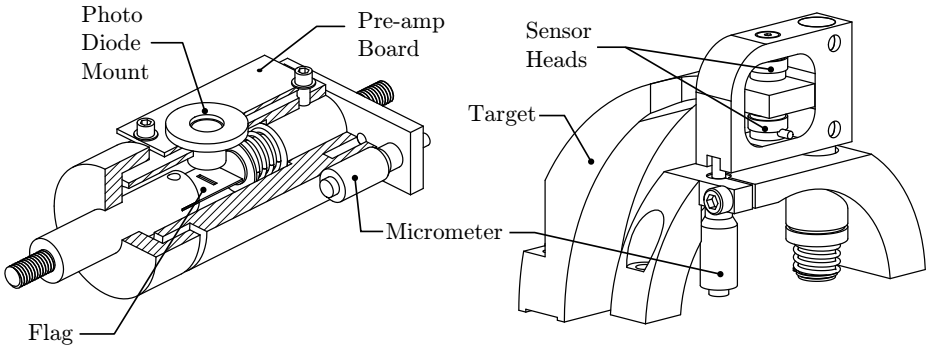

Optical Displacement Sensor DIT5200

Figure D.1: The displacement sensors used in this experiment.

Similar to the optical sensor, the mount for the DIT5200 also incorporates a micrometer head. Arched geometry is used throughout the design of the mount to avoid placing material directly in bending (especially important for a feedback sensor).

D.1.2 Seismometers

Four seismometers are used in this experiment: the Geospace manufactured HS1, the Sercel manufactured L4C, the Streckheisen manufactured STS2 and the Geotech manufactured S13. With exception of the STS2, all of these instruments are passive with one sensitive axis. The sensitivity and noise performance of the passive instruments scales with the size of both the suspended mass (magnet) and coil, hence the larger instruments exhibit better performance.

Alternatively, the STS2 is a 3-axis active device whose output is based on the control effort required to maintain a constant mass position. With an actively controlled suspension, the effective natural frequency may be set to very low frequencies. The natural frequency of the STS2 is 0.008 Hz.

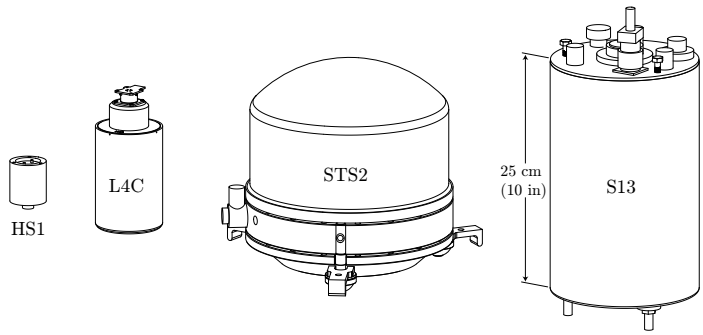

Figure D.2: The seismometers used in this experiment.

The output of a seismometer is derived from differential motion between the case and the suspended mass. At frequencies below the natural frequency, the suspended mass begins to move with the case and the motion ceases to be detectable. This characteristic is evident in the sensitivity curves of the seismometers where at frequencies below the natural frequency, the position sensitivity drops three orders of magnitude per decade.

D.1.3 Sensor Noise

The noise profiles for all of the sensors is shown in figure D.4. The sensor noise profiles for the displacement sensors are measured by rigidly fixing both ends of the displacement sensor to a single piece of material and taking long duration spectral measurements while the sensor is stored in an insulated box.

Alternatively, the noise profiles for the passive seismometers are predicted based on the published noise performance of the readout op-amps and the resistance of the sensor coil. The op-amp for each instrument is chosen to be impedance matched to the coil such that the voltage noise divided by current noise is approximately equal to the coil resistance in the frequency band where the seismometer is to be used. If more than one op-amp match this criteria, the selection may be further refined by choosing the op-amp which has the lowest noise in the band of interest. The three op-amps commonly used in this experiment are the OP-27, LT1001 and the LT1124. Once an

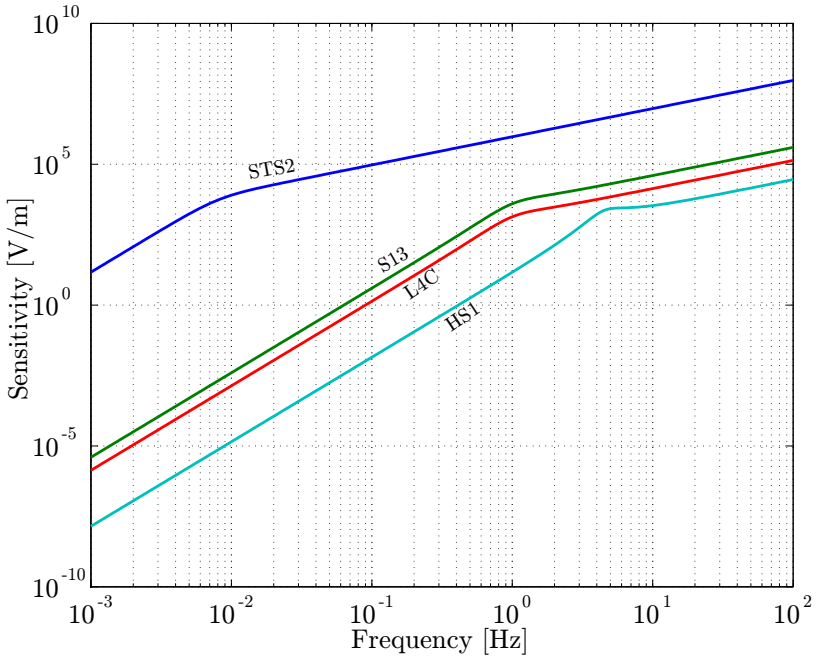

Figure D.3: The position sensitivities for seismometers used in this experiment. Note that seismometer sensitivities are more commonly shown in terms of velocity.

op-amp is selected, the noise profile is calculated by summing the Johnson noise, op-amp voltage noise and op-amp current noise in quadrature. This sum is multiplied by the inverse of the sensitivity plot to give the final noise for the instrument. Note that the predicted noise for the HS1 shown in figure D.4 agrees well with the measured noise of the HS1 shown in figure 6.11.

The noise profile of the STS2 is calculated base on values provided by the manufacturer. Because of the low-frequency and high sensitivity of this instrument, it is difficult to verify these numbers experimentally.

D.2 Tilt-Horizontal Coupling

Tilt-horizontal coupling is a fundamental problem that plagues all horizontally oriented seismometers (inertial sensors). The difficulty arises when it is not possible to distinguish between the effect of accelerations and the Earth's gravitational field. If

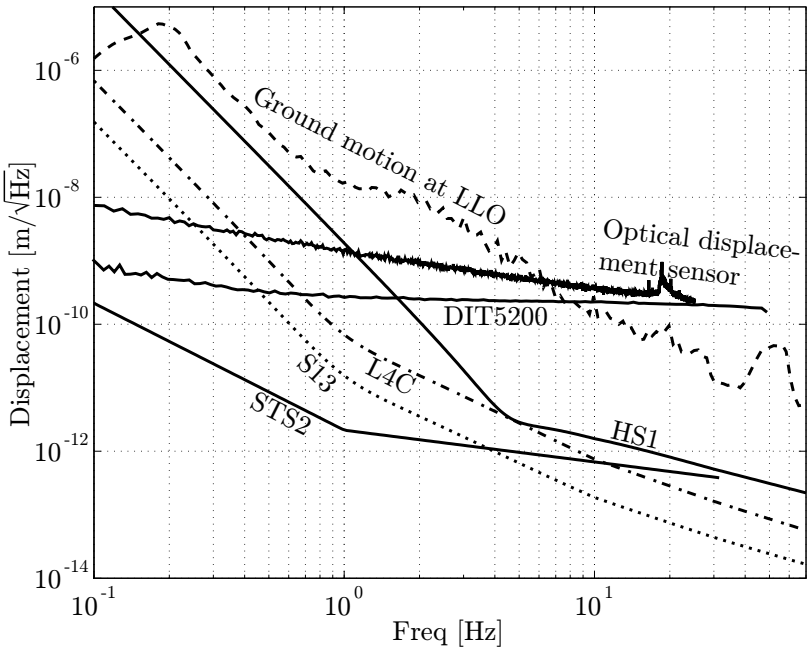

Figure D.4: Sensor noise profiles for sensors used in this experiment.

an inertial sensor is represented by a simple pendulum, figure D.5 demonstrates how an inertial sensor may give the same output for a horizontal acceleration and a tilt.

An inertial sensor detects motion with respect to inertial space by measuring the apparent force acting on its proof mass in the sensor's sensitive direction. When the sensor is operating in the Earth's gravitational field, the apparent force on the sensor is:

$$F = m(\ddot{x} - g\cos(\Theta)) \tag{D.1}$$

where m is the mass of the proof mass, x denotes the sensor's position in inertial space, g is the gravitational acceleration constant on Earth, and Θ denotes the angle of the sensor's sensitive direction with respect to the local vertical, defined by the gravitational field. Hence, the sensor's output is not only a function of the sensor's position, but also a function of the sensor's orientation. The sensitive direction, Θ, can be decomposed as the sum of its nominal direction, Θ_0, and its incremental

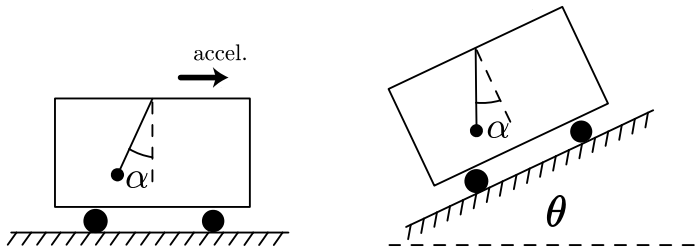

Figure D.5: Tilt-horizontal coupling as described by a simple pendulum in a cart. It is impossible for an observer inside of the cart to distinguish between the effect of horizontal acceleration (a) and tilt (b).[35]

change, θ:

$$\Theta = \Theta_0 + \theta \tag{D.2}$$

This is substituted into equation D.1, and the double angle formula is applied. Then, the incremental change to the apparent force, f, is:

$$f = m(\ddot{x} + g\theta \sin(\Theta_0)) \tag{D.3}$$

Vertically oriented sensors are not coupled to the gravitational force because $\Theta_0 = 0$. However, in the horizontal orientation, $\Theta_0 = \frac{\pi}{2}$ and the apparent force is:

$$f = m(\ddot{x} + g\theta) \tag{D.4}$$

Hence, the motion of the proof mass is a result of both horizontal motion and the local gravitational field. The response function of the sensor in the frequency domain is given by the Laplace transformation of equation D.4:

$$F(s) = \beta(s)(X(s)s^2 + g\theta(s)) \tag{D.5}$$

where the gain and phase relationships of the sensor dynamics and electronics are implicit in $\beta(s)$. Equation D.5 shows that the sensitivity ratio of tilt to horizontal acceleration scales as g/s^2. Hence, tilt measurement dominates the output of a

seismometer at low-frequencies.

Bibliography

[1] R. Hulse and J. Taylor (1975), "Discovery of a pulsar in a binary system," Astrophysical Journal,195,L51-L53.

[2] J. Taylor and J. Weisberg, "Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16," Astrphysical Journal, 345, pp.434-450.

[3] R. Weiss (1972), "Quarterly Progress Report, MIT Research Lab of Electronics", 105, 54.

[4] J R Smith et al (2004) "Commissioning, characterization and operation of the dual-recycled GEO 600", Class. Quantum Grav., 21, S1737-S1745.

[5] K. Tsubono and the TAMA collaboration, "TAMA Project," Proc. of Conference on Gravitational Wave Detection, Tokyo, eds. K. Tsubono, M. K. Fujimotot, K. Kuroda, Universal Academy Press Inc., Tokyo (1997) pp. 183-191.

[6] J. Y. Vinet et al, Proc. of Gravitation and Cosmology, ICGC-95 Conference, Pune, India 13-19 Dec. 1995, Astrophysics and Space Library, Kluwer Academic Publishers 211, (1997) pp. 89-93.

[7] D. G. Blair et al, "Progress in the development of technology for advanced LIGW detectors," Gravitational Wave Detection, Universal Academic Press Inc., 75 (1997).

[8] K. Danzmann, A. Rüdiger, "LISA technology – concept, status, prospects", Class. Quantum Grav., 20(2003), S1-S9.

[9] P Fritschel, D Coyne, J Giaime, B Lantz, D Shoemaker (2001). "Seismic Isolation Subsystem Design Requirements Document", http://www.ligo.caltech.edu/docs/E/E990303-03.pdf

[10] Fritschel P (2001). "Advanced LIGO Systems Design," http://www.ligo.caltech.edu/docs/T/T010075-00.pdf

[11] E J Daw, J A Giaime, D Lormand, M Lubinski and J Zweizig (7 May 2004). "Long-term study of the seismic environment at LIGO,"Class. Quantum Grav., 21 No 9, 2255-2273.

[12] R. Abbott, R. Adhikari, G. Allen, S. Cowley, E. Daw, D. DeBra, J. Giaime, G. Hammond, M. Hammond, C. Hardham, J. How, W. Hua, W. Johnson, B. Lantz, K. Mason, R. Mittleman, J. Nichol, S. Richman, J. Rollins, G. Stapfer, and R. Stebbins (7 April 2002). "Seismic Isolation for Advanced LIGO," Class. Quantum Grav., 19 No 7, 1591-1597.

[13] B. Lantz (1999). "Quantum Limited Optical Phase Detection in a High Power Suspended Interferometer", Phd Thesis, Massachusetts Institute of Technology.

[14] P. Saulson (1994). "Fundamentals of Interferometric Gravitational Wave Dectectors", World Scientific, Salem, MA.

[15] B. Abbot et al (2004). "Limits on gravitational wave emission from selected pulsars using LIGO data",

[16] D.E. McClelland, D.G. Blair, D. Ottaway, et al (1999). "Status of the Australian Consortium for Interferometric Gravitational Astronomy" Amaldi Conference Proceedings

[17] E. Hecht (2002). "Optics" 4^{th} Edition, Addison Wesley, San Francisco, CA

[18] Z. Geng and L. Hayes (1994). "Six Degree of Freedom Active Isolation Using the Stewart Platforms," IEEE Transactions on Control Systems Technology, Vol. 2, No. 1, pp.45-53.

[19] H.-C. Wu (2000). "Ultra-precision carriage positioning of a quiet hydraulics testbed", Stanford PhD Thesis, Stanford, Ca.

[20] D. Thayer and J. Vagners (1998). "Six-Axis Vibration Isolation System using Soft Actuators and Multiple Sensors", 21st Annual AAS Guidance and Control Conference, Breckenridge, CO, February.

[21] A. Slocum (1992). "Precision Machine Design", Society of Manufacturing Engineers.

[22] D. DeBra (1984). "Design of Laminar Flow Restrictors for Damping Pneumatic Vibration Isolators", Presented at the CIRP 34th General Assembly, Paris, France.

[23] C. Steele and Khaliq ur-Rehman Shad (1995). "Asymptotic-numeric solution for shells of revolution", Applied Mechanics Reviews, 48(11, part 2).

[24] Roark and Young (1994). "Roark's Formulas for Stress and Strain", McGraw-Hill, Inc

[25] Clark, R. A. (1950). "On the Theory of Thin Elastic Toroidal Shells", Journal of Mathematics and Physics, vol. 29, no.3

[26] Expansion Joint Manufacturers Association (1993). "Standards of the Expansion Joint Manufacturers Association, Inc.", Tarrytown, New York.

[27] White, Franklin (1994). "Fluid Mechanics", Hightstown, New Jersey.

[28] Oppenheim, Schafer, Buck (1999) "Discrete-Time Signal Processing", Prentice-Hall, Inc., New Jersey.

[29] Gere, Timoshenko (1990). "Mechanics of Materials", PWS Publishing Company, Boston Ma.

[30] Hong, Sae Bae (1999). "Active Vibration Isolation and Alignment Issues for LIGO", Engineers' Thesis, Stanford University.

[31] Hua, Wensheng (2005). "Low-Frequency Vibration Isolation and Alignment System for LIGO", Thesis, Stanford University.

[32] Viersma, Taco (1990) "Analysis, Synthesis and Design of Hydraulic Servosystems and Pipelines", Delft University, Delft, Holland.

[33] Rowe, W. B. (1983) "Hydrostatic and Hybrid Bearing Design", London.

[34] Block, P. (1976) "Das Hydrostatische Keilspaltlager", Thesis, Delft University

[35] W Hua et al (2002) "Low Frequency Active Vibration Isolation for Advanced LIGO", The Internation Society of Optical Engineering.